Powering the Workforce: The Role of Community Colleges in Energy Transformations*

Chloé Tanaka
Harvard University

November 12, 2026

Link to latest version

Abstract

Energy transformations change the skills demanded by local employers. Whether regions can meet those demands depends on the capacity of local training institutions to adapt alongside. I study how U.S. community college programs and graduates respond to energy-induced shifts in local labor demand. I focus on two energy transformations, the shale oil and gas boom of 2000–2014 and the ongoing expansion of solar power. In both, resource endowments determine the location of economic activity, while advances in technology and policy determine the timing. Relative to colleges outside shale plays, shale-exposed colleges see a 30% increase in energy programs (e.g., Petroleum Technology, Well Drilling) and a doubling of energy graduates, while non-energy programs and graduates contract. Over the subsequent decade, the additional energy graduates exceed one third of the estimated boom-induced increase in local middle-skill energy employment. Local solar investment also adds related programs and graduates, but responses are smaller per job created than for shale. At a more granular level, solar course sections nearly double within California community colleges. Turning to student outcomes, in individual microdata from Texas, energy graduates earn 14% more over ten years than observably similar peers in other technical programs. These findings demonstrate how energy transformations shape local training opportunities, human capital accumulation, and career paths.

*chloetanaka@g.harvard.edu. I am deeply grateful to Gordon Hanson, Joseph Aldy, and Joshua Goodman for their invaluable advice and support. I thank Luis Armona, Ed Glaeser, Todd Lensman, Mark Shepard, Clare Suter, Connie Xu, Ruby Zhang, the participants of the Harvard/MIT Environmental Economics and Policy Seminar, Harvard Kennedy School Economics and Social Policy Workshop, Harvard Economics Labor and Public Finance Workshop, Harvard Environmental Economics Workshop, Carnegie Mellon Doctoral Student Participatory Workshop on Climate and Energy Decision Making, and CERAWeek NextGen Cohort for their constructive comments. I gratefully acknowledge the generous financial support from the Harvard University Graduate Prize Fellowship and the Reimagining the Economy project. I would also like to thank the Texas Education Research Center for providing data. All errors are my own. The conclusions of this research do not necessarily reflect the opinion or official position of the Texas Education Research Center, the Texas Education Agency, the Texas Higher Education Coordinating Board, the Texas Workforce Commission, or the State of Texas.

1 Introduction

How do energy endowments shape local economies? Energy booms create jobs and raise income in resource-rich regions (Michaels, 2011; Feyrer et al., 2017; Allcott and Keniston, 2018), altering the education and employment decisions of students and workers (Cascio and Narayan, 2022; Kovalenko, 2023). A substantial literature documents how energy shocks affect local labor markets and individual human capital decisions; however, far less is known about the institutions that shape the supply of skilled labor and whether they adjust when local energy resources become more valuable.

Booms raise both the payoff to training and its opportunity cost. They reshape local economies by reallocating workers into expanding occupations, and individuals can acquire the skills required for these new positions on the job or via training institutions. From 2005 to the peak of the U.S. shale boom in 2014, employment in shale-exposed commuting zones grew 3.4 percentage points more than in unexposed areas[1]. Of that added employment, 14% was high-skill, a bachelor's degree or above; 45% middle-skill, some college but less than a bachelor's degree; and 41% low-skill, a high school degree or less. Middle-skill workers accounted for nearly half of the gain, suggesting that the boom created substantial demand for subbaccalaureate skills. Community colleges are particularly well positioned to supply this workforce. They enroll about 30% of undergraduate students in the United States (NCES, 2025a), and their vocational orientation suggests a natural alignment with local labor markets. Yet, existing research indicates that community colleges expand capacity slowly, if at all, in response to labor demand (Grosz, 2022, 2020).

Since these institutions often serve populations on the margin of college attendance, a boom can work directly against enrollment. Upward pressure on wages increases the opportunity cost of attendance and often reduces educational attainment (Cascio and Narayan, 2022; Kovalenko, 2023; Goodman and Winkelmann, 2025). At the same time, energy transformations also generate large swings in skill demand, creating mismatches between the qualifications workers possess and those employers require. Growing demand for trades workers gives community colleges an incentive to expand their vocational offerings. Thus, an energy boom produces opposing forces on local human capital formation that may give rise to a human capital resource curse, crowding out educational investments (Gylfason, 2001), or

[1]Based on American Community Survey microdata (Ruggles et al., 2025) using Autor and Dorn (2013) and Autor et al. (2019) crosswalks from public use microdata areas to 1990 commuting zones. Shale-exposed commuting zones (115) overlap plays booming between 2005 and 2010 and unexposed commuting zones (490) are outside plays (EIA, 2011). Each education group's contribution is the change in its employment between 2005 and 2014 as a share of total 2005 employment.

bolster human capital accumulation at the middle-skill layer. This paper asks how energy transformations change local labor markets and, in particular, what role these educational institutions play in that adjustment.

I address these questions using two major transformations in U.S. energy production: the 2000–2014 shale oil and gas boom and the ongoing expansion of solar energy. In both settings, preexisting variation in local resource potential determined where energy production could expand, and technological and policy shocks determined when development became profitable. I exploit the interaction between these fixed resource endowments and timing shocks to estimate the causal effects of energy transformations on vocational program offerings and graduates at nearby community colleges. I use institutional data from U.S. community colleges to measure these responses and individual-level administrative records from Texas to estimate earnings returns to energy credentials. Shale provides a long-run context for institutional adaptation to a persistent labor demand shock, while solar captures the early stages of adjustment as the current clean energy transition slowly unfolds.

I first turn to the U.S. shale boom. Beginning in the early 2000s, advances in hydraulic fracturing ("fracking") and horizontal drilling combined with rising commodity prices suddenly made it profitable to recover oil and gas that had been trapped in shale formations for millions of years. This resulted in a large, abrupt, and plausibly exogenous surge in demand for skilled trades workers in energy production. Shale plays are the geographic areas where these formations exist, and the Energy Information Administration (EIA) provides data on their boundaries. To assess the impact of the boom on educational institutions, I compare colleges inside booming plays to similar colleges in the same state but outside the play.

Exploiting staggered boom timing across plays, driven by geological variation in ease of extraction, I employ stacked difference-in-differences and event-study designs to estimate the causal effect of the boom on vocational program offerings and graduates at local community colleges. I leverage the National Center for Education Statistics' (NCES) Integrated Postsecondary Education Data System (IPEDS) to measure these outcomes. The data comprise annual, institution-level surveys on programs offered and graduates by program at the universe of public two-year colleges from 1990 to 2024. I map programs to related occupations and industries to analyze sectoral trends. My sample covers 391 colleges in 19 states, with 100 shale-exposed colleges across 17 plays that boomed in 2001–2012. The median community college adds 28 vocational programs per decade, which train 24% of graduates.

I find that the boom increases the number of energy programs by 30% of the baseline mean (0.16 programs, 0.33 SD) at community colleges inside a shale play compared to

colleges outside the play. Energy graduates also increase significantly. Shale-exposed colleges more than double their energy credentials awarded each year, adding 5 more graduates annually (0.66 SD) than their non-shale-exposed counterparts. Across the 100 shale-exposed colleges in my sample, this amounts to about 16 additional energy programs offered in the average post-boom year and 550 additional energy graduates per post-boom year. Graduate completions in preexisting energy programs show no signs of expansion, consistent with growth concentrated in new programs–an indication that vocational offerings may be an important margin of adjustment. In line with the existing literature, responses are slow, with effects beginning about five years after onset and continuing for more than a decade. The effects remain large through the end of the event window, suggesting that this is not simply short-term reallocation, but instead reflects a sustained increase in institutional capacity to produce energy-aligned graduates. Community colleges supplied enough graduates to account for over one third of the 15,300 middle-skill energy workers added by the boom in the decade following onset. However, the expansion arrives slowly, with graduates covering only about 5% of new middle-skill jobs at the employment peak. Adjusting two years sooner or as strongly as the average top-quartile college would each roughly triple that share.

The expansion on the energy side is accompanied by a relative decline in non-energy programs and graduates, as colleges outside plays grow their non-energy offerings more quickly than those of shale-exposed colleges. The boom resulted in 10% smaller non-energy catalogs (3 programs, 0.26 SD) on average at shale-exposed colleges compared to their non-exposed counterparts. Again, the graduate response is even stronger, with 31% fewer graduates (105 graduates, 0.43 SD) in non-energy programs per year, likely reflecting the rising opportunity cost of enrollment. In absolute terms, the magnitude of the effect for non-energy is materially larger than for energy programs and graduates. This is unsurprising given the comparative sizes of the groups–there are far more non-energy than energy programs at baseline. Event study results also reveal an asymmetry in timing, with the relative reduction in non-energy programs and graduates beginning almost immediately after the boom. I do not find any significant change in total spending, instructional spending, or tuition per full-time-equivalent student, which is consistent with a reallocation of existing resources toward energy.

Community colleges are budget-constrained institutions that vary substantially in measured quality (Kurlaender et al., 2016) and career and technical orientation, which is correlated with their ability to produce high-wage credentials (Holzer et al., 2026). Their financial resources, administrative and instructional capacity, existing program catalogs, and regional labor market conditions are measurably different. Although colleges in the same shale play may face similar demand, their capacities to create and sustain new programs are less ho-

mogeneous. I find that some shale-exposed colleges more than double their energy offerings while others, equally exposed, do not adjust. I provide descriptive evidence on the correlates of this heterogeneity in response, exploring four key metrics of capability that correspond to three inputs to the community college production function: spending, administrative quality, and instruction. Colleges with higher administrative pay expand energy programs more, while those with higher faculty pay expand less, and neither of the other two metrics, instructional spending per student and prior NSF grants, predicts the response. Unlike faculty compensation, the administrative pay relationship does not appear to operate through differential revenue growth, suggesting a role for a less tangible factor such as quality.

I then turn to the ongoing expansion of solar energy and provide evidence on an institutional response that is still unfolding. Starting in the 2000s, falling solar module costs and growing national and state policy incentives made solar development increasingly profitable, particularly in areas with abundant solar resources. As in the shale setting, this transformation combines predetermined geographic resource variation with subsequent shocks that altered the value of the resource. Following Chan and Zhou (2025), I use a Bartik-style exposure design, interacting local solar potential with changes in policy incentives and module prices, to instrument for local solar investment across distributed and utility-scale solar. I focus on the 2002–2024 period during which solar investment expanded substantially, using the same IPEDS data to measure community college responsiveness. My sample includes 1,165 colleges across 47 states and 478 commuting zones.

As with shale, colleges expand solar training where capacity grows. A one-gigawatt increase in solar investment raises solar programs at the average college by 10% (0.02 programs, 0.08 SD) and graduates by 21% (0.14 graduates, 0.05 SD) annually, relative to their sample means. The average college's commuting zone received 0.7 gigawatts in solar investment over the period. Summed over the 1,165 community colleges in my sample, that implies an average of 4 additional solar programs on offer and 29 additional solar graduates per year. A back-of-the-envelope calculation suggests that the expansion-induced increase in solar graduates amounts to less than 1% of the roughly 7,000 solar jobs added each year in sample commuting zones. Notably, program and graduate responses are qualitatively similar but somewhat smaller than for shale. For programs, the effects in standard deviations cannot be distinguished at the 5% level, while the graduate response remains more modest for solar. One potential explanation is that labor demand for solar is more front-loaded and transitory. Nearly two thirds of solar employment is in installation and development (IREC, 2024), which primarily occur prior to operation. By the time a college creates a program and students complete it, local hiring may have subsided, weakening the incentives of colleges

and students to invest in new solar-specific credentials.

If colleges can modify instruction within related programs at a lower cost and on a shorter timeline than creating new ones, programs and graduates would understate the response. To assess this more granular margin, I supplement the national analysis with a case study of California's community colleges, using the California Community Colleges Chancellor's Office (CCCCO) course-level data. These administrative data cover all 116 community colleges from 1993 to 2025. My sample includes 113 of these colleges across 38 counties. I find that a one-gigawatt increase in solar investment raises solar course sections by 93% (0.40 sections, 0.59 SD), indicating that this may be a more accessible margin of response for community colleges. The average college's county added 1.5 gigawatts of solar capacity between 2002 and 2024. Across the 113 community colleges, the average shock-induced increase is 19 solar sections on offer each year.

The solar results demonstrate that the institutional responsiveness observed in the shale context extends to the clean energy transition; however, the scale and form of the response may depend on the nature of the underlying labor demand shock. The two shocks differ in size and duration, so I compare responses per 1,000 new jobs across all skill levels[2]. The shale boom added about 1,455 total energy jobs to the average exposed commuting zone, against an implied 1,270 new solar jobs. Per 1,000 new jobs, shale produced about 7 graduates annually and 0.2 programs, compared with solar's 0.7 graduates annually and 0.1 programs across the U.S., and 1.2 course sections in California. The new energy programs, such as Petroleum Technology and Heavy Equipment Operation, are distinct from the solar programs, Solar Technician and Energy Systems Technician. Therefore, community colleges exhibit general dexterity across various classes of training. Still, graduate responses are stronger to shale, suggesting that other factors may be affecting adjustment to solar. As AI and electrification increase electricity demand, these differences raise questions about whether community colleges can better contribute to the expansion of generation capacity and the clean energy transition. Understanding when local training constrains energy deployment and how colleges might respond more quickly is an important direction for further research.

A natural question is how this response affects the future earnings of energy graduates. Energy jobs tend to be high paying, but given the boom-and-bust nature of the sector, the future demand for the skills acquired in energy programs is uncertain (Katovich et al., 2026). I use linked administrative microdata from the Texas Education Research Center (ERC), which follows Texas students from elementary school into the labor market, to descriptively

[2]Middle-skill jobs cannot be estimated for solar in the way that they can for shale because solar jobs are not separately identified in local employment data.

assess the returns to energy credentials and investigate the impact of the shale boom on earnings for graduates of community colleges. In line with Abboud and Betz (2021), I define the Texas shale boom as 2005–2014 and the first bust as 2015–2017, when oil prices fell by more than 70%. The second bust is during the Covid-19 pandemic of 2020. One obvious concern is that any observed premium reflects worker characteristics that would have been rewarded in the labor market absent an energy credential. To limit this, I control for pre-college test scores and compare energy graduates to those of the closest alternative programs; however, I still interpret the estimates as descriptive. Conditioning on a broad set of observables, I find that energy graduates earn a 14% premium on the net present value (NPV) of 10-year earnings relative to graduates of other engineering and technician programs, which remains stable at the 15-year horizon. In fact, relative to all but one non-energy field, energy is associated with an earnings premium. Estimates for each cohort reveal that cross-cohort differences track general labor market conditions at entry more closely than the shale cycle. They also do not appear to reflect changes in selection into energy fields. Together, the results suggest that community colleges responded to the shale boom by expanding energy programs that may offer lasting earnings benefits for graduates.

This paper relates to three strands of the literature. The first is on the resource curse and, more specifically, its human capital dimension. A large cross-country literature documents a negative relationship between natural resource dependence and economic growth (Sachs and Warner, 2001; Van der Ploeg, 2011), related in part to reduced educational investment (Gylfason, 2001). Studies of modern U.S. energy booms, however, fail to find the classic symptoms, documenting higher income and population growth as well as non-mining employment gains (Michaels, 2011; Weber, 2014; Feyrer et al., 2017; Allcott and Keniston, 2018). The effects on human capital are similarly mixed. While Maniloff and Mastromonaco (2014) find that shale development dampened growth in the share of college-educated adults, Weber (2014) reports no decline in local educational attainment. Yet, much evidence shows that fracking raised high school dropout rates and lowered university enrollment among young men by increasing the opportunity cost of schooling, suggesting that the human capital costs of resource booms may operate on margins not captured by aggregate measures (Morissette et al., 2015; Zuo et al., 2019; Cascio and Narayan, 2022; Kovalenko, 2023). The Texas boom reduced school test scores and student attendance even as it tripled local tax bases (Marchand and Weber, 2020; Kovalenko, 2023). My findings speak to both sides of this evidence, focusing on an understudied margin: adjustment of middle-skill training institutions. Existing work predominantly asks how energy booms change high school students' incentives to stay in school, taking the set of educational options as given. I examine middle-

skill training, which has received little attention despite the importance of these skills for the oil and gas sector. And, I show that both the students and the institutions supplying training respond, changing the type of human capital the boom produces as well as the amount. Energy booms do not simply reduce educational investment. They shift it toward skills demanded by the booming sector, which may help explain why these booms do not produce the broad-based deskilling predicted by the resource curse.

Second, this paper contributes to work on local labor market adjustment to structural economic shocks and the role of human capital in that process. A substantial literature examines responses to negative shocks. Labor markets recover slowly (Autor et al., 2016) and reallocation largely comes through demographic turnover rather than the reskilling of existing workers (Autor et al., 2025; Krause, 2025a,b). Yet, when workers do retrain, there are large returns to community college career and technical programs (Jacobson et al., 2005), particularly for workers switching industries (Leung and Pei, 2020). Shocks affect individual human capital investment through two channels: changing the opportunity cost of education and altering the skills demanded by local employers. Negative shocks align these forces, pushing individuals into schooling, as both wages and demand for existing skills fall (Betts and McFarland, 1995; Goodman and Winkelmann, 2025). Community college enrollment rises during downturns and following mass layoffs (Betts and McFarland, 1995; Foote and Grosz, 2020; Goodman and Winkelmann, 2025), and students shift away from declining fields (Acton, 2021; Weinstein, 2022) and toward high-return ones (Foote and Grosz, 2020). Positive shocks put these forces in opposition. The net effect is ambiguous, as wages rise while shifting skill demand increases the return to training. Research on gradual changes in occupational demand reveals the skill-building channel, as students reallocate toward high-return programs across community colleges (Grosz, 2022) and four-year institutions (Conzelmann et al., 2023). On the extensive margin, community college enrollment declines in strong labor markets (Goodman and Winkelmann, 2025), and Atkin (2016) shows that the expansion of low-skill export manufacturing in Mexico similarly raises high school dropout rates. A closely related paper examining a large, positive labor demand shock to the two-year market is Charles et al. (2018). They find that the 2000s housing boom reduced enrollment, as young adults took low-skill jobs in related sectors. However, the extent to which a housing boom raises demand for middle-skill workers is unclear, leaving an open question about how educational investment responds when these opportunities expand.

Students follow local labor market signals. In contrast, existing work suggests that institutional supply is far more rigid, though I find surprising dexterity. Grosz (2022) shows that community college course sections and faculty do not appear to adjust to shifting em-

ployment patterns, with the alignment in completions driven primarily by student demand. Supply constraints limit access to high-return fields, such as nursing (Grosz, 2020), and shape responsiveness in four-year degrees (Conzelmann et al., 2023). I make two contributions to this literature. First, I examine two large, positive labor demand shocks and demonstrate that community colleges respond by creating programs and reorienting their offerings toward the booming sector. Second, I show that these supply-side responses extend to graduates, altering the composition of vocationally trained workers.

Finally, this paper extends a growing literature on how regions respond to the clean energy transition and which are positioned to gain or lose. Recent work characterizes the geography and composition of green energy investment and employment. Wind and solar resources combined with inexpensive land create opportunities for renewable expansion (Huang and Kahn, 2024), which raises local employment (Chan and Zhou, 2025). Green jobs appear to pay a premium (Vona et al., 2019; Curtis and Marinescu, 2023) and generate multiplier effects in other local sectors (Vona et al., 2019; Chan and Zhou, 2025). Declining solar costs are projected to benefit places with stronger solar resources and more electricity-intensive industries (Arkolakis and Walsh, 2024). Evidence on fossil fuel communities is more mixed (Curtis and Marinescu, 2023; Hanson, 2023; Lim et al., 2023). Moves from carbon-intensive to green jobs are rare (Curtis et al., 2024), reflecting substantial frictions in the reallocation of workers across energy sectors, including locational mismatch (Hanson, 2023; Lim et al., 2023) and possible skill and training gaps (Popp et al., 2021). Popp et al. (2021) find that a region's ability to benefit from the transition may depend on the skills of its existing workforce; yet, less is known about the role of local training in building these skills. Descriptive work suggests that related programs remain small (Greenspon and Raimi, 2025). I provide evidence on the complementary role of worker training in helping resource-rich regions capitalize on this transition.[3] The solar analysis demonstrates how community colleges adjust programs and course offerings and train more graduates as renewable capacity expands. These responses can help build the workforce needed to support the energy transition.

The remainder of the paper proceeds as follows. Sections 2 and 3 describe the setting and the data and outcomes. Section 4 reports the effects of the shale boom on community college offerings and graduates, the margin of adjustment, and heterogeneity in the response. Section 5 turns to the clean energy transition and presents the results from the national and California solar build-outs. Section 6 explores earnings effects for energy fields. Section 7 draws out the magnitude of the responses and policy implications, and Section 8 concludes.

[3]How local training institutions support regions adversely impacted by the energy transition remains an open question for future research.

2 Background

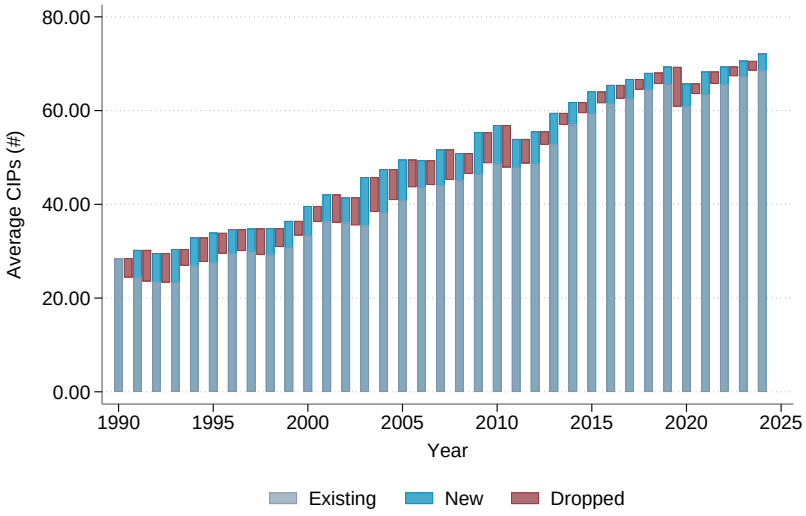
2.1 Community Colleges

Between 1990 and 2024, more than 1,300 community colleges have operated across the United States, accounting for approximately 30% of total Fall enrollment in institutions of higher education (NCES, 2025a). These open-access, locally governed two-year public institutions have evolved over the twentieth century, shifting from primarily transfer-oriented missions toward career and technical education driven by funding mandates to meet the needs of the local community (Kane and Rouse, 1999; Cohen and Brawer, 2003). Beyond this, two features in particular suggest community colleges should be closely tied to their local labor markets (Grosz, 2022; Cohen and Brawer, 2003). First, they serve less-mobile and lower-income students who often enroll close to home and enter the local labor force directly upon completion (Acton, 2021; Hillman and Weichman, 2016; Sentz et al., 2018), meaning that both the demand they face and the workers they supply are largely regional. Second, they draw 50–60% of their revenue from state and local governments (NCES, 2024), potentially tying their incentives to regional industrial and policy conditions rather than the national market.

Although the composition of most funding sources is stable over time for the average institution, cyclical shifts in federal and state support occur (Barr and Turner, 2013; NCES, 2024). This stability also masks declines in resources per student when appropriations fail to keep pace with enrollment growth (Bound, 2007; Bound et al., 2010). Levels can fluctuate from year to year and there is substantial heterogeneity across institutions and regions. While some colleges rely on state and local funding for more than 90% of their revenue, others receive less than 20% (NCES, 2024). Panel A in Figure A1 shows average funding composition over 2004–2019 and Panel B shows the 2019 distribution across institutions. Variation in their allocation mechanisms adds an additional layer of complexity. A majority of states (37) rely on some type of funding formula that accounts for factors including enrollment and program costs. The remainder receive funding based on prior year levels (6) or institutional budget request (7) (InformEd States, n.d.). Thus, two community colleges facing identical demand shocks may have very different marginal incentives and constraints.

These features give community colleges a possible motivation to respond to local demand, but the question of whether their program structure gives them the means remains. CIP codes are standardized categories of fields of study used to report information on programs. They typically span several related majors. Figure 1 documents program churn,

Figure 1: Average Annual Program Stock and Flow at Community Colleges



Notes. This figure shows the average number of programs (CIP codes) across award levels (associate's degrees, one-to-two-year certificates, and certificates of less than one year) offered at community colleges from 1990-2024 using NCES IPEDS completions data. New programs are programs that are reported in the current year, but not in the previous year. Dropped programs are those that are not reported in the current year, but were reported in the previous year. New and dropped program definitions exclude 1990 and 2024. The sample is restricted to institutions that remained in the data throughout the period.

plotting the average number of existing, new, and dropped CIP codes per institution over time across certificates of less than one year, one-to-two-year certificates, and associate's degrees. On average, community colleges add about 6 CIPs per year and drop around 4.5. In a typical year, 1.4 programs are added on net with some years of outsized program creation and discontinuation seemingly tied to funding cycles.

Although Figure 1 illustrates that program turnover is a routine feature of the community college landscape, it also makes clear that the bulk of the existing program offering is maintained from year to year, with more than two thirds of programs carried over on average. Churn, when a college drops a program and later offers it again, is also uneven, with certain programs and award levels turning over much more frequently than others. Table 1 reports the top 10 highest and lowest estimated churn[4] programs across award levels. Associate degrees are the least likely to churn, about once every 33 years a college offers them, compared with about once every 20 years for certificates of both levels. Although I do not directly observe start-up costs, the pattern is consistent with lower churn where programs are most expensive to establish. The highest churn programs are largely classroom-based business and

[4]Churn rates for programs offered in few program-years are noisy, so I report empirical Bayes estimates that shrink each program's observed rate toward the average for other programs in the same two-digit CIP family and award level. Table 1 also reports observed churn.

Table 1: Program Churn at Community Colleges

(a) Highest Churn

Award Level	CIP Code	CIP Title	Churn per 10 Program-Years Est.	Obs.
1–2 yr	52.1803	Retail Operations	1.28	1.39
AA	50.0706	Multimedia	1.19	1.36
1–2 yr	52.1902	Fashion Merchandising	1.14	1.31
1–2 yr	52.1501	Real Estate	1.13	1.20
<1 yr	01.0000	Agriculture	1.11	1.64
1–2 yr	50.0301	Dance	1.10	2.02
<1 yr	50.0501	Drama/Theatre	1.08	1.32
<1 yr	09.0401	Journalism	1.07	1.26
1–2 yr	52.0803	Financial Support Svcs	1.06	1.13
1–2 yr	09.0401	Journalism	1.05	1.20

(b) Lowest Churn

Award Level	CIP Code	CIP Title	Churn per 10 Program-Years Est.	Obs.
AA	51.0917	Polysomnography	0.05	0.00
AA	51.3801	Registered Nursing	0.05	0.05
AA	51.0001	Health and Wellness	0.06	0.00
AA	51.1503	Clinical Social Work	0.06	0.00
AA	51.0602	Dental Hygiene	0.07	0.06
AA	51.0711	Medical Assistant	0.08	0.00
AA	51.0911	Radiologic Technology	0.08	0.08
AA	01.8301	Veterinary Technology	0.08	0.07
AA	51.0810	EMT	0.09	0.00
AA	51.0920	MRI Technology	0.09	0.00

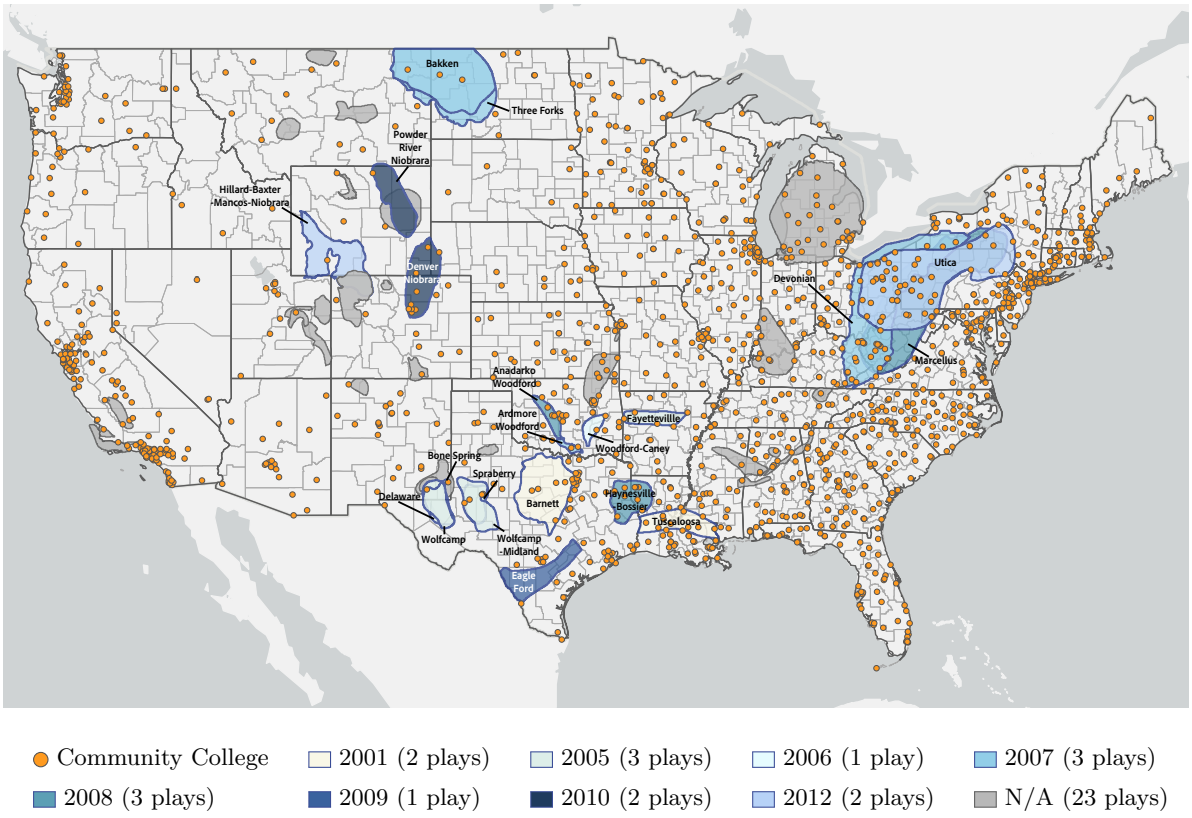
Notes. This table denotes the ten programs with the highest (panel a) and lowest (panel b) estimated churn. Churn is defined as programs that are dropped and added at the same college to exclude genuine new program creation and discontinuation. Estimated churn uses empirical Bayes shrinkage, pulling each program's observed rate toward the average for other programs in the same two-digit CIP family and award level and placing more weight on the observed rate for programs with more program-years. Observed churn is the number of churn events per 10 program-years, where a program-year is a year in which a college reports the program. Across all programs, churn averages 0.30 per 10 program-years for associate's degrees, 0.52 for one-to-two-year certificates, and 0.50 for certificates of less than one year.

arts fields, while the lowest-churn ones are health associate's degrees that require specialized equipment, clinical placements, and accreditation. The median energy program sits near the middle of the churn distribution, with energy associate's degrees exhibiting higher churn (75th percentile) and one-to-two-year and less-than-one-year certificates churning less (41st and 37th percentiles, respectively).

2.2 The Shale Boom

The shale boom is a good setting to study due to the unexpected nature of its takeoff combined with the sustained labor required to recover shale oil and gas. Shale is a dense (or "tight") sedimentary rock that is rich in organic materials. Unlike in traditional oil fields, shale oil and gas live in small deposits between the dense rock such that they cannot be extracted with traditional drilling techniques. The location of shale formations is often defined by the plays and basins in which they sit, and in this paper, shale plays will be used to define treatment. A play is an underground, geographic area containing a set of shale formations that are geologically similar. And, a basin is a large depression in the earth's surface in which a number of shale plays with similar attributes, for example, hydrocarbon intensity and age, sit. Both the location and timing of extraction make the shale boom exogenous to community college location and program choice.

Figure 2: U.S. Shale Plays and Community Colleges



Notes. This figure depicts U.S. shale play boundaries using shapefiles from the EIA. Community colleges from NCES IPEDS data have been overlaid in orange. Play shading indicates treatment year. Four labeled plays overlap a sample play in the same basin and are combined with it as one play for treatment and in the legend counts. Wolfcamp and Bone Spring overlap Delaware, Spraberry overlaps Wolfcamp-Midland, and Three Forks overlaps Bakken. Plays in gray are not included in my sample either due to insufficient treatment colleges, never reaching the treatment threshold, or confounding effects from the financial crisis (Michigan Antrim). Commuting zone boundaries are shown in light gray. Basemap sourced from ESRI.

The US has the second most technically recoverable shale oil and the fourth most gas in the world (EIA, 2013), but it was not economically feasible to extract prior to the early 2000s. Only after innovations in hydraulic fracturing[5] and horizontal drilling[6] techniques, could the shale boom take off. These advancements combined with favorable economic conditions, including rising oil prices, suddenly unlocked vast reserves of these organic materials. Since the geology of a shale play determines the costs and benefits of production based on the density of the rock and the quantity of the reserves to be recovered, different plays will experience their "booms" at different times, giving rise to a staggered treatment design.

[5]Fracking involves pumping water and sand down through wells and into the shale rock to blast open pockets of hydrocarbons.

[6]In horizontal drilling, a drill is inserted down into the earth that then makes a 90 degree turn to access pockets of hydrocarbons.

Geographically, the spatial distribution of shale plays, including their boundaries and thickness, is fixed by ancient geology. Figure 2 shows that there does not appear to be any correlation between the location of community colleges and shale plays themselves. Furthermore, knowing the boundary of these plays is highly unlikely given that geospatial sensing techniques are required to determine the location and characteristics of these plays.

Shale extraction is also considerably more labor intensive than traditional oil and gas due to the nature of its production cycle. Shale oil is characterized by rapid initial production followed by a steep decline related to the geology of the formations, necessitating constant drilling. This translates into persistent, high-volume demand for skilled trades workers, for example, drilling crews, truckers, welders, and construction workers. The labor intensity of the industry means these booms are large, long-lasting shocks to local labor demand.

Approximately 4% of workers in the mining, oil and gas sector (NAICS 21) hold an occupational associate's degree, and an additional 18% have obtained some college but no degree (Flood et al., 2025)[7]. This suggests that a material share of the industry (around 22%) enters through subbaccalaureate pathways for which community colleges are a natural vocational training provider. Throughout, I refer to these workers, with some college but less than a bachelor's degree, as middle-skill. Examples include certificates and associate degrees in oil and gas production technology, natural gas compression, and petroleum field service. These programs prepare students to work as production operators, compressor technicians, and field-service technicians.

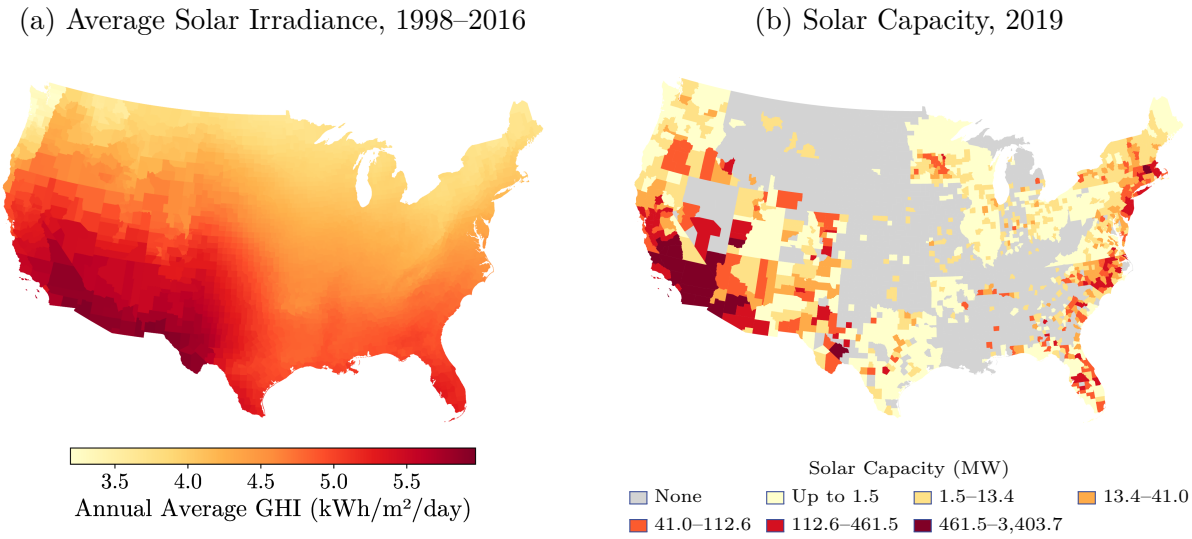
Community colleges may also serve the roughly 45% of workers in the sector without any postsecondary education through short, non-credit credentialing. Entry-level workers typically receive industry safety certificates, commercial driver's licenses, or other occupation-specific training. Employers may provide this on the job, through private vendors, or by partnering with local community colleges. However, much of this activity cannot be observed in my data because non-credit credentialing does not lead to formal awards.

2.3 The Solar Expansion

The roll-out of solar power is a second and ongoing energy transformation. Like shale, the location of a predetermined resource, in this case sunshine, shapes the location of development, and technology and policy shocks determine when. The productivity of solar power depends on "irradiance", or the amount of sunlight that reaches the ground in a given area.

[7]Based on the average share of NAICS 21 employment with an occupational associate's degree or some college, no degree in the 2010–2019 CPS.

Figure 3: Solar Irradiance and Installed Solar Capacity by County



Notes. Panel (a) displays average annual global horizontal irradiance (GHI) by county over 1998–2016 from the NREL National Solar Radiation Database (Sengupta et al., 2018). Panel (b) displays utility-scale and distributed solar capacity by county in 2019 from EIA Form 860 (EIA, 2025b) and Lawrence Berkeley National Laboratory's (LBNL) Tracking the Sun (LBNL, 2025b).

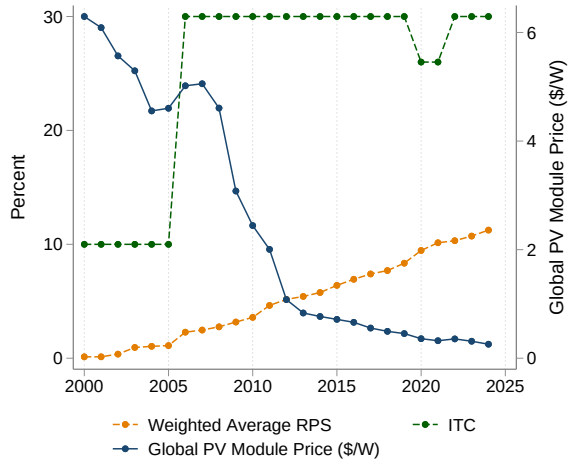
Figure 3a shows that irradiance is highest in the southern U.S., particularly the Southwest, and generally declines moving north and east. Installed capacity follows irradiance in the Southwest, but is also high in the Northeast and the Carolinas, where state policy rather than sunshine drove investment (Figure 3b).

Through the mid-2000s, solar power was too expensive to compete with other sources of electricity. However, over the subsequent two decades, two policy changes coincided with a dramatic decline in technology costs to close that gap (see Figure 4). First, state Renewable Portfolio Standards (RPS), which require suppliers to obtain a specified share of their electricity from renewable sources, expanded materially. In 2005, only 12 states had RPS obligations in effect; by 2024 that number rose to 29 states. Throughout the same period, the renewable electricity required under those obligations rose from about 1% to over 11% of total electricity sales. Second, the Energy Policy Act of 2005 increased the federal investment tax credit (ITC) for solar projects from 10% to 30% for systems starting service in 2006, and has since remained near that level. At the same time, the global price of solar panels fell more than 90% in real terms from roughly $4.60 per watt in 2005 to about $0.26 in 2024 (Our World in Data, 2026). Against this backdrop, total installed U.S. solar capacity expanded from less than half a gigawatt in 2000 to nearly 155 gigawatts in 2024 (EIA, 2025b; LBNL, 2025b).

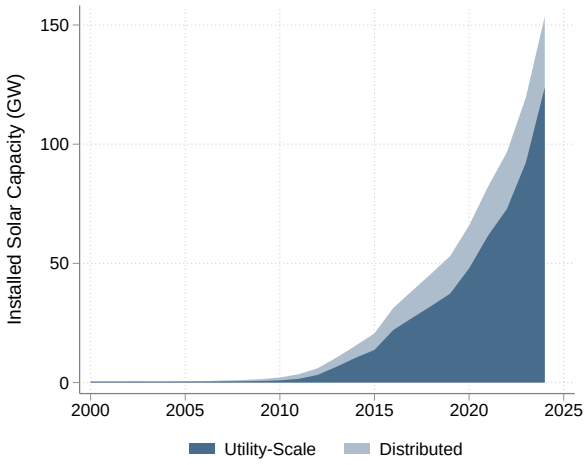
Yet, unlike shale, the labor demand generated by solar power is smaller and less per-

Figure 4: Solar Cost and Policy Shifters and Installed Solar Capacity

(a) Cost and Policy Shifters



(b) Installed Solar Capacity



Notes. Panel (a) displays national Renewable Portfolio Standard obligations as a percent of retail electricity sales (LBNL, 2025a), the federal investment tax credit (ITC) rate, and the global price of solar photovoltaic modules in inflation-adjusted dollars per watt (Our World in Data, 2026). Panel (b) displays installed solar capacity, combining utility-scale capacity from EIA Form 860 and distributed capacity from LBNL's Tracking the Sun (EIA, 2025b; LBNL, 2025b).

sistent. Nearly two thirds of solar employment is in installation and project development compared to less than one-tenth in operations and maintenance (IREC, 2024). For any given project, employment peaks during construction. Without continued investment in new projects, much of the local labor demand declines once a facility begins operating. This transience is especially pronounced for distributed solar where individual installations typically take days rather than months or years. National employment and capacity growth between 2010 and 2022 implies that on average solar generates roughly 1,800 jobs for each additional gigawatt of capacity (IREC, 2024). A gigawatt is enough to supply the average annual electricity consumption of almost 200,000 homes[8]. The average community college's commuting zone received about 0.7 gigawatts of solar investment over the sample period, equivalent to approximately 1,270 jobs, supporting less employment than the 1,455 jobs the shale boom added to the average exposed commuting zone (Section 4.2).

The solar workforce includes a number of occupations for which community colleges provide training, including electricians, solar photovoltaic installers, and operations and maintenance technicians. Relevant programs include certificates and associate's degrees in electrical and renewable energy technology, as well as courses in solar installation, system

[8]Based on the EIA's estimate of annual average electricity consumption at 10,500 kWh per household and a utility-scale solar capacity factor of 23%.

design, and maintenance. These credentials prepare workers to install panels and electrical connections, test system performance, and troubleshoot equipment. Entry requirements vary across occupations. PV installers typically enter with a high school diploma and receive on-the-job training, sometimes supplemented by community college coursework (BLS, 2026c). Electricians typically combine classroom instruction with supervised work experience and state or local licensing requirements. Community colleges, employers, and unions all provide this instruction (BLS, 2026b). Associate's degrees or apprenticeships are relevant pathways into solar operations and maintenance (IREC, n.d.).

These arrangements may generate a different institutional response than shale. Solar-specific instruction can build on established electrical and construction skills (BLS, 2026c). As a result, colleges may accommodate demand by adding courses or specializations within existing programs. Training may also expand through employer- or union-sponsored apprenticeships, with colleges supplying related classroom instruction.[9] Combined with the smaller and less persistent employment shock described above, these features suggest that colleges may respond less strongly to solar than to shale. They may also adjust within existing programs or partnerships with employers and unions rather than create new programs.

3 Data

In this section I describe the main data sources used in this analysis, their summary statistics and the construction of the key outcome variables.

3.1 Data Sources

US Community Colleges. I obtain data on 1,324 community colleges operating between 1990 and 2024 in the US from the NCES IPEDS (NCES, 2025b). All schools receiving Title IV funding are required to participate in annual IPEDS surveys, providing data on nearly all public postsecondary institutions as well as many private ones. IPEDS includes information on the institution's characteristics, such as size, location, ownership, and population served. The shale shock analysis focuses on 391 community colleges across 19 "shale states", and the solar analysis covers 1,165 colleges across 47 states and 478 commuting zones.

I measure community college responsiveness as two broad categories of outcomes: breadth

[9] Union-represented workers accounted for 13.4% of solar employment in 2023 (IREC, 2024). The Inflation Reduction Act also linked enhanced tax credits to registered apprenticeship participation for covered projects. For projects subject to these requirements and beginning construction in 2024, registered apprentices were generally required to perform at least 15% of construction, alteration, or repair labor hours (IRS, 2024).

of vocational offerings and graduation counts by program, for public two-year colleges. Previous research similarly relies on the IPEDS data to study the responses of postsecondary enrollment and degree completions to changing labor market conditions (Long et al., 2015; Foote and Grosz, 2020; Conzelmann et al., 2023). The year here refers to the end of the academic year. Each institution is required to report the number of completions ("graduations") by six-digit CIP code ("program") for all active programs including those with no graduates in a given year. My data includes 1,275 unique six-digit CIP codes. I supplement this data with IPEDS institutional characteristics, including funding, enrollment, staffing, and expenses, which have commonly been used in existing literature (see Bound et al., 2010; Barr and Turner, 2013; Deming and Walters, 2017).

Shale Activity. In order to assign community colleges to treatment and control groups, I use shale play boundary shapefiles from the EIA (EIA, 2011). These files have been used in previous analyses of the shale boom to determine exposure to fracking (Bartik et al., 2019; Cascio and Narayan, 2022). To assess the onset of a shale boom, I rely on two data sources. The first is Bartik et al. (2019)'s list of the date a play's fracking potential "became public knowledge". Second, I leverage proprietary, well-level microdata from Enverus DrillingInfo (Enverus, 2025), which covers the near-universe of U.S. oil and gas wells including information on well location, type, production, and date completed. Horizontal oil and gas wells for production are geographically matched to shale plays to characterize drilling activity similar to Feyrer et al. (2017) and Bartik et al. (2019). I use the date each play reaches 10 completed wells as the boom date for plays that are not available in Bartik et al. (2019)'s data as well as in robustness specifications. Of the 44 US shale plays in the EIA data, 28 experience a "boom" during my sample period (1990–2024) of which 17 are included in my analysis[10].

Solar Investment. Extending the analysis to renewable energy shocks, I construct measures of local solar potential from remote-sensing data in the National Renewable Energy Laboratory's (NREL) National Solar Radiation Database (NSRDB) (Sengupta et al., 2018). Following Chan and Zhou (2025), I measure local solar potential as the average solar global horizontal irradiance (GHI) of each commuting zone multiplied by its land area. I pair this geographically fixed potential with two time-varying shocks. The first is the Lawrence Berkeley National Laboratory's (LBNL) state-level Renewable Portfolio Standard (RPS) obligations, in gigawatt-hours, available annually from 2000 to 2024 (LBNL, 2025a). The second shifter is the inflation-adjusted global price of solar photovoltaic modules, or panels,

[10] 6 plays are excluded due to insufficient treatment colleges or years and 4 due to complete overlap with plays that experience earlier booms. I also drop Michigan Antrim whose boom coincided with the Great Recession and the collapse of the state's manufacturing base, making it difficult to separate these concurrent shocks from the boom's effect.

in dollars per watt, obtained from Our World in Data (Our World in Data, 2026). The shift-share-style measure is an instrument for solar investment, which I define as newly installed solar capacity. Utility-scale capacity comes from the EIA's Form 860, a federally mandated, annual survey of power plant operators (EIA, 2025b)[11]. I obtain distributed (residential and small commercial) capacity[12] from LBNL's Tracking the Sun database, which compiles national project-level data from utilities, state agencies, and other sources (LBNL, 2025b).

California Community Colleges. The California Community Colleges Chancellor's Office (CCCCO) provides data on all 116 community colleges operating in California from 1993 to 2025 through its public DataMart portal (California Community Colleges Chancellor's Office, 2026). The administrative data are compiled from the standardized Management Information System (MIS) submissions that each of the system's districts files with the Chancellor's Office annually. They cover students, courses, student services, outcomes, and faculty and staff. In addition to information on graduates and programs offered, CCCCO data include course-level information, tracking the number of sections offered for each course as well as enrollment at the TOP-code[13] level. Previous research has relied on the CCCCO data to study enrollment and course outcomes, including their responses to changing labor market conditions (Grosz, 2020, 2022).

Linked Administrative Records. For individual-level evidence on the labor-market returns to energy credentials, I use linked administrative microdata from the Texas ERC following Andrews et al. (2024). The data connects: (i) K–12 records from the Texas Education Agency (TEA), including high-school course-taking and career-and-technical participation, graduation and exit, and standardized achievement test scores; (ii) postsecondary records from the Texas Higher Education Coordinating Board (THECB), including enrollment, course-taking, and completion, together with each college's career-and-technical offerings; and (iii) quarterly unemployment-insurance earnings from the Texas Workforce Commission (TWC), available from 1990. I leverage these records for a descriptive analysis of the long-run earnings impacts of energy credentials for Texas students. Like Kovalenko (2023), I also use the individual-level data to study impacts of the shale boom, replicating my main

[11] All power plants with at least 1 megawatt (MW) of total nameplate capacity and connected to the grid are required by law to file annually under the Federal Energy Administration Act of 1974.

[12] To avoid double counting with the EIA data, I combine capacity reported at the same address and omit installations of 1MW (AC) or larger.

[13] The Taxonomy of Programs (TOP) is the six-digit classification the California Community Colleges Chancellor's Office uses to identify the discipline of a course or program; it is maintained by and used only within the California Community College system and is the state analogue of the federal CIP codes used elsewhere in this paper. The Chancellor's Office publishes a TOP-to-CIP crosswalk and applies it to convert program-level data from TOP to CIP for reporting to national databases such as IPEDS.

results for plays within Texas.

Other Variables. To supplement this data, I draw annual county population estimates from the National Cancer Institute's Surveillance, Epidemiology, and End Results (SEER) program (SEER Program, 2026). I measure county-level industry employment to construct a local employment shock using data from the Bureau of Economic Analysis (BEA) county employment by industry tables (CAEMP25N, 2001–2022) (U.S. Bureau of Economic Analysis, 2023). To approximate vocationally trained employment, I rely on state-level shares of middle-skill workers from the Current Population Survey (CPS) (Flood et al., 2025). I also leverage National Science Foundation (NSF) award records (1990–1995) (National Science Foundation, 2026) to measure administrative quality in my heterogeneity analysis. I use data on state-level funding policies for community colleges from InformEd States (InformEd States, n.d.), a higher-education policy data initiative whose datasets code state funding mechanisms, performance-based-funding policies and local and equalization funding sources. Other local demographic and economic characteristics are calculated using the Census Bureau's American Community Survey (ACS) (Ruggles et al., 2025).

3.2 Outcome Construction

For the bulk of my analyses, I do not have individual-level data on the occupations or industries in which graduates work, but rather community college reported data on programs and their graduates. Therefore, in order to analyze the industrial alignment of community colleges with local employment demand, I must first map programs to their corresponding occupations and industries.

Energy-Related Programs and Graduates. IPEDS reports information on graduates and programs according to their classification of instructional program (CIP) codes. CIP codes are standardized categories of fields of study maintained by the NCES and, as such, do not correspond directly to a single major or degree program[14]. Throughout the paper I refer to CIP codes as "programs" for ease of exposition, but it is important to note that they are not a one-to-one mapping. They often encompass multiple majors or programs and should be thought of as a coarser measure of educational offering. They are also used

[14]Six-digit CIP codes vary considerably in their level of detail. Examples from the IPEDS 2010 CIP code descriptions include: Mining Technology or Technician (15.0901) - "A program that prepares individuals to apply basic engineering principles and technical skills in support of engineers and other professionals engaged in the development and operation of mines and related mineral processing facilities."; and Education General (13.0101) - "A program that focuses on the general theory and practice of learning and teaching, the basic principles of educational psychology, the art of teaching, the planning and administration of educational activities, school safety and health issues, and the social foundations of education."

Table 2: Energy-Related CIP Codes

CIP Code	CIP Title	Weight
46.0504	Well Drilling/Driller	47.3%
15.0901	Mining Technology/Technician	40.7%
15.0999	Mining and Petroleum Technologies/Technicians, Other	33.3%
49.0202	Construction/Heavy Equipment/Earthmoving Equipment Operation	24.0%
41.0000	Science Technologies/Technicians, General	22.3%
15.0903	Petroleum Technology/Technician	21.8%
41.9999	Science Technologies/Technicians, Other	17.1%
41.0399	Physical Science Technologies/Technicians, Other	17.1%
14.2501	Petroleum Engineering	16.6%
14.2101	Mining and Mineral Engineering	15.1%
46.0505	Blasting/Blaster	13.0%
14.3901	Geological/Geophysical Engineering	11.5%
41.0303	Chemical Process Technology	7.3%
46.0000	Construction Trades, General	4.9%
15.0701	Occupational Safety and Health Technology/Technician	3.5%
47.0302	Heavy Equipment Maintenance Technology/Technician	3.5%
46.0502	Pipefitting/Pipefitter and Sprinkler Fitter	2.7%
46.0415	Building Construction Technology/Technician	2.6%
46.0403	Building/Home/Construction Inspection/Inspector	2.6%
46.0402	Concrete Finishing/Concrete Finisher	2.5%
46.0201	Carpentry/Carpenter	2.5%

Notes. This table lists the CIP codes classified as energy programs under the binary 2.5% threshold. Weight is the share of a program's graduates expected to enter the energy sector (NAICS 21) constructed using the CIP-to-NAICS mapping described in Appendix C.

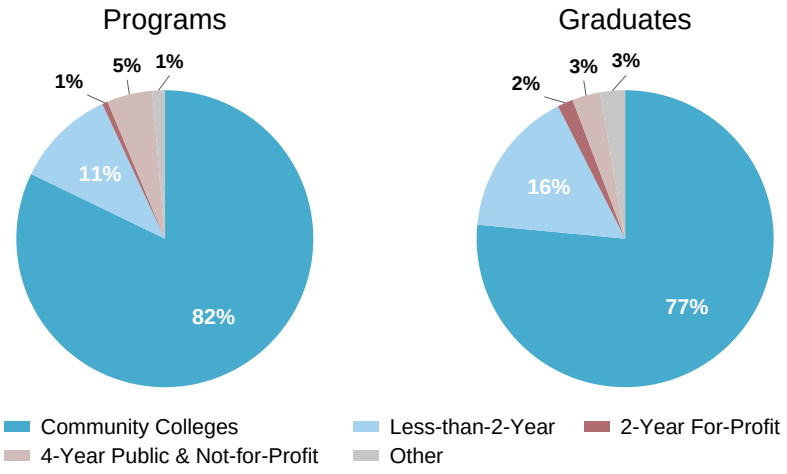
to capture information on instructional programs and degree or certificate completions for many educational surveys and datasets.

Similar to existing research mapping programs to occupations, I begin by mapping each six-digit CIP code to the set of occupations (SOC codes) that would be theoretically suited to the graduates of that program using a crosswalk jointly created by the NCES and Bureau of Labor Statistics (Long et al., 2015; Acton, 2021; Grosz, 2022). This crosswalk is formed on the basis that a program provides the skills necessary to perform the tasks required to work in a given occupation. I give equal weight to all occupations associated with a program, consistent with Long et al. (2015) who report minimal sensitivity to differing specifications.

I then match occupations to industries (2-digit NAICS codes) using employment shares from the BLS' National Occupation-Industry Employment Matrix[15] (NOIEM, 2019). The resulting weights yield a mapping from each CIP code to the set of industries its graduates are expected to enter with each weight representing the theoretical share of graduates from a given program (six-digit CIP code) that work in a given industry. See Figure C1 for more

[15]I use the 2019 BLS NOIEM in my final analysis but have run sensitivity analyses to other years (2016, 2018, 2020). Results remain consistent in directionality and size.

Figure 5: Energy Programs and Graduates by Institution Type



Notes. This figure shows the share of energy programs (left) and energy graduates (right) by institution type for certificates and associate degrees across all energy CIP codes. A program is a CIP code and award level reported by an institution in a year. Shares are computed in each year and averaged over 1990–2024. Community colleges are public two-year institutions, and Other includes four-year for-profit and two-year not-for-profit private institutions. Data are from NCES IPEDS Completions.

details on my mapping procedure.

Programs are classified as "energy" (NAICS 21[16]) or "non-energy" (all other industries) using a binary decision rule. A program is an energy program if at least 2.5%[17] of its graduates are expected to enter the energy sector, corresponding to approximately the top 20 CIP codes expected to send graduates to the energy industry (see Table 2 for the full list). Figure 5 shows that community colleges are the main providers of this training. Over 1990–2024, they account for 82% of energy certificate and associate programs and 77% of energy graduates at these levels, with most of the remainder at less-than-two-year institutions (Figure A2 details the institutional breakdown for each energy program). Non-energy programs are all other career and technical programs that are not categorized as energy programs. The binary rule has the advantage of transparency and a clear interpretation, with each program assigned to a single industrial category. As a robustness check, Appendix E.1 reports results from an alternative "weighted-CIP" measure that uses the sum of the fractional shares of all programs sending graduates to each industrial category to calculate the outcome measures. Results are consistent across specifications.

[16]North American Industry Classification System (NAICS) code 21 includes mining, quarrying, and oil and gas extraction.

[17]Threshold was determined by manual analysis of the programs categorized as energy and their corresponding association with oil and gas employment.

My primary outcomes are simple counts of the number of energy and non-energy programs and the number of energy and non-energy completions offered at an institution in a given year. Let $E_m \equiv \mathbb{1}\{\omega_{m,\text{Energy}} \geq 0.025\}$ denote the energy decision rule for program m. Then

$$M_{i,t}^{\text{Energy}} = \sum_{m \in \mathcal{M}_{i,t}} E_m, \tag{1}$$

$$M_{i,t}^{\text{Non-Energy}} = \sum_{m \in \mathcal{M}_{i,t}} (1 - E_m), \tag{2}$$

where $M_{i,t}^{\text{Energy}}$ is the count of energy programs (CIPs) offered at institution i in year t and $M_{i,t}^{\text{Non-Energy}}$ is the count of non-energy programs. $\mathcal{M}_{i,t}$ is the set of all programs offered at institution i in year t and $\omega_{m,\text{Energy}}$ is the weight measuring the share of graduates from program m expected to enter the energy industry. Energy and non-energy graduates (completions), $G_{i,t}^{\text{Energy}}$ and $G_{i,t}^{\text{Non-Energy}}$, are constructed similarly by summing the number of graduates, $g_{i,t,m}$, across all energy and non-energy programs, respectively:

$$G_{i,t}^{\text{Energy}} = \sum_{m \in \mathcal{M}_{i,t}} E_m \cdot g_{i,t,m}, \tag{3}$$

$$G_{i,t}^{\text{Non-Energy}} = \sum_{m \in \mathcal{M}_{i,t}} (1 - E_m) \cdot g_{i,t,m}. \tag{4}$$

I also assess compositional changes by looking at the share of energy programs and graduates relative to the total number of programs and graduates, respectively. This allows me to determine whether the boom leads to a reorientation of the institutional portfolio toward energy fields even if the total number of programs or graduates does not change. Finally, I explore the margin of expansion whether through existing or new programs by analyzing the number of graduates from new and existing programs. Graduates of new programs are defined as completions within the first three years a program is offered to allow for the time it takes students to enroll and complete a program. Existing programs are those that were offered within the two years preceding a shale boom, but created three or more years before the boom.

Solar Sections, Programs, and Graduates. Unlike energy programs, solar programs (and courses) are easier to distinguish given the limited number of occupations and CIP codes historically sending graduates to solar jobs. Recent research measures green employment using the task content of occupations (Vona et al., 2019) and identifies renewable employment using job titles, skill requirements, and employer information (Curtis and Marinescu, 2023; Curtis et al., 2024). I build on this literature by identifying solar training in community

colleges using program classifications (CIP codes) and course titles. There are only two programs in the IPEDS data that can credibly be assigned to a "solar" category and existed prior to 2020. Thus, I use a narrow definition of solar programs comprising Solar Technician and Energy Systems Technician.

In the CCCCO data on section offerings, I manually identify "solar" courses using keywords in their titles, including "solar" and "photovoltaic". Because the CIP classification for solar is limited by historical categorization, programs in the CCCCO data are designated as solar based on the course-level data. This approach captures solar instruction within broader vocational programs that may not have a solar-specific CIP code. The top 10 programs in terms of their share of solar courses are classified as solar programs. Non-solar courses and programs are all career and technical courses and programs that are not designated as solar excluding those that are closely aligned with solar energy activities (e.g. electricians, linemen, etc.) and other programs that include smaller shares of solar courses.

The main outcomes for the solar analysis are the number of course sections offered, programs, and graduates across solar and non-solar designations. Similar to the shale analysis, I also assess composition changes using the share of solar sections, programs, and graduates relative to the total number of sections, programs, and graduates, respectively.

4 The Shale Boom and Community College Responses

Do community colleges respond to the boom and if so, how? In this section, I describe the research design and establish its validity, then trace the response through program offerings and completions. I assess the observed heterogeneity in the response, investigating the role of institutional characteristics and local labor market conditions, as well as the fields that give way to make room for energy expansion.

4.1 Identification Strategy and Estimation

I estimate the effect of the shale boom on community college outcomes using stacked difference-in-differences and event study designs in which each play constitutes a separate quasi-experiment.

Treatment is assigned based on whether a community college lies inside the geographic footprint of a shale play. Treated colleges are those physically located within a play's boundary and controls are located outside the play but within the same state(s) as the play, ensuring that treated and control institutions share common state-level policy environments

and funding regimes. As a robustness test, I also implement donut style designs where colleges within buffer zones of 10, 25, 50, and 75 miles from the shale play are excluded from the analysis to address concerns about potential contamination of the control group. The energy results in Appendix E.2 do not change materially.

Following Bartik et al. (2019), treatment timing is play-specific and defined as the first year a play's "fracking potential" became "public knowledge" for the 15 plays that are included in their paper. For the remaining two plays, I use the first year the play reaches 10 completed horizontal wells, which is relatively similar to the "public knowledge" date for the plays where I observe both. Because shale plays became economically viable at different times related to their geological characteristics and the ease with which oil or gas could be extracted, treatment onset ranges from 2001 to 2012. Table B1 provides a detailed description of treatment designation, college counts, and timing for each play.

In the style of Michaels (2011) and similar to prior work using the shale boom as a natural experiment (Feyrer et al., 2017; Allcott and Keniston, 2018; Bartik et al., 2019; Cascio and Narayan, 2022; Kovalenko, 2023), my baseline specification is:

$$y_{i,p,t} = \alpha_{i,p} + \gamma_{p,t} + \beta \left(\text{Treat}_{i,p} \times \text{Post}_{p,t}\right) + \varepsilon_{i,p,t} \tag{5}$$

where $y_{i,p,t}$ is an outcome such as energy programs for community college, i, in play, p, and year t; $\alpha_{i,p}$ are institution-by-play fixed effects controlling for the characteristics of an institution that are fixed across time (for example, founding mission, geographic location within a play, and campus size); and $\gamma_{p,t}$ are play-by-year fixed effects absorbing play-specific annual shocks that are constant across institutions within a play (for example, gas price shocks, fuel export costs, and local economic conditions).

Because it takes time for colleges to start new programs, I estimate a version of Equation 5 that replaces the single treatment indicator, $\text{Treat}_{i,p} \times \text{Post}_{p,t}$, with separate indicators for shale-exposed colleges in the short run (0–4 years after onset), medium run (5–9 years), and long run (10–14 years):

$$\begin{aligned} y_{i,p,t} = \alpha_{i,p} + \gamma_{p,t} + \beta^{SR}\left(\text{Treat}_{i,p} \times \text{Post}^{SR}_{p,t}\right) + \beta^{MR}\left(\text{Treat}_{i,p} \times \text{Post}^{MR}_{p,t}\right) \\ + \beta^{LR}\left(\text{Treat}_{i,p} \times \text{Post}^{LR}_{p,t}\right) + \varepsilon_{i,p,t} \end{aligned} \tag{6}$$

where $\text{Post}^{SR}_{p,t}$, $\text{Post}^{MR}_{p,t}$, and $\text{Post}^{LR}_{p,t}$ indicate years 0–4, 5–9, and 10–14 after onset.

To assess full dynamics and pre-trends, I also estimate the analogous event-study spec-

ification:

$$y_{i,p,t} = \alpha_{i,p} + \gamma_{p,t} + \sum_{k \neq -1} \beta_k \left(\text{Treat}_{i,p} \times \mathbf{1}\{t - \text{onset}_p = k\} \right) + \varepsilon_{i,p,t} \tag{7}$$

where $k = -1$ is the omitted reference period. Pre-treatment coefficients β_k for $k < 0$ test the parallel-trends assumption. To address the negative weighting and bias from heterogeneous treatment effects that can arise in conventional two-way fixed effects models with staggered adoption (Goodman-Bacon, 2021; Sun and Abraham, 2021; de Chaisemartin and d'Haultfœuille, 2020), I implement the Gardner (2021) two-stage estimation procedure in all my specifications. Standard errors are clustered at the play-level.

Identification. In this setting, each play p defines a stack, and G_p is the year play p starts its shale boom. Treated, or shale-exposed, colleges are located within a play's boundaries, with $B_{i,p}$ the binary indicator for whether college i is located within play p. Control, or comparison, colleges are those in the same states as a play p but outside the play's boundaries, $B_{i,p} = 0$. Because comparison institutions for a given play p can be located within a separate play p', $\mathcal{B}_i = \max_p B_{i,p}$ indicates whether a college is located within any play, and $p(i)$ denotes the play containing a shale-exposed college i.

For each play p, treated colleges are the set:

$$\mathcal{T}_p = \{\, i \;:\; B_{i,p} = 1 \,\}$$

and control colleges are the set:

$$\mathcal{C}_{p,t} \;=\; \{\, i \;:\; s(i) \in S_p, \;\; B_{i,p} = 0, \;\; (\mathcal{B}_i = 0 \;\text{ or }\; t < G_{p(i)}) \,\},$$

where $s(i)$ denotes college i's state and S_p the set of states play p intersects. As is common in a stacked framework, control colleges for play p comprise both never-treated colleges across all years, $\mathcal{B}_i = 0$, and late-treated colleges in other plays in their pre-boom period, $t < G_{p(i)}$. Treatment for college i in play p at time t is defined as: $D_{i,p,t} = B_{i,p} \cdot \mathbb{1}\{t \geq G_p\}$.

Let $Y_{i,t}(g)$ denote the potential outcome for college i in year t if a local shale boom begins in year g, and $Y_{i,t}(0)$ the potential outcome absent a boom. The causal effect of the boom on college i in year t is $Y_{i,t}(g) - Y_{i,t}(0)$, the difference in its potential outcomes. However, observed outcomes satisfy:

$$Y_{i,t} \;=\; (1 - \mathcal{B}_i) \cdot Y_{i,t}(0) \;+\; \mathcal{B}_i \cdot Y_{i,t}(G_{p(i)})$$

In words, only one potential outcome can be observed for each college ($Y_{i,t}(G_{p(i)})$ for shale-exposed colleges and $Y_{i,t}(0)$ for never-treated colleges), so the causal effect for each institution can never be directly observed. Instead, I estimate the average effect of the boom on shale-exposed colleges, or the average treatment effect on the treated (ATT), at event time $k = t - G_p$ relative to boom onset:

$$ATT(k) \;=\; E\Big[\, Y_{i,\,G_{p(i)}+k}(G_{p(i)}) \;-\; Y_{i,\,G_{p(i)}+k}(0) \;\Big|\; \mathcal{B}_i = 1 \,\Big] \tag{8}$$

I compare shale-exposed colleges to control colleges in the same state(s) but outside the play. The average difference in outcomes between shale-exposed and control colleges identifies the causal effect of the boom on shale-exposed colleges if control colleges trace the counterfactual path for outcomes absent the boom, $Y_{i,G_{p(i)}+k}(0)$. The assumptions below describe the conditions under which β_k in Equation 7 recovers the $ATT(k)$.

Assumption 1 (No anticipation). A college's outcomes do not respond to a shale boom prior to its onset. That is, for every college i, boom year g, and year $t < g$: $Y_{i,t}(g) = Y_{i,t}(0)$.

Assumption 2 (No spillovers). A play's boom has no impact on outcomes at colleges outside its boundary. Accordingly, the outcomes of shale-exposed colleges in a given play are unaffected by booms in other plays, and the outcomes of never-treated colleges are unaffected by booms in any play.

Assumption 3 (Parallel trends). Absent the boom, outcomes at shale-exposed colleges would have evolved in parallel with outcomes at control colleges. Formally, for each play p and all years t, t':

$$\mathbb{E}\Big[\, Y_{i,t}(0) - Y_{i,t'}(0) \,\Big|\, i \in \mathcal{T}_p \,\Big] \;=\; \mathbb{E}\Big[\, Y_{i,t}(0) - Y_{i,t'}(0) \,\Big|\, i \in \mathcal{C}_{p,t} \cap \mathcal{C}_{p,t'} \,\Big]$$

Together, Assumptions 1 and 2 imply that the observed outcomes of all untreated colleges are their potential outcomes absent a boom. That is, for never-treated control colleges in any year t: $Y_{i,t} = Y_{i,t}(0)$, and for shale-exposed colleges in pre-boom years $t < G_{p(i)}$: $Y_{i,t} = Y_{i,t}(0)$. From this, it follows that the difference in observed outcomes between shale-exposed and control colleges in pre-boom years is equal to the difference in their potential outcomes absent a boom. Assumption 3 then implies that, on average, this difference is constant over time, so the counterfactual untreated potential outcome for shale-exposed colleges in post-boom years can be reconstructed from the observed outcomes of control colleges and the constant differential. The average difference between observed post-boom outcomes for shale-exposed colleges and this counterfactual identifies the causal effect of the boom on

shale-exposed colleges for each play p and event time k:

$$ATT(p,k) \;=\; \mathbb{E}\Big[\, Y_{i,\,G_p+k}(G_p) \;-\; Y_{i,\,G_p+k}(0) \;\Big|\; i \in \mathcal{T}_p \,\Big] \tag{9}$$

which averaged across plays gives the $ATT(k)$:

$$ATT(k) \;=\; \sum_{p} \omega_p \, ATT(p,k),$$

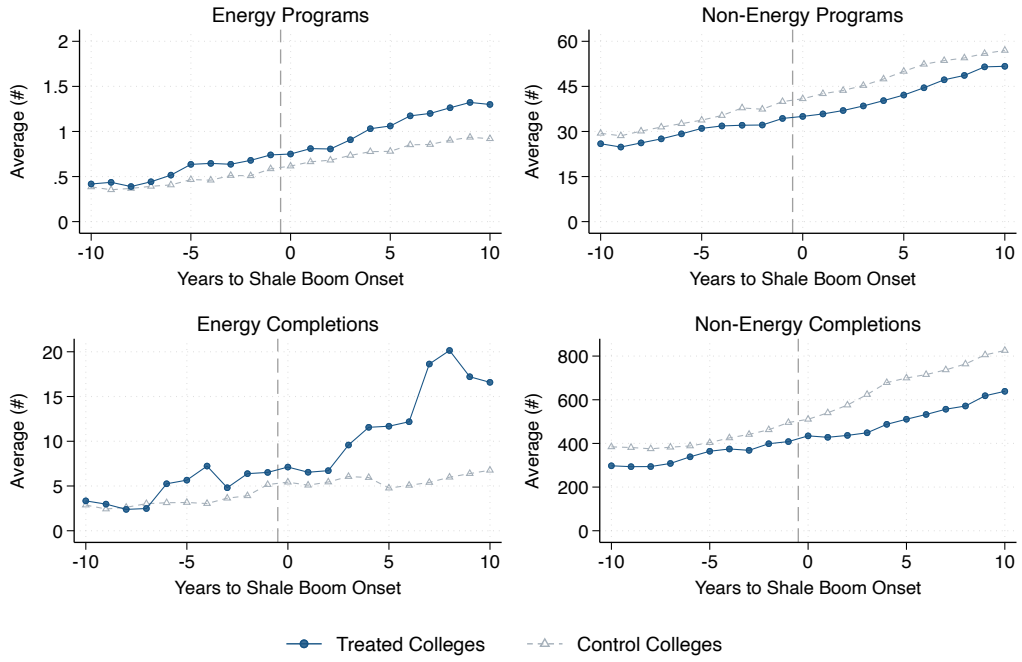
where ω_p is the share of shale-exposed colleges in play p. The Gardner (2021) estimator weights each play by its share of treated observations, in this case college-years, avoiding the variance based weights of single-stage stacked regressions that require correction to recover the $ATT(k)$ (Wing et al., 2024).

A separate challenge in staggered designs is that conventional two-way fixed effects estimators use already-treated units as controls. When effects change over time, as they do here, these "forbidden comparisons" contaminate the estimate, introducing bias (de Chaisemartin and d'Haultfœuille, 2020; Goodman-Bacon, 2021; Sun and Abraham, 2021). My design addresses this issue in two ways. First, the stacked structure treats each play as a separate experiment with its own never- and late-treated, in-state control group. As such, a college that has already been treated never serves as a comparison unit. Second, the Gardner (2021) two-stage procedure only uses untreated observations to estimate fixed effects, so treated post-period observations never actually enter the counterfactual. Results are robust to specifications that exclude late-treated units from the control group altogether.

There are a number of reasons the above identifying assumptions may not hold. If students and colleges are aware that shale development is set to begin, institutions may adjust their offerings and students may adjust their enrollment decisions prior to the boom onset in anticipation of the labor demand shock. Furthermore, the economic effects of the shale boom were not confined to the counties in which drilling occurred (Feyrer et al., 2017; Allcott and Keniston, 2018). If the spillover income and employment effects are strong enough, they might affect community colleges and students outside play boundaries. Finally, educational outcomes at colleges inside and outside shale plays may have followed different trajectories even in the absence of the boom.

Starting with the final concern, parallel trends cannot be tested directly; however, the pre-boom event study coefficients from Equation 7 provide evidence that outcomes were on similar trends prior to the boom. Figures 8 and 9 report these results where I find null effects across outcomes in most pre-treatment periods. Figure 6 plots the unadjusted means

Figure 6: Raw Means by Event Time, Treated vs. Control Colleges



Notes. The top panels plot the unadjusted mean number of energy (left) and non-energy (right) programs offered per community college, separately for treated (in-play) and control (out-of-play, in-state) institutions, by event time relative to the play's boom-onset year. The bottom panels plot the same for the unadjusted mean number of energy and non-energy completions.

of key outcomes for shale-exposed and control colleges over event time. Consistent with the event study results, energy offerings and completions evolve similarly in most pre-boom years. Additionally, Table B2 documents that community colleges in and outside plays were largely similar along observable baseline characteristics.

The evidence against pre-trends can also be used to assess anticipation. The lack of significant effects in the years just prior to the boom suggests shale-exposed colleges were behaving similarly in the lead up to shale drilling. The results are also robust to specifications that restrict the comparison group to never-treated colleges, mitigating concerns about late-treated comparisons. Finally, I address potential spillovers by successively excluding control colleges near play boundaries at 10, 25, 50, and 75 mile rings. The estimates, reported in Appendix E.2, do not change materially for energy outcomes, while the non-energy program effect weakens and the non-energy completion effect strengthens at wider buffers. Although these exercises cannot directly verify the identifying assumptions, they provide evidence on their observable violations and support a causal interpretation of the estimates.

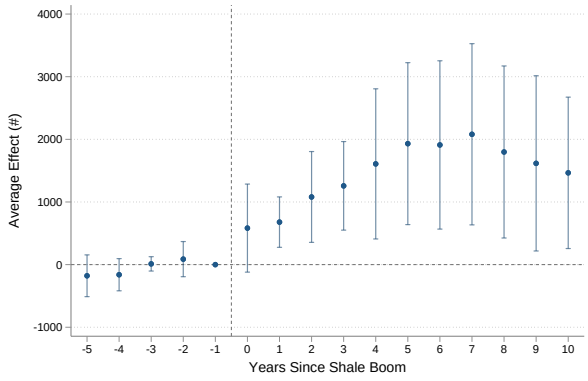
Table 3: Effect of Shale Boom on Local Energy Employment

Outcome:	Energy Employment	Middle-Skill Energy Employment	Energy Share	Middle-Skill Energy Share
Treat × Post	1,455***	264**	0.99*	0.89
	(527)	(108)	(0.56)	(0.69)
CZ FE	Y	Y	Y	Y
Play × Year FE	Y	Y	Y	Y
Observations	8,547	8,547	8,547	8,547
Plays	17	17	17	17
Baseline Mean	3,146	695	3.89	3.37

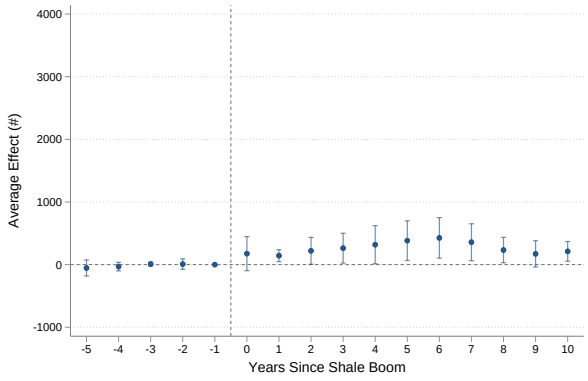
Notes. This table presents the Gardner (2021) two-stage difference-in-differences estimates from Equation 5 for commuting-zone employment outcomes. Columns 1 and 2 report the effect on energy (NAICS 21) employment and middle-skill energy employment, respectively, and columns 3 and 4 report energy's share of total employment and of middle-skill employment in percentage points. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are clustered at the play level in parentheses.

Figure 7: Shale Effects on Local Energy Employment Over Time

(a) Energy Employment



(b) Middle-Skill Energy Employment



Notes. This figure plots the Gardner (2021) two-stage event study estimates for local energy employment (NAICS 21) (a) and middle-skill energy employment (b) outcomes. Standard errors are clustered at the play level.

4.2 Labor Demand Shock

Before turning to community college outcomes, I show that the shale boom raised local demand for energy workers. I estimate Equations 5 and 7 for total energy employment and middle-skill energy employment at the commuting zone level. I approximate middle-skill energy employment by applying state-level CPS shares of energy workers with an occupational associate's degree or some college but no degree to commuting zone energy employment. Table 3 and Figure 7 present the results. On average, shale-exposed commuting zones employ 1,455 more energy workers than non-shale exposed commuting zones in the years after the

Table 4: Effect of Shale Boom on Programs Offered

Outcome:	Energy CIPs	Non-Energy CIPs	Energy's Share of CIPs
Raw Outcome:			
Treat × Post	0.16**	-2.93***	0.25
	(0.06)	(1.05)	(0.24)
Treat × 0–4 Years	0.05	-2.58***	0.14
	(0.07)	(0.83)	(0.18)
Treat × 5–9 Years	0.20**	-3.30**	0.47*
	(0.08)	(1.30)	(0.28)
Treat × 10–14 Years	0.23**	-2.62	0.33
	(0.10)	(1.66)	(0.35)
Effect per SD:			
Treat × Post	0.33**	-0.26***	0.10
	(0.14)	(0.09)	(0.10)
Treat × 0–4 Years	0.10	-0.23***	0.06
	(0.14)	(0.07)	(0.08)
Treat × 5–9 Years	0.42**	-0.29**	0.20*
	(0.17)	(0.12)	(0.12)
Treat × 10–14 Years	0.49**	-0.23	0.14
	(0.21)	(0.15)	(0.15)
CC FE	Y	Y	Y
Play × Year FE	Y	Y	Y
Observations	32,566	32,566	32,566
Plays	17	17	17
Baseline Mean	0.52	28	1.80
SD of ΔOutcome	0.47	11	2.40

Notes. This table presents the Gardner (2021) two-stage difference-in-differences estimates from Equation 5 and Equation 6 for community college programs. Columns 1 and 2 report the effect on the number of energy and non-energy programs (CIPs), respectively, and Column 3 reports energy's share of programs (CIPs) in percentage points. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Standard errors are clustered at the play level in parentheses.

boom, an increase of nearly 50% over the baseline mean. Middle-skill energy employment rises by 264 workers, around 40% of its baseline mean. The event study shows no differential trends before the boom, after which energy employment rises steadily, peaks five to seven years after onset, and remains elevated a decade later. These results establish that the shale boom created a large and sustained increase in local demand for energy workers, including the middle-skill workers that community colleges train. In the next subsection, I explore whether those institutions responded.

4.3 Effects on Program Offerings

A direct measure of institutional supply is the count of energy programs a college offers, and on this margin I find that community colleges respond meaningfully to the shale boom. Table 4 presents the estimates of Equation (5) and Equation (6) for outcomes of program offerings. I report effects in standard deviations (SD) of the pre-boom annual change in the outcome, with standard errors in parentheses. Shale-exposed colleges increase their energy

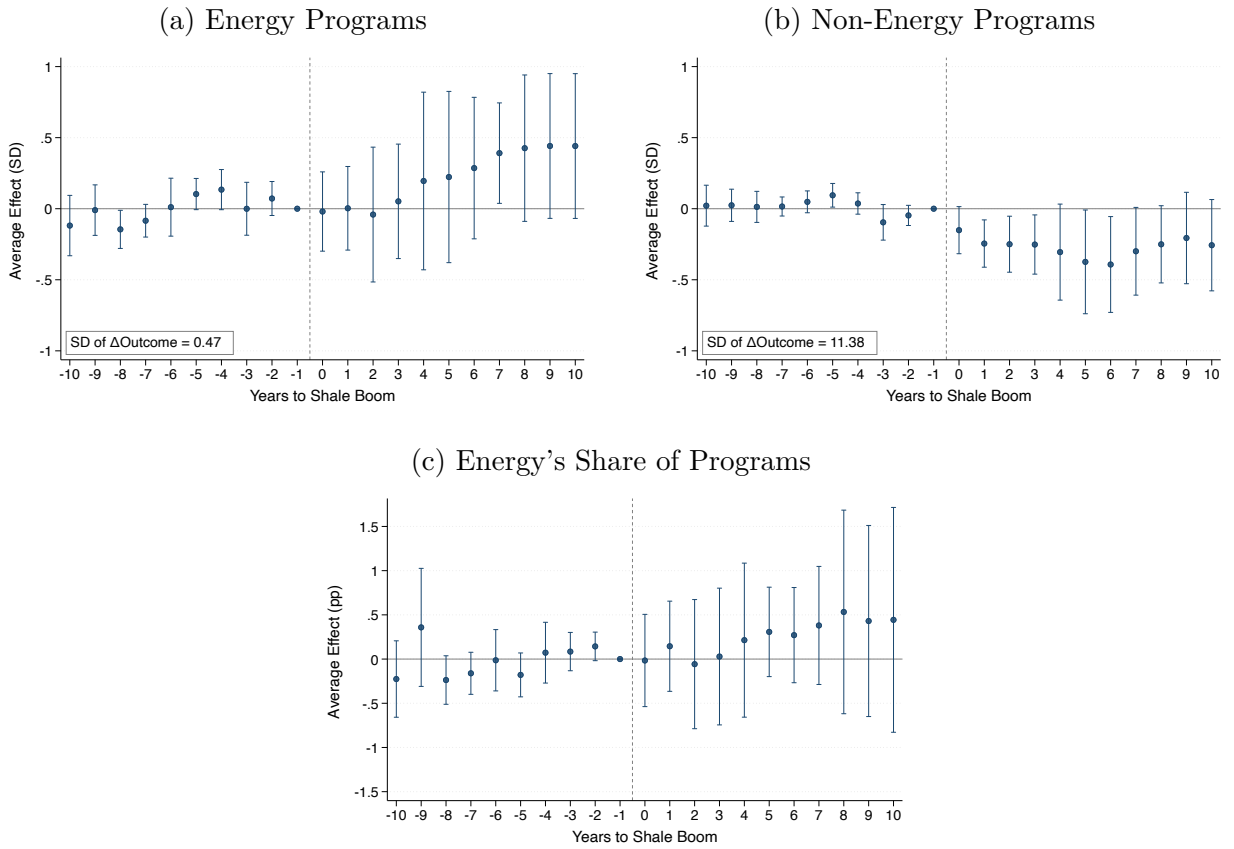
program offerings by 0.33 SD (0.14) on average in the post-boom period, a 30% increase relative to their pre-period mean of 0.52 programs. The size of the effect reflects the narrow base of energy programs at the typical community college. However, in most years, colleges do not change their energy offerings at all. The distribution of annual changes prior to the boom has an interquartile range of zero and a standard deviation of 0.47, with a long right tail. In fact, the change is zero from the 10th to the 90th percentile of the distribution. Against this baseline, a sustained increase of 0.33 SD is large and represents a substantial proportional expansion in institutional capacity directed at the booming sector. The effect builds over time, with no detectable change in the first five years after onset, and remains largely stable in the medium- to long-term at 0.42 SD (0.17) and 0.49 SD (0.21), respectively. This is consistent with new programs that take several years to launch, but persist once established.

The increase in energy programs is accompanied by a relative decline in non-energy fields. Treated colleges offer 0.26 SD (0.09) fewer non-energy programs compared to non-shale-exposed colleges, representing 10% of the pre-period mean of 28 programs. The raw trends in Figure 6 indicate that this gap opens up because treated colleges stop keeping pace with the continued non-energy program growth at control colleges, rather than because they actively cut existing offerings. The non-energy response begins immediately after onset, with a -0.23 SD (0.07) effect in years 0–4, is the strongest in the medium run, -0.29 SD (0.12), and remains similar in magnitude but becomes imprecise after 10 years.

Together, these shifts produce an increase in energy's share of programs of 0.25 percentage points (0.24) over the post-period, though the estimate is not statistically significant. The share rises most in the medium run, by 0.47 percentage points (0.28), significant at the 10% level. Since both margins move in the expected direction, the lack of statistical precision is likely due to the small number of energy programs at the typical community college and the resulting noise in the share measure, producing a large standard error relative to a small point estimate.

The event-study (Figure 8) illustrates the full dynamics of this adjustment. Pre-treatment coefficients are close to zero and jointly insignificant, providing no evidence of differential trends prior to the boom. The treatment effect on energy programs begins to emerge approximately five years after the boom hits, though not statistically significant across the sample except in year seven, and grows steadily over the decade. This is consistent with the lead time required for institutions to develop, approve, staff, and implement new programs. Non-energy programs show an immediate contraction in the boom year that

Figure 8: Shale Effects on Programs Over Time



Notes. This figure plots the Gardner (2021) two-stage event study estimates for program outcomes. Panel (a) displays energy programs, panel (b) displays non-energy programs, and panel (c) displays energy's share of programs. Panels (a) and (b) are in standard deviations of the pre-boom annual change in the outcome, and panel (c) is in percentage points. Standard errors are bootstrapped and clustered at the play level.

progresses modestly until beginning to revert to the baseline in year seven. The event study for energy's share of programs (Figure 8c) is proportionally consistent with these patterns.

An interesting feature of the timing of program response is that it appears to be asymmetric. While the increase in energy programs takes several years to materialize, the shortfall in non-energy programs is immediate. This suggests that the boom may have a more rapid effect on the extensive margin of program discontinuation than on the extensive margin of program creation, which is consistent with the idea that it is easier for colleges to stop offering programs than to start new ones.

Table 5: Effect of Shale Boom on Completions

Outcome:	Energy Completions	Non-Energy Completions	Energy's Share Completions
Raw Outcome:			
Treat × Post	5.53***	-105**	0.57**
	(1.29)	(46)	(0.23)
Treat × 0–4 Years	0.86	-62**	0.35
	(1.20)	(28)	(0.22)
Treat × 5–9 Years	8.33***	-115**	1.08***
	(1.96)	(55)	(0.36)
Treat × 10–14 Years	9.95***	-137*	0.77***
	(2.61)	(75)	(0.27)
Effect per SD:			
Treat × Post	0.66***	-0.43**	0.23**
	(0.15)	(0.18)	(0.09)
Treat × 0–4 Years	0.10	-0.25**	0.14
	(0.14)	(0.12)	(0.09)
Treat × 5–9 Years	0.99***	-0.47**	0.43***
	(0.23)	(0.22)	(0.14)
Treat × 10–14 Years	1.18***	-0.55*	0.31***
	(0.31)	(0.30)	(0.11)
CC FE	Y	Y	Y
Play × Year FE	Y	Y	Y
Observations	32,566	32,566	32,551
Plays	17	17	17
Baseline Mean	4.71	340	1.09
SD of ΔOutcome	8.42	247	2.51

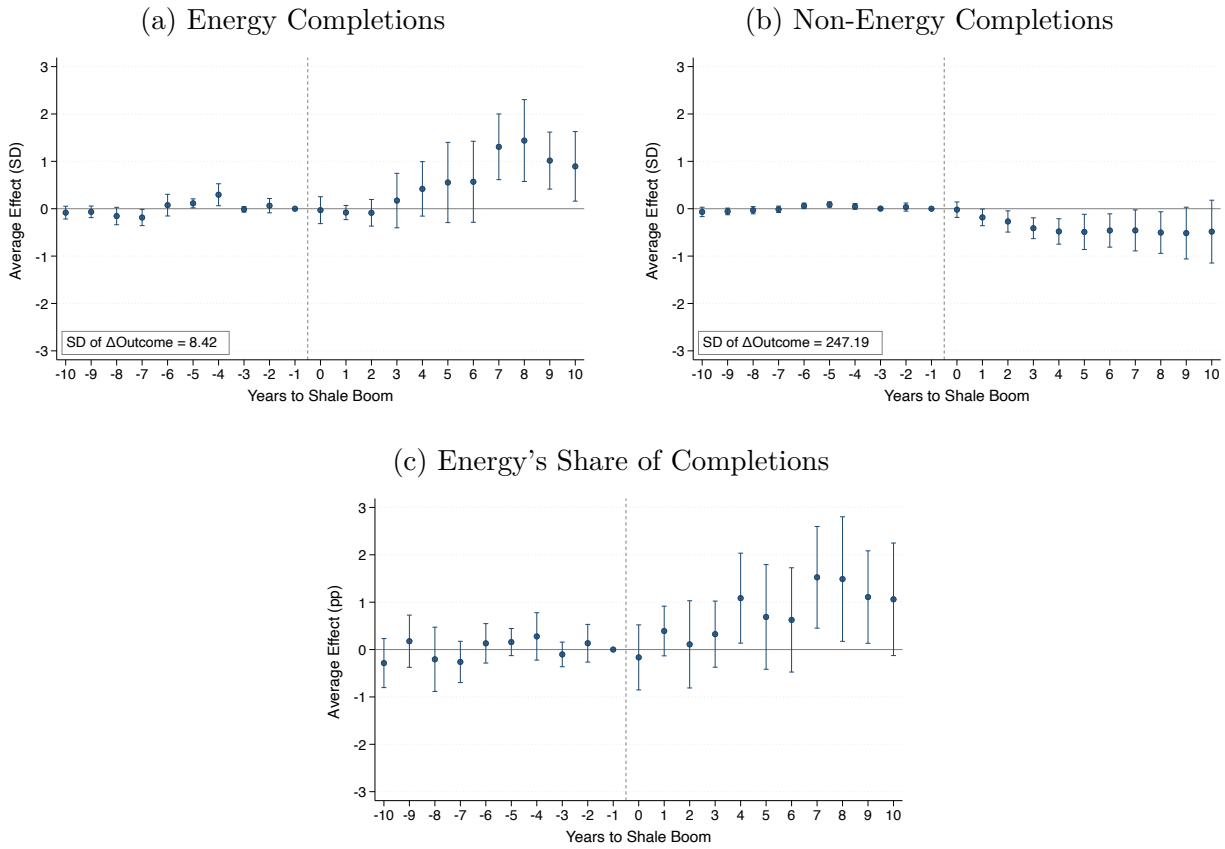
Notes. This table reports the Gardner (2021) two-stage difference-in-differences results of Equation 5 and Equation 6 with community college completions as the outcomes. The outcome in column 1 is the count of energy completions; in column 2, the count of non-energy completions; and in column 3, energy's share of completions (percentage points). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the play level in parentheses.

4.4 Effects on Completions

Unlike program offerings, which reflect institutional supply, completions are the product of both supply and demand. The boom may increase completions by creating new energy programs or expanding existing ones, but it may also reduce them by pulling students out of college or into other fields. Here the pattern mirrors the program offerings results but is sharper. Table 5 reports completion-based results for Equation (5) and Equation (6).

Energy completions rise by 0.66 SD (0.15) on average over the post-period, more than doubling the pre-period mean of 4.7 (117% increase). In raw terms, this is about 5.5 additional energy completions per college per year. This outsized effect is consistent with the existing literature on postsecondary education, which suggests that students are more dynamic and responsive to labor market changes than institutions themselves (Grosz, 2022, 2020; Conzelmann et al., 2023). Similar to programs, there is no detectable effect on completions in the first 5 years after the boom's onset, after which the effect grows to 0.99 SD (0.23) in the medium term, and remains near there after 10 years at 1.2 SD (0.31). This

Figure 9: Shale Effects on Completions Over Time



Notes. This figure plots the Gardner (2021) two-stage event study estimates for completions outcomes. Panel (a) displays energy completions, panel (b) displays non-energy completions, and panel (c) displays energy's share of completions. Panels (a) and (b) are in standard deviations of the pre-boom annual change in the outcome, and panel (c) is in percentage points. Standard errors are bootstrapped and clustered at the play level.

path is consistent with students completing energy programs after new ones are established.

On the non-energy side, treated colleges have 0.43 SD (0.18) fewer completions compared to non-shale-exposed colleges, a 31% decline relative to the pre-period mean of 340. Like non-energy programs, the change in non-energy completions begins in the short run, but unlike programs, it continues to grow through the long run, reaching 0.55 SD (0.30) 10–14 years after onset. The net effect shifts the composition of graduates toward energy fields, raising energy's share of completions by 0.57 percentage points (0.23) on average. The shift is largest in years 5–9, when energy's share of graduates doubles, rising by 1.08 percentage points (0.36).

The decline in non-energy completions is large, and consistent with the opportunity-cost channel documented in the resource-boom literature (Cascio and Narayan, 2022; Weinstein,

Table 6: Effect of Shale Boom on Completions in Pre-Existing Programs

Outcome:	Completions in Existing Energy CIPs	Completions in Existing Non-Energy CIPs
Raw Outcome:		
Treat × Post	-0.01	-38
	(0.88)	(30)
Effect per SD:		
Treat × Post	0.00	-0.42
	(0.18)	(0.33)
CC FE	Y	Y
Play × Year FE	Y	Y
Observations	32,218	32,218
Plays	17	17
Baseline Mean	4.03	285
SD of ΔOutcome	4.97	90

Notes. This table reports the Gardner (2021) two-stage difference-in-differences results of Equation 5 with community college completions in programs that existed prior to the boom as the outcomes. The outcome in column 1 is the completions in existing energy programs, and in column 2 it is the completions in existing non-energy programs. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the play level in parentheses.

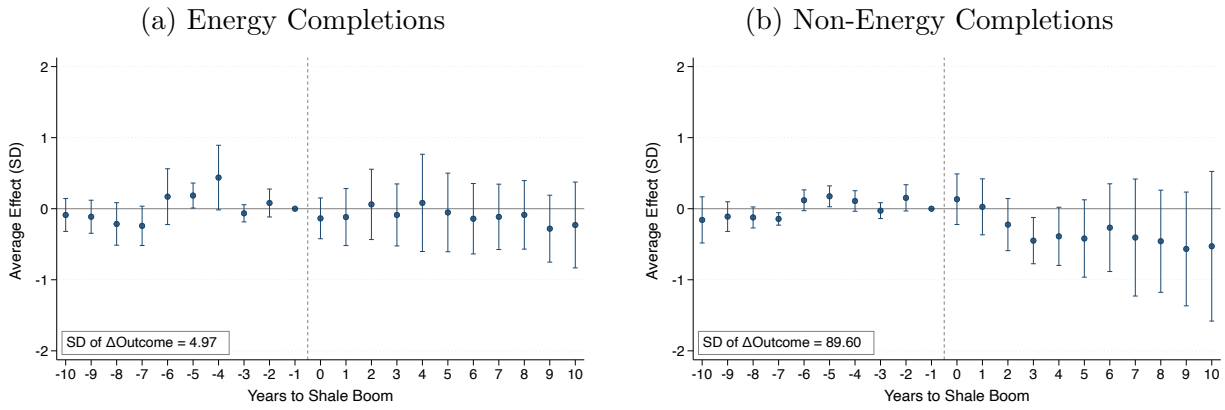
2022; Kovalenko, 2023). High wages in the booming sector draw students out of non-energy programs and out of college altogether. What is striking is that energy completions rise even as total completions fall. The boom pulls students out of the classroom on average, yet community colleges still graduate substantially more energy-aligned workers, providing evidence that the institutional response is strong enough to withstand the downward pressures on overall enrollment. The event study (Figure 9) shows a similar pattern to the program response with relatively flat pre-trends and a steady post-boom rise in energy completions, peaking around years seven to eight, while non-energy graduates respond almost immediately with a gradual decline that levels off about four years after the boom. Again, it is clear that energy programs must exist and students must enroll and complete them before an increase in graduates can be observed, but students can leave non-energy programs at any time, creating the asymmetry in the timing of the response.

4.5 Margins of Adjustment

4.5.1 New versus Existing Energy Programs

Where do the additional energy graduates come from? Colleges could meet new demand by graduating more students from programs they already offer or by standing up new ones. Here, the evidence points to the latter. If the boom primarily expands existing programs, we would expect to see a rise in energy completions from pre-existing energy programs. However, looking only at completions in energy programs that predate the boom, I find no

Figure 10: Shale Effects on Completions in Existing Programs Over Time



Notes. This figure plots the Gardner (2021) two-stage event study estimates for completions in existing energy (a) and non-energy (b) programs. Estimates are in standard deviations of the pre-boom annual change in the outcome. Standard errors are bootstrapped and clustered at the play level.

detectable increase. From Table 6, the coefficient is −0.01 on a pre-period mean of 4.0, and its confidence interval excludes the bulk of the 5.5-completion headline effect, reaching 1.71 at its maximum, so existing programs cannot account for the rise. Figure 10a confirms that there is no discernible increase in completions from existing energy programs in the years following the boom and pre-trends are parallel.

On the other hand, energy program creation is positive on every margin I examine—the count of new energy programs, completions in new energy programs, and energy's share of new programs all carry positive point estimates—but each is imprecisely estimated, reflecting how rare program creation is in any single college-year (Tables 7 and 8). Moreover, completions in "new" programs only span the 3 years post program creation, limiting their overall effect. The event study results in Figure 11a show what looks like a gradual albeit noisy upward trend in completions from new energy programs. Read together, the precise null on existing programs and the positive but noisy new-program estimates indicate that colleges accommodate the boom primarily by building new capacity rather than by expanding the programs they already run, which is also supported by the lag in energy completions discussed above. The marginal energy graduate is largely the product of a program that would not have existed absent the boom, indicating that colleges meaningfully impact the way students respond to the shock.

Table 7: Effect of Shale Boom on New Programs

Outcome:	Energy New CIPs	Non-Energy New CIPs	Energy's Share New CIPs
Raw Outcome:			
Treat × Post	0.01	0.14	0.18
	(0.01)	(0.30)	(0.73)
Effect per SD:			
Treat × Post	0.01	0.01	0.02
	(0.04)	(0.03)	(0.06)
CC FE	Y	Y	Y
Play × Year FE	Y	Y	Y
Observations	32,566	32,566	22,700
Plays	17	17	17
Baseline Mean	0.07	3	2.20
SD of ΔOutcome	0.38	10	12

Notes. This table presents the Gardner (2021) two-stage difference-in-differences estimates from Equation 5 for new programs, where a program is new in the first year a college offers it. Columns 1 and 2 report the effect on the number of new energy and non-energy programs, respectively, and column 3 reports energy's share of new programs in percentage points. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are clustered at the play level in parentheses.

Table 8: Effect of Shale Boom on Completions in New Programs

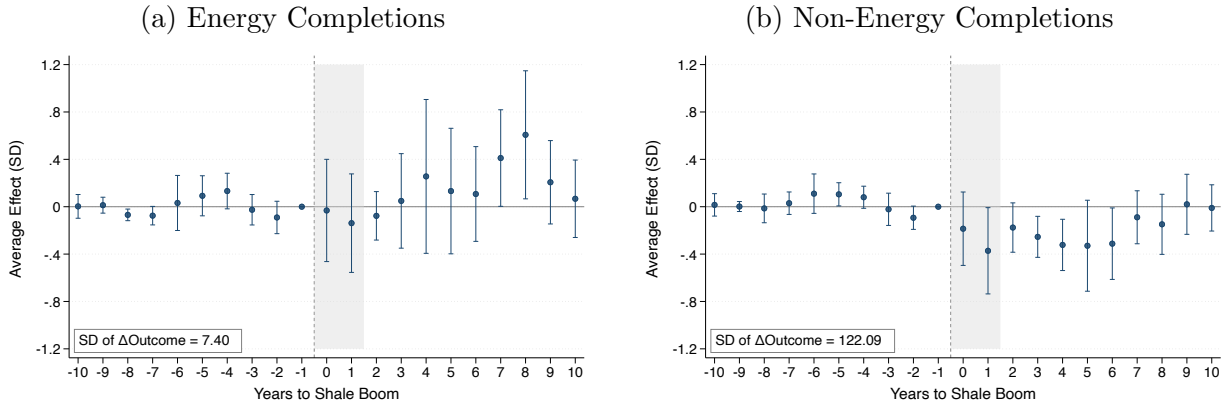
Outcome:	Completions in New Energy CIPs	Completions in New Non-Energy CIPs	Energy's Share Completions in New CIPs
Raw Outcome:			
Treat × Post	0.67	-9.96	0.47
	(0.51)	(7.79)	(0.66)
Effect per SD:			
Treat × Post	0.09	-0.08	0.07
	(0.07)	(0.06)	(0.11)
CC FE	Y	Y	Y
Play × Year FE	Y	Y	Y
Observations	32,566	32,566	28,509
Plays	17	17	17
Baseline Mean	1.10	57	1.50
SD of ΔOutcome	7.40	122	6.20

Notes. This table presents the Gardner (2021) two-stage difference-in-differences estimates from Equation 5 for completions in new programs, counted in the first three years a college offers each program. Columns 1 and 2 report the effect on completions in new energy and non-energy programs, respectively, and column 3 reports energy's share of completions in new programs in percentage points. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are clustered at the play level in parentheses.

4.5.2 Budgetary Implications

This response could operate through an expansion of resources, with colleges hiring additional staff or raising tuition to expand energy capacity, or through the reallocation of existing resources as colleges shift their specialization toward energy. To examine these margins of adjustment, I estimate effects on total expenditures per full-time-equivalent (FTE) enrollment, instructional expenditures per FTE, and gross tuition per FTE. I restrict this analysis

Figure 11: Shale Effects on Completions in New Programs Over Time



Notes. This figure plots the Gardner (2021) two-stage event study estimates for completions in new energy (a) and non-energy (b) programs. Estimates are in standard deviations of the pre-boom annual change in the outcome. Standard errors are bootstrapped and clustered at the play level.

Table 9: Effect of Shale Boom on Budgetary Measures

Outcome:	Expenditures per FTE	Instructional Expenditures per FTE	Gross Tuition per FTE
Raw Outcome:			
Treat × Post	-88	39	-68
	(669)	(306)	(107)
Effect per SD:			
Treat × Post	-0.011	0.012	-0.100
	(0.086)	(0.095)	(0.156)
CC FE	Y	Y	Y
Play × Year FE	Y	Y	Y
Observations	12,392	12,392	12,392
Plays	11	11	11
Baseline Mean	14,534	5,290	4,029
SD of ΔOutcome	7,792	3,222	684

Notes. This table presents the Gardner (2021) two-stage difference-in-differences estimates from Equation 5 for budgetary outcomes. Column 1 reports the estimates for total expenditures per FTE, column 2 for instructional expenditures per FTE, and column 3 for gross tuition per FTE. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are clustered at the play level in parentheses.

to the 11 plays whose booms began in 2007 or later since public institutions transitioned to new financial accounting standards between 2001-2003, making certain financial data comparable only from 2004 onward. I find no evidence of response on any of these margins, with no detectable effect for any of the outcomes (see Table 9). Although the null estimates cannot rule out modest changes in resources, they are consistent with a reallocation of existing resources toward energy programs. This interpretation also aligns with the opposing changes in energy and non-energy programs documented above. As discussed in Section 2.1, certain programs are more costly to establish and run than others, and new energy programs, which

require specialized equipment and instructors from a booming industry, are likely among the more costly to establish. This is also consistent with the asymmetry between the energy program gain and the non-energy program decline because expanding energy programs without additional resources per student would likely divert funds away from other programs.

4.6 Heterogeneity Across Institutions

Community colleges operate under tight budgets and differ vastly in their measured quality (Kurlaender et al., 2016) and their career and technical orientation, which is linked to their ability to produce credentials with high labor market value (Holzer et al., 2026). Even when colleges in the same play face a similar rise in labor demand, these differences may shape how nimbly they adapt. Estimates across and within plays suggest substantial heterogeneity in responsiveness. Why do some colleges expand energy training while others do not? I examine three broad inputs to the community college production function: spending, administrative quality, and instruction. I proxy for each with pre-boom characteristics available in my data. These proxies motivate hypotheses about the extent of adjustment, rather than provide direct measures of the underlying mechanisms.

First, as discussed in Section 2.1, colleges differ substantially in how much funding they have, where it comes from, and the mechanisms used to allocate it. I measure a college's resources using average pre-boom instructional expenditures per FTE student. Instructional spending is a common measure of college quality, including at community colleges (Stange, 2012; Cohodes and Goodman, 2014), and Deming and Walters (2017) find that higher spending per student raises degree completions at public colleges, but evidence specific to community colleges is mixed (Stange, 2012). Colleges with more instructional resources per student may be higher quality and have more room to staff and equip new programs. However, higher instructional spending can also reflect a more costly mix of programs, leaving less to redirect. Accordingly, the heterogeneity estimates provide suggestive evidence on the role of resources rather than a test of either channel.

Second, administrative capacity or quality may influence colleges' ability to implement programmatic changes. McCormack et al. (2014) find that management practices in UK university departments are associated with teaching and research performance, conditional on resources and past performance. Institutional leadership may also matter for securing outside sources of funding, coordination with local employers, or identifying emerging training needs. Administrator quality is difficult to measure, so I use two proxies. One is an indicator for whether a college received any NSF grants in the pre-boom period (1990–1995). Winning

these competitive grants requires identifying training requirements, writing proposals, and managing outside funds, the same capabilities necessary to build new programs. The other is the pre-boom average Executive, Administrative, and Managerial (EAM) salary normalized by local per capita income to strip out the effect of high income areas. Existing research suggests that compensation reflects the ability to attract high quality workers in public sector employment (Dal Bó et al., 2013; Gagliarducci and Nannicini, 2013). If this applies in the community college setting, higher EAM compensation may be associated with a stronger community college response. However, neither NSF grant receipt nor compensation is a direct measure of administrator quality or experience, so this analysis cannot isolate their contributions to adjustment.

Third, to launch a new technical program, a college must hire instructors from the very occupations whose labor demand is rising during the boom. When local private-sector wages rise, colleges whose pay is not competitive in their local labor market may be unable to recruit or retain them. Evidence from K–12 education shows that teacher pay relative to local outside wages shapes who teaches and how well schools perform (Hoxby and Leigh, 2004; Falch, 2010). The Texas shale boom itself widened the gap between teacher and private-sector wages, increasing teacher turnover and the share of inexperienced teachers (Marchand and Weber, 2020). Similar constraints bind in postsecondary programs where nursing schools turn away tens of thousands of qualified applications each year due to a lack of faculty and other capacity (AACN, 2026). Interviews with community college program leaders identify difficulties attracting and retaining clinical instructors as a constraint on program expansion (Murnane et al., 2024). I therefore approximate a college's ability to attract qualified instructors with its pre-period average faculty salary relative to county per capita income. I expect a stronger response where instructor compensation is more competitive with local incomes. However, higher average salaries could also reflect a more senior, tenured faculty, which may make them more difficult to reassign or replace and slow a shift toward new fields. I do not control for tenure directly because more than 40% of public two-year colleges have no tenure system (NCES, 2023), so doing so would drop a large share of the sample. Since faculty salary is fixed before the boom, it cannot capture the rise in wages that the boom itself causes, so it only partly reflects recruitment constraints.

Although I cannot causally identify the mechanisms, I test these predictions in a single pooled specification that augments Equation (5):

$$\begin{aligned} y_{i,p,t} = \alpha_{i,p} + \gamma_{p,t} + \beta\,(\text{Treat}_{i,p} \times \text{Post}_{p,t}) + (\text{Treat}_{i,p} \times \text{Post}_{p,t})\,\mathbf{C}_i'\boldsymbol{\delta} + (\text{Post}_{p,t})\,\mathbf{C}_i'\boldsymbol{\zeta} \\ + \phi\,(\text{Treat}_{i,p} \times \text{Post}_{p,t})\,\text{Enroll}_i + \psi\,\text{Post}_{p,t}\,\text{Enroll}_i + \varepsilon_{i,p,t}, \end{aligned} \tag{10}$$

Table 10: Heterogeneity in Shale Response by Pre-Boom Characteristics

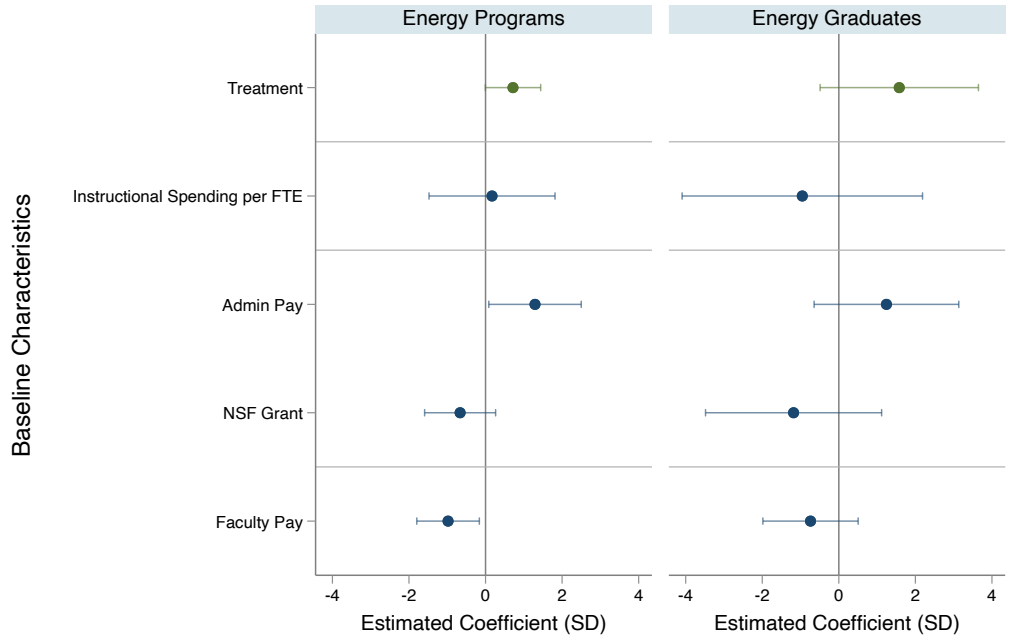
Outcome:	Energy CIPs	Energy Completions	Total Revenue	Enrollment (FTE)
Treat × Post	0.72*	1.58	-1.79***	-1.05**
	(0.33)	(0.95)	(0.52)	(0.38)
× Instructional Spending per FTE	0.17	-0.95	-3.25***	-0.03
	(0.76)	(1.44)	(0.72)	(0.61)
× Admin Pay	1.29**	1.24	0.57	0.67*
	(0.55)	(0.87)	(0.42)	(0.32)
× NSF Grant	-0.66	-1.18	1.46**	1.23**
	(0.43)	(1.06)	(0.58)	(0.50)
× Faculty Pay	-0.98**	-0.74	-0.63*	-0.53
	(0.38)	(0.57)	(0.34)	(0.30)
CC FE	Y	Y	Y	Y
Play × Year FE	Y	Y	Y	Y
Stacking Weights	Y	Y	Y	Y
Observations	15,187	15,187	12,696	15,177
Plays	13	13	13	13
SD of ΔOutcome	0.45	5.50	3,738,764	364

Notes. Estimates of Equation (10). Outcomes are expressed in standard deviations of their annual change among pre-period control colleges. The first row reports Treat × Post; each remaining row reports the coefficient on (Treat × Post) × C_i for the standardized pre-boom characteristic C_i in the row label (NSF Grant is an indicator). All columns control for pre-boom enrollment interacted with Treat × Post and with Post, include the Post × C_i terms, which are not displayed, and use stacking weights. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the play level in parentheses.

where $\mathbf{C}_i$ is the vector of standardized baseline characteristics for institution i, including instructional spending per FTE, NSF grant receipt, administrative pay, and faculty pay. Enroll_i is standardized pre-boom FTE enrollment, which I include as a separate control, interacted with both $\text{Treat}_{i,p} \times \text{Post}_{p,t}$ and $\text{Post}_{p,t}$, so that each characteristic is compared across colleges of similar size. The elements of $\boldsymbol{\delta}$ give each characteristic's marginal contribution to the energy-program response, holding the others fixed. The $\text{Post}_{p,t} \times \mathbf{C}_i$ terms allow colleges with different characteristics to follow different paths after the boom regardless of treatment. I estimate Equation (10) for energy programs and energy graduates, and, to examine whether the characteristics operate through resources or student demand, for total revenue and full-time-equivalent (FTE) enrollment. The sample is restricted to institutions that report these characteristics to IPEDS and to plays that boom after 2003. It includes 70 treated and 249 control colleges across 13 plays (69 in the total revenue specification).

Table 10 and Figure 12 report the results from Equation 10. Administrative pay is the strongest predictor of the energy response. A one-standard-deviation increase in pre-boom administrative pay relative to local income raises the program effect by 1.3 standard deviations of the annual change (0.55). The graduate effect is similar in size, 1.2 SD (0.87), but noisy. Colleges with higher faculty pay adjust less, −1.0 SD (0.38), contrary to the instructor-hiring hypothesis. Administrative and faculty pay are correlated across colleges (0.67), so the coefficients are identified from colleges where the two diverge. Much of the

Figure 12: Average Shale Effects by Pre-Boom Characteristic

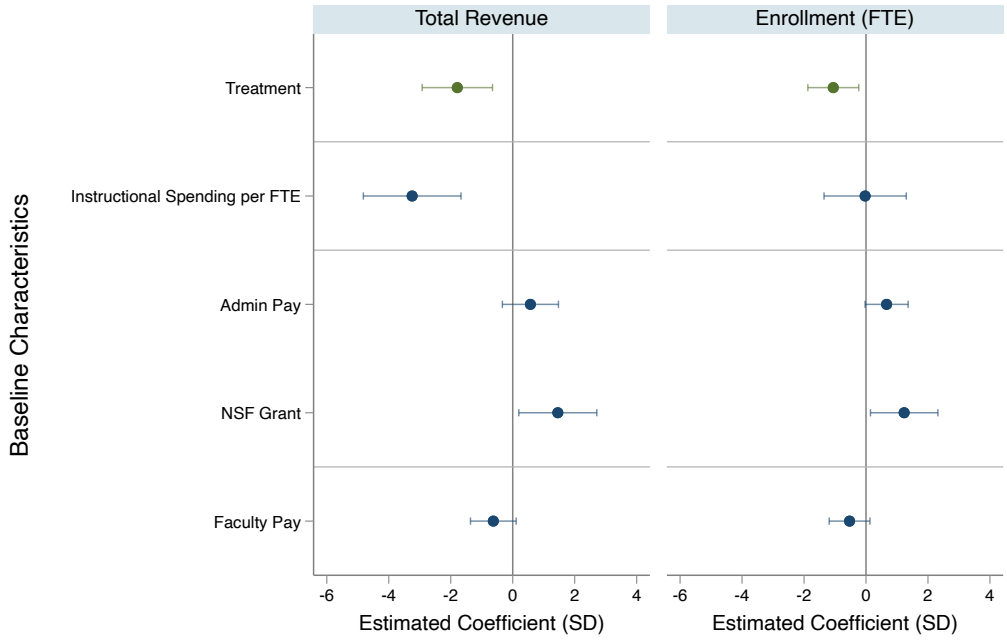


Notes. This figure plots the coefficient on Treat × Post and on (Treat × Post) × C_i for each pre-boom characteristic C_i from Equation (10), with 95% confidence intervals. Characteristics are grouped as instructional spending, administrative quality (administrative pay and NSF grants), and faculty pay. Outcomes are expressed in standard deviations of their annual change. Standard errors are clustered at the play level.

variation in faculty salaries occurs across states, so this estimate may reflect state policies rather than a college's ability to recruit or replace instructors. Higher salaries may also reflect a more senior or tenured faculty, making replacement more difficult, and they make each additional instructor more expensive for a budget-constrained institution. Neither instructional spending per FTE nor prior NSF grants predicts the program or graduate response.

The revenue and enrollment outcomes reported in Figure 13 help separate these channels. Administrative pay has no detectable relationship with revenue growth, so its link to program expansion does not appear to run through additional resources and may instead reflect administrative quality. Its association with enrollment is positive but only marginally significant. Prior NSF grants predict larger gains in revenue and enrollment, but not in programs, consistent with grant-seeking capacity bringing in resources rather than new offerings. Colleges with higher faculty pay see slower revenue growth during the boom. Because a new program carries fixed start-up costs in instructors and equipment, slower revenue growth combined with more expensive instructors may leave these colleges less room to expand,

Figure 13: Average Shale Effects on Revenue and Enrollment by Pre-Boom Characteristic



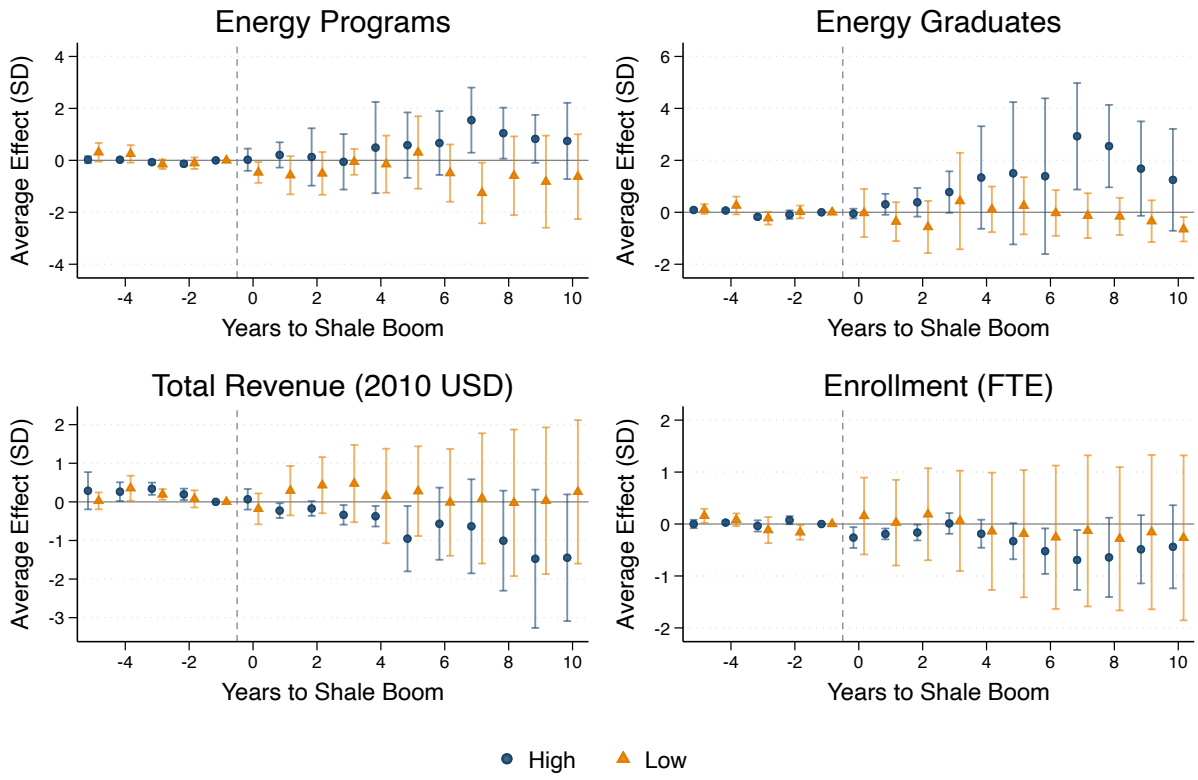
Notes. This figure plots the coefficient on Treat × Post and on (Treat × Post) × C_i for each pre-boom characteristic C_i from Equation (10), with 95% confidence intervals, for total revenue and FTE enrollment. Characteristics are grouped as instructional spending, administrative quality (administrative pay and NSF grants), and faculty pay. Outcomes are expressed in standard deviations of their annual change. Standard errors are clustered at the play level.

which is more consistent with a cost interpretation of the faculty pay coefficient than with a recruitment constraint.

To examine timing, I estimate Equation (7) separately for colleges in the top and bottom quartiles of administrative pay (Figure 14). The two groups follow similar paths in energy programs and graduates prior to the boom. Colleges with high administrative pay add energy programs after the boom, by about one standard deviation within seven years, and their energy graduates rise by two to three standard deviations over the same period, while colleges with low administrative pay show no increase in either. Revenue at high administrative pay colleges does not rise and, if anything, falls relative to low-pay colleges, which reinforces that their expansion is not financed by additional resources.

Event studies for the other three characteristics are reported in Appendix Figures A6–A8 and are noisier. Considered alone, colleges with high faculty pay do not add fewer energy programs than colleges with low faculty pay, although their revenue and enrollment decline after the boom, so the negative pooled coefficient on faculty pay appears only after

Figure 14: Shale Effects Over Time by Pre-Boom Administrative Pay

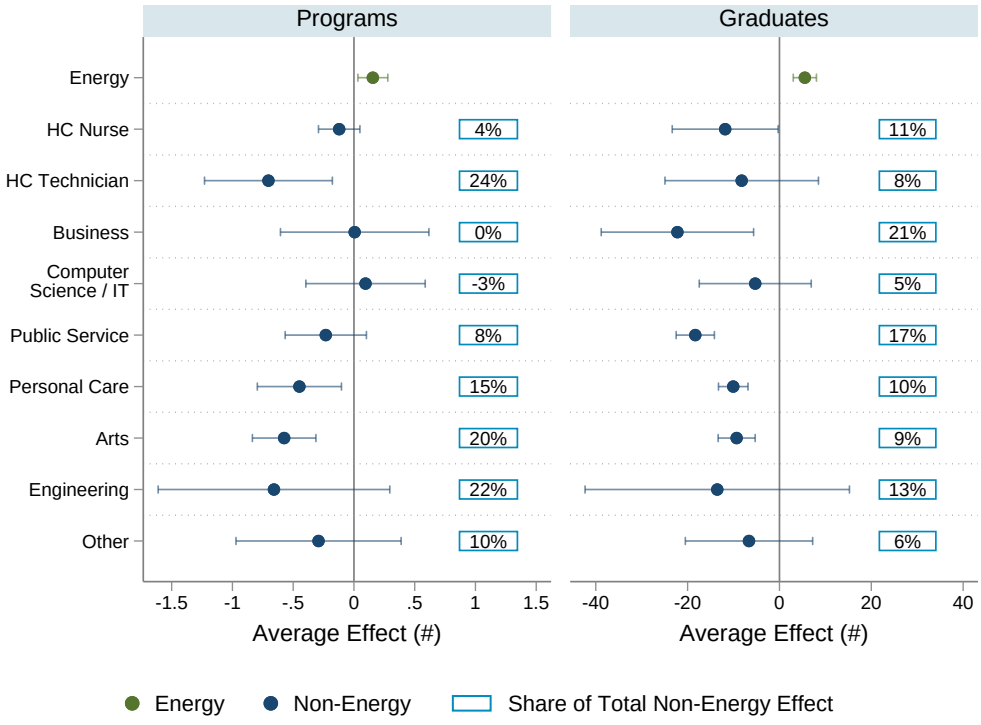


Notes. This figure plots Gardner (2021) two-stage event study estimates separately for colleges in the top quartile (High) and bottom quartile (Low) of pre-boom administrative pay relative to county per capita income, for energy programs, energy graduates, total revenue, and FTE enrollment. Colleges in the middle two quartiles are excluded. Year effects vary with pre-boom enrollment. Outcomes are expressed in standard deviations of their annual change. Standard errors are clustered at the play level.

conditioning on the other metrics. Colleges with high instructional spending add energy programs late in the window, but do so at the same time as low-spending colleges, and the two groups cannot be distinguished. Colleges with prior NSF grants add fewer energy programs in the first years of the boom but catch up within roughly six years.

Since a college's pay structure and spending may themselves reflect local conditions or other omitted factors, these are associations rather than causal relationships, but their direction is telling. They show that responsiveness is not orthogonal to observable characteristics. Still, much of the variation across colleges remains unexplained even after these characteristics are accounted for, and the rest likely lies in features of these institutions that are harder to measure, such as leadership and relationships with local employers, which remains a question for future research.

Figure 15: Shale Effects on Non-Energy Programs and Completions by Field



Notes. This figure plots the Gardner (2021) two-stage difference-in-differences estimates of Equation (5) by field for non-energy programs (left) and completions (right). Boxes report each field's share of the total non-energy effect of −2.93 programs and −105 completions. Shares sum to 100, and negative shares indicate fields that move in the opposite direction to the overall decline. Bars show 95% confidence intervals. Standard errors are clustered at the play level.

4.7 Decomposition of Non-Energy Effects

The negative effect on non-energy programs raises questions about the programs that community college entrants in treated areas may be losing access to and whether the associated reduction in completions is driven by those institutional choices. The aggregate non-energy slowdown could reflect a broad, shallow effect spread across many fields or a sharp one concentrated in a few. To distinguish between these possibilities, I re-estimate Equation (5) separately for a set of program-industry groupings[18] of non-energy programs.

Figure 15 plots the coefficients for each industry grouping for comparison. The slowdown across programs is relatively small and spread across a variety of fields. The most significant

[18]Programs are generally classified into industries in two forms: First, if the CIP code clearly denotes a specific industrial focus, for example, education or health care, all programs in that classification are assigned to that industry group. Second, if this is not possible, as was the case with energy, then a threshold based on the mapping described in Section 3.2 is used for classification.

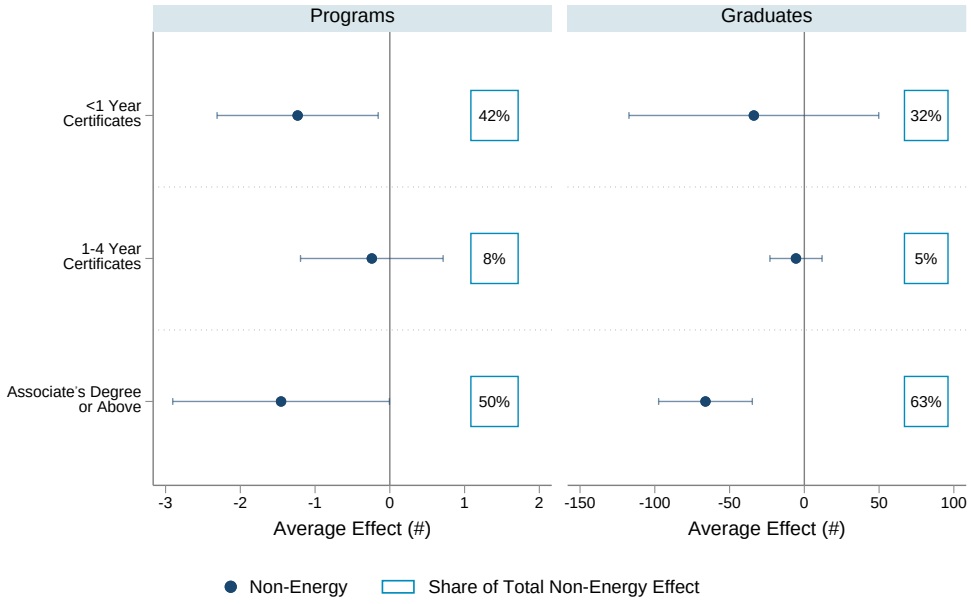
shortfalls are in healthcare technicians at 24% of the total effect on non-energy programs, the arts at 20%, and personal care at 15%. Engineering also represents about 20% of the effect, but the estimate is noisy. Measured in completions, the declines are broad and largely track program changes. Certain fields exhibit outsized responses compared to program effects, for example, business, which accounts for 21% of the decline in graduates compared to no change in programs, and public service, which represents 17% of the total decline, roughly double the program response. Across the remaining fields, contractions in completions are fairly even, though with varying degrees of precision.

The results imply that these are either the credentials that students are most likely to forgo as the opportunity cost of education rises, or that they compete most directly with energy programs for students and instructional resources. This pattern echoes the headline results, indicating that students are more responsive to labor market demand than colleges. On the non-energy side, the shortfall often operates through completion shifts within a stable set of offerings rather than through program closure or forgone creation. Nursing in particular stands out since it is one of the highest return vocational tracks that community colleges offer. The sensitivity of high-return programs and completions to the boom calls into question whether or not the specialization induced by the boom leaves potential students better off.

As discussed in Section 2.1, program churn differs across award levels, and is likely tied to the start-up cost of each program. If the boom slows the growth of non-energy offerings, either through fewer new programs or more program closures, this may be reflected in the programs that turn over most often and cost the least to stop and restart. Alternatively, if colleges require funds to start up or expand costly energy programs, they may instead scale back other costly programs. I examine the effect of the boom on the distribution of non-energy programs across credential types, grouping them into short certificates of less than one year, long certificates of one to four years, and associate degrees and above, which are almost entirely associate degrees.

Figure 16 depicts the estimates of Equation (5) for each group. Associate degrees account for half of the decline in programs (-1.45) and 63% of the decline in completions (-66), roughly in line with their approximately 55% share of non-energy program-years and completions in the sample. However, this is slightly outsized compared to the expected churn from an associate's degree. These credentials exhibit the least churn, and weighting each award level's program-years by its churn rate, associate degrees would be expected to account for about 42% of program turnover. Short certificates, on the other hand, are

Figure 16: Shale Effects on Non-Energy Programs and Completions by Award Level

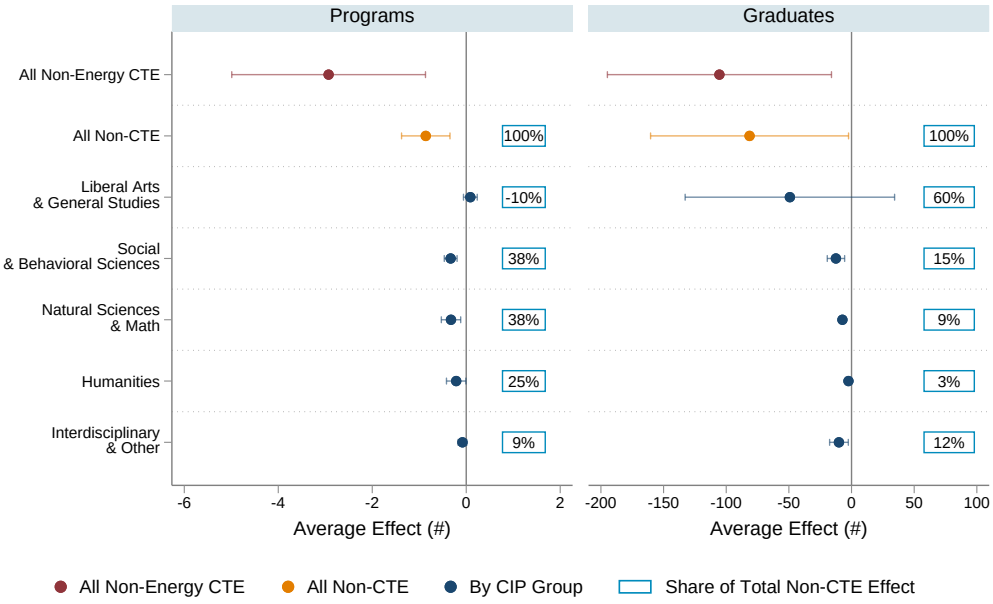


Notes. This figure plots the Gardner (2021) two-stage difference-in-differences estimates of Equation (5) by award level for non-energy career and technical programs (left) and completions (right). Short certificates are awards of less than one academic year, and long certificates are awards of one to four years. Associate and above includes associate degrees and the few bachelor's and graduate awards at these colleges. Boxes report each award level's share of the total non-energy effect of −2.93 programs and −105 completions, and the shares sum to 100. Bars show 95% confidence intervals. Standard errors are clustered at the play level.

overrepresented in terms of expected turnover and their presence in the sample. They make up 16% of non-energy program-years and based on churn would be expected to contribute about 20% of the decline, but they account for 42% of the decline in programs (−1.23). Their share of the completions decline (32%) is imprecise and largely in line with their overall share of completions in the sample. Long certificates contribute little, 8% of the program effect and 5% of the completions effect despite accounting for 29% of program-years and 20% of completions, and neither estimate is distinguishable from zero. Based on their churn, they would be expected to account for about 38% of program turnover.

Across award levels, the non-energy decline in programs is concentrated at both ends of the cost distribution. Short certificates, the cheapest programs to start, fall about twice as much as their normal turnover would predict, and associate degrees, the costliest, fall somewhat more than predicted, while long certificates barely change. This pattern is consistent with colleges slowing the growth of their lowest cost programs while also scaling back expensive ones as their energy portfolios expand. It also mirrors the results by field, with declines in several low-cost, classroom-based fields as well as higher cost, healthcare fields

Figure 17: Shale Effects on Non-CTE Programs and Completions by CIP Family



Notes. This figure plots the Gardner (2021) two-stage difference-in-differences estimates of Equation (5) for academic (non-CTE) programs (left) and completions (right), in total and by two-digit CIP family. The top row shows the total effect on non-energy career and technical programs (−2.93) and completions (−105) for comparison. Liberal arts and general studies covers CIP 24, social and behavioral sciences CIP 05, 42, and 45, natural sciences and mathematics CIP 26, 27, and 40, and humanities CIP 16, 23, 38, 39, and 54. Interdisciplinary and other covers the remaining families. Boxes report each group's share of the total non-CTE effect of −0.86 programs and −81 completions, and the shares sum to 100. Non-CTE outcomes cover the same colleges and years as the non-energy estimates, with zeros for college-years without non-CTE programs. Bars show 95% confidence intervals. Standard errors are clustered at the play level.

that require specialized equipment and training. The loss of graduates is broader and runs mainly through associate degrees with considerable variation in short certificates.

The decomposition thus far only includes career and technical programs; however, academic programs draw on the same budgets and instructors, so I also explore how the boom affects them. These credentials are predominantly associate's degrees (85-90% of program-years and completions in the sample), which typically serve as the transfer track to 4-year institutions. I group the non-CTE programs by two-digit CIP family into liberal arts and general studies, social and behavioral sciences, natural sciences and mathematics, humanities, and other fields. Figure 17 reports the results of Equation 5 on each category. Exposed colleges offer 0.86 (0.26) fewer non-CTE programs (0.53 SD, s.e. 0.16) and award 81 (40) fewer non-CTE completions (2.1 SD, s.e. 1.0). The relative reduction in graduates compared to non-exposed colleges is of similar size to the 105 in non-energy career and technical fields, but the effect on non-CTE programs is smaller. The program gap is concentrated in the social and behavioral sciences (38%), natural sciences and mathematics (38%), and humanities

(25%), while liberal arts programs show no relative change. On the other hand, liberal arts account for the vast majority of the shortfall in completions (60%), although the estimate is noisy. The remaining fields each show smaller declines across completions. Event studies show non-CTE programs and completions at exposed colleges declining relative to controls in the five years prior to the boom, so these estimates likely overstate non-CTE effects. Taken at face value, the boom draws students away from the academic transfer track as well as from other vocational fields.

5 Solar Investment and Community College Response

The shale boom is just one of many different types of energy transformations, and whether community colleges respond the same way to the clean-energy shocks is a separate and highly topical question. Rapid buildout of data centers and growth in artificial intelligence are driving a structural shift in the need for electricity. Meeting the demand as well as decarbonization goals will require sustained expansion of electrical generation capacity, particularly from renewables. The U.S. Energy Information Administration (EIA) forecasts 50% growth in total electricity consumption over the next 25 years with generation from renewables alone expected to increase by nearly 200%, accounting for almost 90% of total consumption by 2050 (EIA, 2025a). Yet, persistent shortages of key skilled workers, for example electricians, welders, solar installers, and wind technicians, are delaying projects across the U.S. and the world (International Energy Agency, 2025; U.S. Department of Energy, 2025). The constraint to growth is not only generation capacity, but also labor. In order to provide evidence on the applicability of the shale findings to this setting, I apply an analogous design to the recent expansion of renewable-energy investment, focusing on solar.

The solar shock differs from the shale boom in three ways that may shape how colleges respond (Section 2.3). First, solar generates less persistent labor demand. Most of that demand is in installation and project development (IREC, 2024), so local hiring falls once construction is over. Second, solar demand may be less geographically concentrated. Distributed solar follows population, and installers can travel across local labor markets, so the demand for trained workers may spread beyond the places where capacity is built. And, third, solar-specific instruction builds on existing electrical and construction skills (BLS, 2026c). Colleges can add it within existing programs, and apprenticeships with employers and unions both offer another route and introduce additional constraints. These features suggest that colleges may respond less to solar than to shale, at least on the margins of program creation and graduates. I therefore examine programs and completions nationally,

as well as at the level of course sections in California, where I can observe finer adjustments.

5.1 Empirical Design

Renewable investment is not exogenous and it flows to places with favorable resources, supportive policy, and, potentially, a workforce already equipped to build and maintain it. I therefore instrument local solar investment with an exposure instrument somewhat analogous to the shale shock approach. Following Chan and Zhou (2025), I interact a fixed and plausibly exogenous measure of local solar "potential", which determines the location of the solar shock, with aggregate shifters unrelated to any single college's program decisions, which determine the timing.

$$Z_{c,s,t} = \text{SolarPotential}_c \times \text{Shock}_{s,t-1} \tag{11}$$

where local solar potential is the product of average solar irradiance (Figure 3a) and land area in commuting zone c, essentially a measure of the area's solar resource usable for generating electricity, and the shifters are Renewable Portfolio Standard (RPS) targets in state s, the portion of a state's electricity that is required to come from renewable generation, and the global price of solar photovoltaic modules per watt (Figure 4a), each lagged one year relative to investment. The first stage is:

$$\text{SolarInvestment}_{c,t} = \pi\, Z_{c,s,t} + \gamma\, population_{c,t} + \delta_t + \theta_s + \nu_{c,s,t}, \tag{12}$$

where solar investment is the annual megawatts of solar capacity installed in commuting zone c in year t and the specification includes population controls, and year and state fixed effects. The second stage relates community-college outcomes to the predicted investment:

$$\Delta Y_{i,c,s,t} = \beta\, \widehat{\text{SolarInvestment}}_{c,t-2} + \delta_t + \theta_s + \lambda\, population_{c,t} + \varepsilon_{i,c,s,t}, \tag{13}$$

where ΔY is the change in solar programs or completions from the IPEDS data at institution i in year t. Predicted solar investment enters with a lag of 2 years, and the specification controls for population at the commuting zone c level and δ_t year and θ_s state fixed effects. Standard errors are clustered at the commuting zone level. The national sample covers 1,165 community colleges in 478 commuting zones across 47 states, with outcomes observed from 2002 to 2024 and lagged investment from 2000 to 2022.

In order to explore more detailed effects on course-level offerings, I also implement a similar design within California. The first stage is the same as above but with the shifter

limited to the California RPS targets and at the county j level, and the second stage is:

$$\Delta Y_{i,j,s,t} = \beta \, \widehat{\text{SolarInvestment}}_{j,t-1} + \delta_t + \theta_{z(j)} + \lambda \, population_{j,t} + \varepsilon_{i,j,s,t}, \tag{14}$$

where ΔY is the change in solar sections from the CCCCO data at institution i in year t. Predicted solar investment enters with a lag of 1 year, and the specification controls for population at the county j level and δ_t year and $\theta_{z(j)}$ commuting zone fixed effects. Standard errors are clustered at the county level. The California sample includes 113 colleges in 38 counties, observed from 2002 to 2024.

Identification. I follow Chan and Zhou (2025)'s strategy, which is similar to a shift share or exposure design (Bartik, 1991; Autor et al., 2013; Allcott and Keniston, 2018; Goldsmith-Pinkham et al., 2020). However, there is only one continuous exposure[19], solar potential, which varies across places, and is interacted with aggregate shifters, rather than sub-group variation within a local area as is the case in the canonical shift-share designs. I rely on exogeneity of exposure rather than the shock for identification where, as in Chan and Zhou (2025), local solar potential is taken to be as good as randomly assigned conditional on the controls. Solar potential must also affect community-college solar outcomes only through local solar investment and not through any other channels. The identifying variation is the differential response of high- versus low-potential commuting zones, or counties, to the common, time-varying changes in RPS and module prices.

Let $D_{c,t}(z)$ denote potential solar investment in commuting zone c and year t given instruments of value z; in other words, the solar investment commuting zone c would receive in year t if its instruments (solar potential interacted with policy and price shifters) took value z. Similarly, let $\Delta Y_{i,t}(d, z)$ be the change in solar outcomes at college i in year t if two years earlier, at $t-2$, its commuting zone, $c(i)$, had received solar investment d and the instruments to predict that investment took the value z. The 2SLS estimate has a causal interpretation under the four standard IV assumptions with heterogeneous effects.

Assumption IV1 (Relevance). The instruments shift local solar investment, meaning the instruments are correlated with observed solar investment. Conditional on population and

[19] A standard Bartik instrument aggregates shocks across industries or other groups, weighted by each location's initial shares of those groups (Bartik, 1991). Here, as in Chan and Zhou (2025), a single exposure, solar potential, is interacted with state or national shifters. Following Goldsmith-Pinkham et al. (2020), identification rests on the conditional exogeneity of this exposure; that is, solar potential must not predict differential changes in college outcomes through channels other than solar investment, conditional on controls and fixed effects. An alternative approach bases identification on as-good-as-random shifters (Borusyak et al., 2022), with recentering to account for nonrandom exposure (Borusyak and Hull, 2023). I do not take this approach because RPS targets may covary with other state policies and module price declines may coincide with broader economic trends, so the shifters are unlikely to be as good as random.

the fixed effects in Equation 12, at least one coefficient in π is not zero.

Assumption IV2 (Independence). Conditional on population and the fixed effects, solar potential is as good as randomly assigned with respect to potential investment and potential outcomes:

$$\Big\{D_{c(i),t-2}(z),\ \Delta Y_{i,t}(d,z)\Big\}_{d,z} \perp\!\!\!\perp \text{SolarPotential}_{c(i)}$$

In other words, if high- and low-potential places had the same solar potential, they would on average have the same solar investment and college outcomes.

Assumption IV3 (Exclusion). Solar potential affects solar training only through solar investment,

$$\Delta Y_{i,t}(d,z) = \Delta Y_{i,t}(d) \quad \text{for all } d \text{ and } z$$

Put simply, holding investment fixed, a college's solar offering and completions do not depend on the instrument, z.

Assumption IV4 (Monotonicity). Annual solar investment is weakly increasing in local solar potential, holding fixed the policy and price shifters as well as any other determinants of investment. That is, for any commuting zone and investment year, greater solar potential cannot reduce solar investment, but the strength of this response may vary across locations and years. For all $\text{SolarPotential}^H \geq \text{SolarPotential}^L$:

$$D_{c,t}\Big(\text{SolarPotential}^H \times \text{Shock}_{s,t-1}\Big) \ \geq\ D_{c,t}\Big(\text{SolarPotential}^L \times \text{Shock}_{s,t-1}\Big)$$

Under these assumptions, β is a weighted average effect of solar investment per gigawatt in the places where investment responds to the instruments (Imbens and Angrist, 1994; Angrist et al., 2000). With two instruments, monotonicity alone does not ensure that the weights are positive, and negative weights may cause estimates to fall outside the range of effects across places or be incorrectly signed (Goldsmith-Pinkham et al., 2020). The national estimate, β from Equation 13, is a weighted average of the estimates from each instrument, so I compute these individual weights following Goldsmith-Pinkham et al. (2020). Both weights are positive, at 0.79 on the RPS instrument and 0.21 on the price instrument. Therefore, the four assumptions above are sufficient to recover the average causal response of solar offerings and completions to one additional gigawatt of investment. This is the continuous-treatment analogue of a local average treatment effect (LATE) (Angrist and Imbens, 1995).

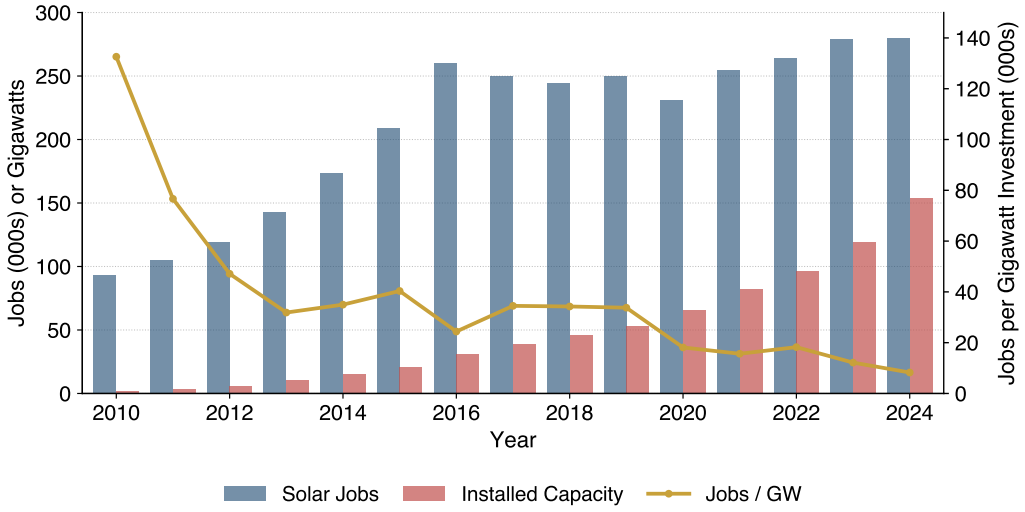
Relevance is the only directly testable assumption, and Figure 19 shows a strong first stage for the national sample; however, within California, the first stage is weaker (Figure 21). Monotonicity cannot be tested directly, but it only requires that higher solar potential not

reduce solar investment, holding all else constant. Holding electricity prices, installation costs, transmission access, and policy fixed, stronger sunlight and more land make solar installation more profitable. Developers therefore have no incentive to invest less where potential is higher.

Turning to Assumption IV2, since I rely on exogeneity of exposure, independence rests on solar potential (Goldsmith-Pinkham et al., 2020; Chan and Zhou, 2025). One potential threat to this assumption is that the RPS shifter varies across states, so concurrent changes in a state, such as new funding for clean-energy training, would bias the estimates. Since I estimate Equation 13 in first differences, the state fixed effects remove a linear trend in each state, but not changes that depart from that trend. I estimate Equation 13 using only national shifters, including module prices and the federal ITC, as well as including state by year fixed effects. As reported in Appendices E.5 and E.6, both approaches produce results consistent with baseline estimates. Another more direct threat is that solar potential is highest in places on different economic trends, for example the Sun Belt, whose colleges may have expanded technical programs and graduates for other reasons. Appendix E.6 reports balance across baseline college and commuting-zone characteristics by solar potential. Although solar potential does not predict baseline solar programs or completions, within states, colleges in higher-potential commuting zones had faster growth in CTE programs, higher enrollment, and a larger Hispanic share before 2002. The effect on completions remains of similar magnitude and significant across all controls; however, the effect on solar programs weakens with the inclusion of Hispanic share. This could reflect other changes at colleges with large Hispanic enrollments, including targeted federal funding, regional economic cycles, or labor-market composition.

Concurrent clean-energy policies in high-potential places, other renewable investment, and responses to announced RPS targets before investment would each violate the exclusion restriction. State-level policies are absorbed by state-by-year fixed effects, and the results are similar using only national shifters (see Appendices E.5 and E.6). Solar and wind potential are negatively correlated (Chan and Zhou, 2025), so any training response to wind investment would, if anything, bias the estimates downward. In California, non-solar sections do not respond, which suggests solar potential does not affect training through broader local changes. Analogous assumptions apply to the California estimates, with counties in place of commuting zones and the single statewide RPS path absorbed by the year fixed effects. These checks cannot confirm the identifying assumptions directly, but they provide evidence against their most likely violations and support a causal interpretation of the solar estimates.

Figure 18: U.S. Solar Employment and Installed Solar Capacity



Notes. This figure plots total U.S. solar employment in thousands from the Solar Jobs Census (IREC, 2024) and installed solar capacity in gigawatts (AC), combining utility-scale capacity from EIA Form 860 and distributed capacity from LBNL's Tracking the Sun for the contiguous United States (EIA, 2025b; LBNL, 2025b). The line plots solar jobs per gigawatt of capacity added that year, in thousands (right axis).

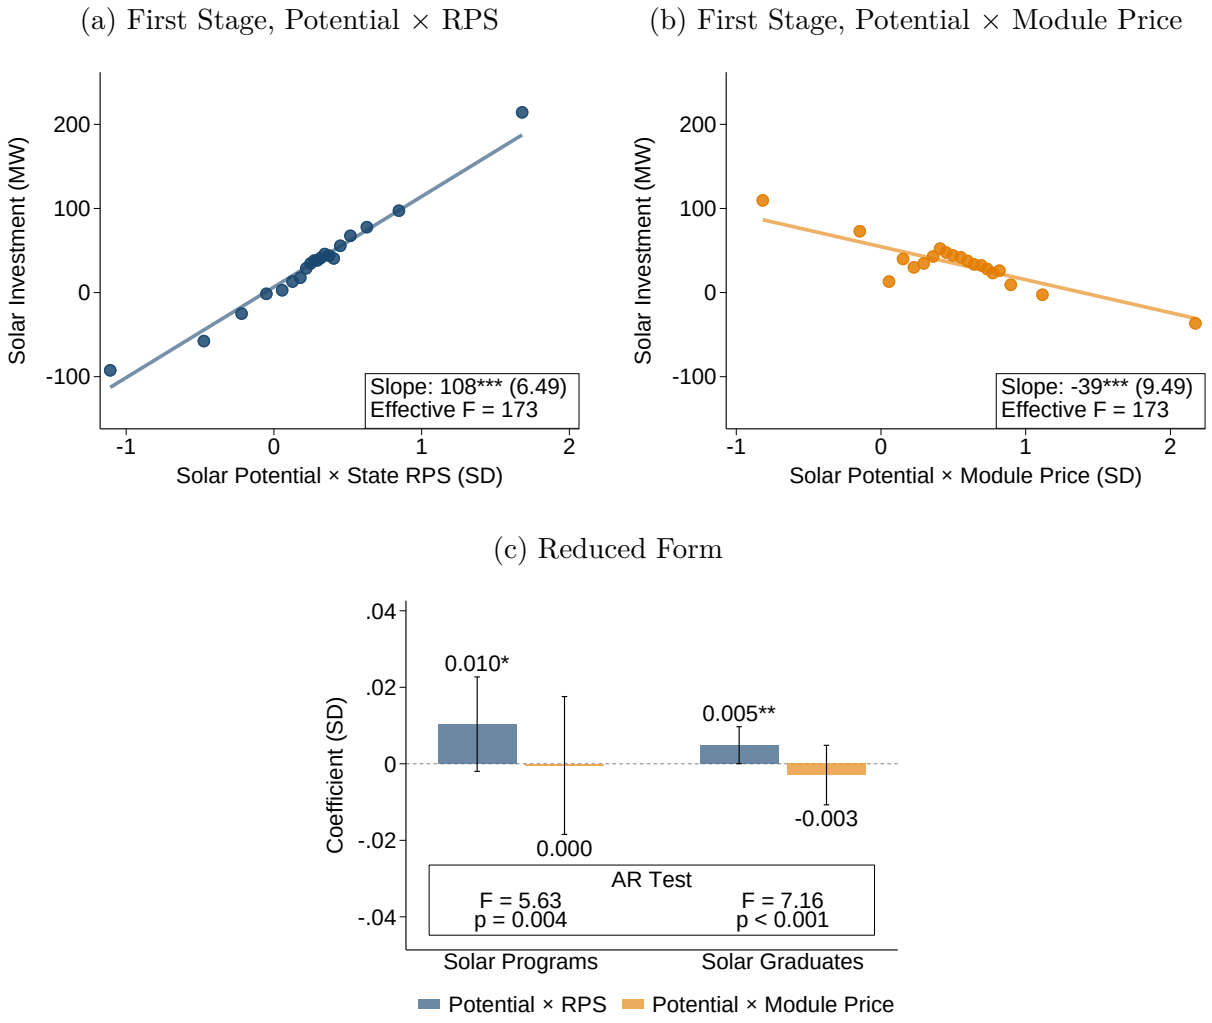
5.2 Labor Demand Shock

Before turning to community college outcomes, I characterize the local labor demand created by solar investment. The solar workforce is not separately identified in local employment data, so I rely on national statistics. Between 2010 and 2022, U.S. solar employment rose from about 94,000 to 264,000 (IREC, 2024), while installed solar capacity grew from about 2 to 97 gigawatts, or roughly 1,800 additional jobs per additional gigawatt (Figure 18). At this ratio, the roughly 0.7 gigawatts of capacity added in the average community college's commuting zone corresponds to about 1,270 jobs. However, additions are highly skewed across local areas, with a median of 0.1 gigawatts. Because most solar jobs are in installation and development (Section 2.3), employment follows the pace of new installations rather than the stock of capacity.

5.3 National Effects on Programs and Completions

The solar instrument is strong in the national sample, where local solar potential interacted with state RPS targets and the global solar module price closely tracks actual solar investment with an effective F-statistic of 173. Solar training also responds directly to the instruments. Solar programs and completions grow faster in commuting zones with high

Figure 19: First Stage and Reduced Form for National Solar Investment



Notes. This figure plots the first stage (a and b) and the reduced form (c) for the national sample. Panels (a) and (b) are binned scatter plots of solar investment against each instrument, after partialling out the other instrument, population, and year and state fixed effects. Both panels also report the effective F-statistic across both instruments (Montiel Olea and Pflueger, 2013). Panel (c) plots the reduced-form coefficient on each instrument, in standard deviations of the annual change in solar programs and completions, with the joint F-statistic and Anderson–Rubin p-value for each outcome. Bars show 95% confidence intervals. Standard errors are clustered at the commuting-zone level. Appendix Table B4 reports the full estimates.

solar potential when their state's RPS targets are higher and solar panels are cheaper. In the reduced form, the instruments are jointly significant for both solar programs and completions (Anderson–Rubin p-values of 0.004 and 0.001). Much of this response comes through the RPS instrument. The two-stage least-squares estimates in Table 11 show that solar investment raises community-college solar offerings. As in the shale analysis, I report effects in standard deviations (SD) of the annual change in the outcome. A one-gigawatt increase in predicted local solar investment increases solar programs by 0.078 SD (0.025), about 10%

Table 11: Effect of Solar Investment on Community-College Solar Outcomes

Outcome:	Δ Solar CIPs	Δ Solar Completions
Raw Outcome:		
$\widehat{Solar Investment}_{t-2}$ (GW)	0.02***	0.14***
	(0.01)	(0.04)
Effect per SD:		
$\widehat{Solar Investment}_{t-2}$ (GW)	0.08***	0.05***
	(0.03)	(0.01)
Year FE	Y	Y
State FE	Y	Y
Population Control	Y	Y
Wald F-Stat	1144	1144
Effective F	173	173
Observations	24,268	24,268
Sample Mean	0.18	0.65
SD of ΔOutcome	0.24	2.72

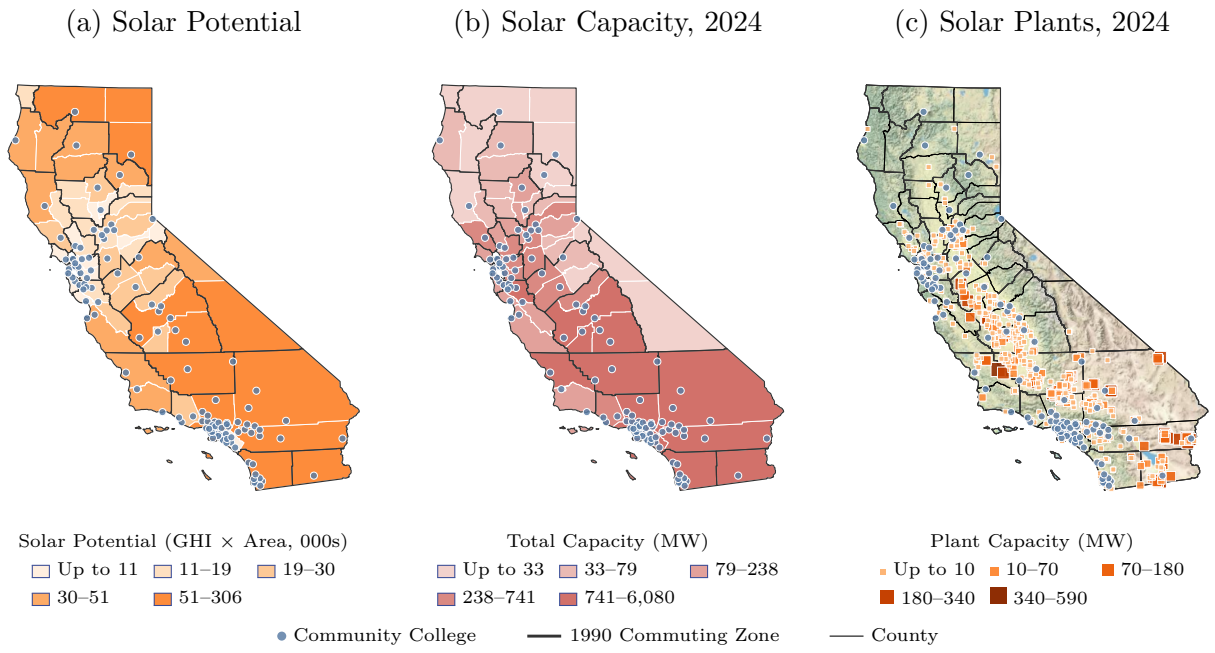
Notes. Column 1 reports the 2SLS effect on the count of solar programs offered; column 2 reports the effect on solar completions. The endogenous regressor is solar investment (instrumented as in the first stage); the lag is two years. The table reports effects on the raw outcome and per standard deviation of the annual change in the outcome. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the commuting-zone level in parentheses.

relative to the sample mean, and solar completions by 0.051 SD (0.015), or about 21%. To provide context, one gigawatt is a large increase in solar investment. It provides enough power for around 200,000 homes annually and is equivalent to about one third of the annual solar capacity installed in the entire state of California from 2015–2024. The average college's commuting zone added 0.7 gigawatts of solar capacity over the sample period, with most of it after 2012. Summed across the 1,165 colleges, the estimates imply an average of about 4 additional solar programs on offer and 29 additional solar graduates each year.

Thus, this provides evidence that the responsiveness documented for the shale boom is not specific to fossil fuels or to energy programs, as the new solar offerings do not overlap with the new energy programs in the shale context. Overall, community colleges expand clean-energy training when local clean-energy investment rises and students sort into solar fields when local demand for those jobs increases. Similar to the shale case, the response from colleges is relatively muted compared to the size of the shock; however, unlike shale, the student response is also relatively small, which may reflect the distribution of solar workers. The solar industry is characterized by a heavily front-loaded workforce in which temporary construction and installation constitute the bulk of jobs. Thus, the lack of large sustained labor demand may affect the number of students sorting into solar-related training.

A second reason the estimated response may appear small is that demand for solar training may be more diffuse than in the shale setting. Distributed solar follows population rather than solar potential, installers travel across local labor markets, and state policy shifts

Figure 20: Solar Potential, Solar Capacity, and Community Colleges in California



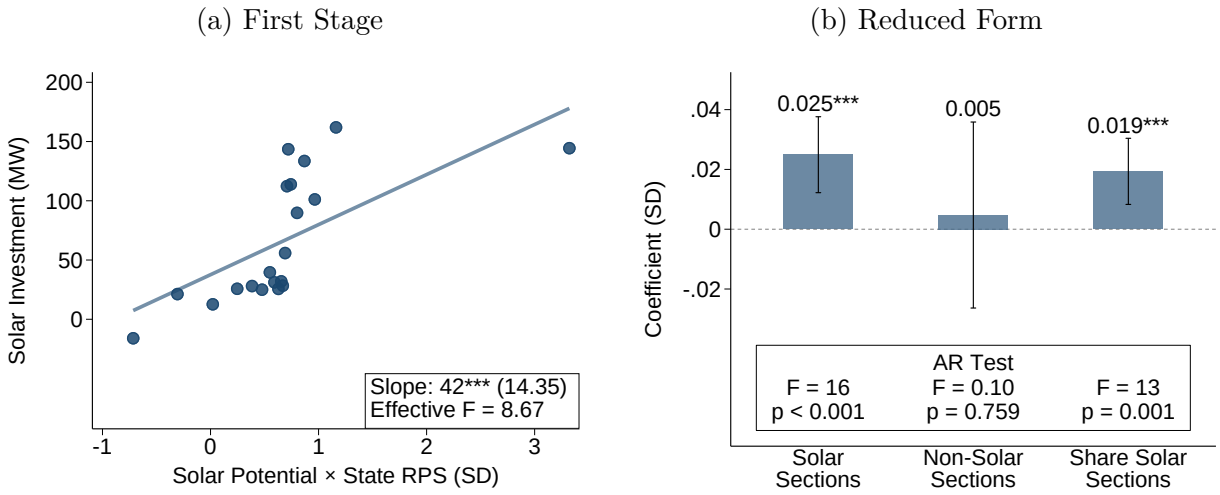
Notes. Panel (a) shows California counties by quintile of solar potential, average annual global horizontal irradiance (GHI) times county land area (Sengupta et al., 2018), which is the exposure component of the California instrument. Panel (b) shows counties by quintile of total installed solar capacity in 2024, utility-scale capacity (EIA, 2025b) plus distributed capacity (LBNL, 2025b). Panel (c) shows each utility-scale solar plant operating or on standby in 2024 from EIA Form 860, with color and size increasing in nameplate capacity. Quintiles are computed across California's 58 counties. Thick lines outline 1990 commuting zones and thin lines outline counties. Dots are public two-year colleges from NCES IPEDS at their most recent reported location. Panel (c) uses the Esri World Physical basemap.

demand statewide. To gauge how much of the expansion in solar training is tied to local investment, I multiply the estimates by the investment in each college's own commuting zone and sum across the sample. Local investment accounts for about 16 additional solar programs and 116 additional solar graduates per year by 2024, roughly 6% of the observed increase in programs (264) and 15% of the observed increase in annual completions (759) over the sample period. Most of the growth in solar training is therefore not explained by local investment, either because colleges respond to regional or statewide demand or because of national trends in the industry. If part of the response is diffuse, the local estimate understates the total training response to solar investment, and the comparison with shale, where the work is tied to sites within the plays, should be read with this in mind.

5.4 California Course-Level Evidence

One constraint to the shale analysis is that colleges often first respond to demand changes by adding courses rather than entire programs to meet students' and employers' needs.

Figure 21: First Stage and Reduced Form for California Solar Investment



Notes. This figure plots the first stage (a) and the reduced form (b) for California. Panel (a) is a binned scatter plot of county solar investment against the instrument, after partialling out population and year and commuting-zone fixed effects. The panel also reports the effective F-statistic. Panel (b) plots the reduced-form coefficient on the instrument, in standard deviations of the annual change in solar sections, non-solar sections, and the solar share of sections, with the F-statistic and Anderson–Rubin p-value for each outcome. Bars show 95% confidence intervals. Standard errors are clustered at the county level. Appendix Table B5 reports the full estimates.

However, the IPEDS data is only available at the program level and there are no shale booms during my sample period in California. The solar shock provides a new context in which I can explore a more granular level of analysis through California community college course offerings. Looking inside programs, I turn to course-section data for all California community colleges (CCCCO DataMart, 1993–2025), in which I can distinguish solar-specific sections from broader trades.

Figure 20 illustrates the California variation in solar potential and installed capacity. With a single statewide RPS, the instrument differs across counties only through solar potential, which is highest in the large desert counties of the south and east (Panel a). Solar capacity is concentrated in those same southern counties and in the Central Valley, where most utility-scale plants sit (Panels b and c). The large counties of the far north also score high on potential, mainly related to their size, but they saw little solar development. This mismatch likely contributes to the weaker first stage in California. Community colleges are concentrated in the Bay Area, the Los Angeles basin, and the Central Valley, so many of them sit close to the state's solar build-out.

California is a harder test. With only 38 counties, the instrument has much less variation to work with, and the first stage is weaker ($F \approx 9$) as Figure 21 shows. The estimates are

Table 12: Effect of Solar Investment on Course Sections in California

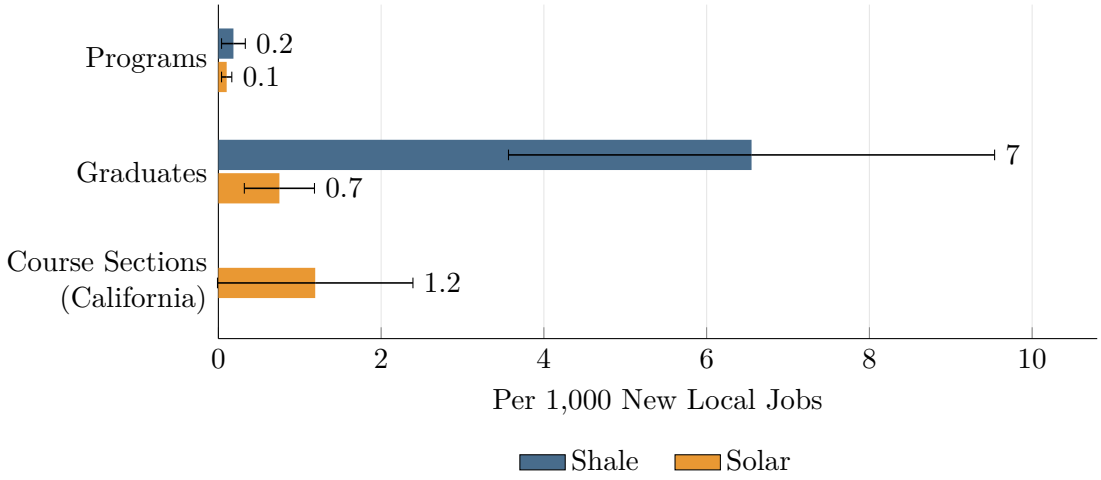
Outcome:	Δ Share Solar Sections	Δ Solar Sections	Δ Non-Solar Sections
Raw Outcome:			
$\widehat{SolarInvestment}_{t-1}$ (GW)	0.05**	0.40*	11
	(0.03)	(0.20)	(34)
Effect per SD:			
$\widehat{SolarInvestment}_{t-1}$ (GW)	0.46**	0.59*	0.11
	(0.21)	(0.30)	(0.35)
Year FE	Y	Y	Y
CZ FE	Y	Y	Y
Population Control	Y	Y	Y
Wald F-Stat	8.67	8.67	8.67
Observations	2,567	2,567	2,567
Sample Mean	0.06	0.43	679
SD of ΔOutcome	0.12	0.67	98

Notes. Outcomes are the share of solar sections, the count of solar sections, and the count of non-solar sections offered by California community colleges in the county. The table reports effects on the raw outcome and per standard deviation of the annual change in the outcome. Solar investment is instrumented using the first stage in Table B5. Lag is one year. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the county level in parentheses.

correspondingly less precise, but they tell a consistent story (Table 12). The reduced form is significant for solar sections and the solar share of sections (Anderson–Rubin p-values of 0.001 or less) but not for non-solar sections ($p = 0.76$). Unlike the typical 2SLS t-test, the Anderson–Rubin test remains valid when the first stage is weak, so it shows that solar investment expands solar offerings; however, the test does not pin down the size of the effect. Therefore, the 2SLS magnitudes below are less certain than in the national sample.

A one-gigawatt increase in predicted solar investment raises the share of solar course sections by 0.054 percentage points (0.03) and solar sections by 0.59 SD (0.30), about 93% of the sample mean, while leaving non-solar sections essentially unchanged. The average college's county added 1.5 gigawatts of solar capacity over 2002–2024, so across the 113 colleges the estimates imply an average of about 19 additional solar sections on offer each year. Neither the choice of policy shifter nor regional growth drives these estimates. Replacing the state RPS with the federal ITC gives nearly identical effects (Appendix E.5). As in the shale setting, the response is a targeted expansion toward the booming field rather than a general increase in activity. This contrasts with Grosz (2022), who finds no response of course sections to gradual shifts in occupational employment in the same California data. As expected, compared to the program-level response in the national sample, the course-level adjustment is both stronger, almost eight times larger in SD terms, and quicker, operating on a one-year lag rather than two.

Figure 22: Shale and Solar Responses per 1,000 New Local Jobs



Notes. This figure compares the community college response to the shale boom and to solar investment per 1,000 new local jobs. Shale effects are the difference-in-differences estimates per college for energy programs (Table 4) and energy completions (Table 5), multiplied by the 100 colleges in the 58 treated commuting zones and divided by the 1,455 total energy jobs across all skill levels that the boom added to the average treated commuting zone (Table 3). Solar effects are the 2SLS estimates per college per gigawatt (Table 11), multiplied by the number of colleges per commuting zone weighted by solar investment (9.8) and divided by 1,800 jobs per gigawatt (Section 5.2). California course sections use the same scaling, with the 2SLS estimate per college per gigawatt (Table 12) multiplied by the number of colleges per county weighted by solar investment (5.4) and divided by about 1,800 jobs per gigawatt. Bars show 95% confidence intervals from the regression standard errors, holding job counts fixed.

5.5 Comparing the Solar and Shale Responses

The solar responses look modest compared to those for shale, but how do we compare a shale "boom" to one gigawatt of solar investment? The shocks are not the same size and in fact measure different margins. What the two shocks have in common, and community colleges and students respond to, is the demand for workers they create. Thus, I compare the responses per 1,000 new local jobs. The distinction matters if the shale results are to inform the energy transition. If colleges respond as strongly to solar as they do to shale per new job, the smaller solar response simply reflects a smaller shock. If they respond less, training may fall behind the demand that the expansion of renewables and growing electricity demand will create. The shale boom added about 1,455 total energy jobs across all skill levels to the average exposed commuting zone (Table 3). Each gigawatt of solar capacity adds about 1,800 jobs (Section 5.2). Figure 22 depicts the response. Per 1,000 new jobs, the colleges in a local labor market produced about 7 additional energy graduates a year and offered 0.2 additional energy programs under the shale boom, compared with 0.7 solar graduates a year and 0.1 solar programs under solar investment. Even holding the size of the shock fixed, shale drew nearly ten times the graduate response of solar and nearly twice the program response, though program responses are not statistically distinguishable.

In California, where I observe courses, solar investment added about 1.2 course sections per 1,000 new jobs, indicating that much of the adjustment to solar may happen inside existing programs rather than through new ones.

6 Program Reallocation and Student Earnings

The results I have presented thus far establish that community colleges adjust their offerings and graduate more workers after energy transformations occur. How this response affects the future earnings of energy graduates is a separate question. This is particularly relevant in the case of shale where energy credentials tend to offer higher wages than alternative credentials, especially at the peak of the boom, but, given the sector's cyclicality, could leave graduates with lower earnings if their skills do not transfer to other occupations once the boom subsides (Katovich et al., 2026).

6.1 Empirical Design

To address this question, I use linked administrative data from the Texas ERC, described in Section 3.1, to estimate the earnings associated with energy credentials. The data connect each student's K–12 history (TEA), college enrollment and completion (THECB), and quarterly unemployment-insurance earnings (TWC) at the individual level.

My sample includes community college completers who have a linked eighth-grade record and earn a credential within eight years of expected high school graduation. Expected graduation is assigned based on the student's eighth-grade cohort and is defined as four years following a student's first eighth-grade record. I refer to the end of the community college completion window as expected completion. Expected completion falls at approximately age 26, the median age of reported community-college completion in my data, and leaves sufficient years afterward to examine longer-run earnings. I exclude incarcerated students. Field of study is assigned based on the first year a student earns any credential. If a student earns multiple credentials in that year, I use a priority hierarchy that places energy first, followed by health care, business, computer science, and the remaining fields to assign field[20].

Earnings are measured using the net present value (NPV) of cumulative quarterly wages in the first 1, 5, 10, and 15 years after expected completion. The NPV combines the level and timing of earnings, including potential unemployment spells during downturns, captur-

[20]Results are robust to assigning field by the first credential earned by month and only using the priority hierarchy for tie breaks in the case of multiple credentials earned in the same month and year.

ing realized downside. I aggregate quarterly earnings into July–June academic years and calculate cumulative earnings over each horizon, discounting annual earnings back to the expected completion date at an annual rate of 3%. Earnings records extend through the second quarter of 2024, providing 15 years of earnings for the 1994–1997 eighth-grade cohorts. I measure these horizons relative to the cohort-specific expected completion date, rather than each student's actual completion date, to keep the earnings windows independent of actual completion timing, which may itself respond to the boom. Each horizon of earnings estimates only includes cohorts for which the full earnings period is available. Within these cohorts, I retain all matched individuals who appear in the UI records at least once regardless of the number of quarters for which earnings are recorded, meaning that some quarters will have no recorded earnings for some individuals. These quarters enter as zeros, so a zero can reflect either non-employment or employment not covered by Texas UI records. I discuss attrition concerns below.

Following Andrews et al. (2024), I estimate the earnings differentials associated with fields of study among community college completers using a selection-on-observables design:

$$Y_{iscj} = \mu + \sum_{k} \theta_k \mathbf{1}\{\text{Field}_i = k\} + X_i'\mathbf{\Omega} + \delta_{cs} + \gamma_{cj} + \varepsilon_{iscj} \tag{15}$$

where Y_{iscj} is the NPV of cumulative recorded earnings for student i from middle school s and eighth-grade cohort c who completed a credential at college j. I estimate the model separately for earnings accumulated over each horizon (1, 5, 10, and 15 years after expected completion). X_i is a vector of academic and demographic characteristics, including standardized eighth-grade math and reading scores, gender, ethnicity, and indicators for special education, limited English proficiency, economic disadvantage, at-risk status, gifted status, and high school CTE coursework. Middle-school-by-cohort fixed effects δ_{cs} account for factors shared by students from the same middle school and cohort, while college-by-cohort fixed effects γ_{cj} account for factors shared by students from the same cohort completing credentials at the same college. The coefficients of interest, θ_k, measure the earnings differential for completers of credentials from field k relative to the omitted category, in this case other engineering and technicians, conditional on these controls. Standard errors are clustered at the community college by cohort level.

Earnings are only observed for employment covered by the Texas unemployment insurance system. Individuals who leave the state entirely never appear in the earnings records and drop out of the sample, while those observed for a portion of the window enter with zeros for unobserved quarters. Prior work on education in Texas suggests that this is of

limited concern for the state (Kovalenko, 2023; Andrews et al., 2024; Foote and Stange, 2022). Overall, Texas has low out-migration of young workers (Foote and Stange, 2022), and Andrews et al. (2024) find little difference in the number of observed earnings quarters across fields of study. For shale specifically, Kovalenko (2023) finds no evidence that the boom changed either out-of-state migration or the probability of inclusion in Texas college or earnings records. In line with Andrews et al. (2024), I estimate Equation 15 for the number of non-missing earnings quarters and find no evidence of differential attrition for energy completers relative to the omitted group in my sample. For other fields, differences in observed quarters are small (see Appendix E.7 for results).

For θ_k to measure the causal effect of selecting into field k on earnings, community college field must be as good as randomly assigned conditional on the controls. The rich set of pre-college measures, particularly test scores, absorbs much of the selection on preparation, but there are a number of reasons that students may select into these fields that are not captured here. Students may have existing connections to the field through their networks, they might simply have a preference for certain types of work, or they may be more tolerant of certain working conditions. If these factors increase a student's quality of work or productivity in her job, they will be positively correlated with her earnings. Since they also make her more likely to choose the field, the result is positive selection, biasing estimated premia upward. As such, I interpret the estimates as descriptive rather than causal. This concern is likely to be the smallest for the comparison that I focus on, energy versus other engineering and technician credentials because students in these closely related fields plausibly share similar characteristics.

6.2 Returns to Energy Credentials

Table 13 reports the estimated differentials by years since expected community college completion, which I use as a proxy for "experience level". Energy completers earn approximately 14% more than completers of other engineering and technician credentials in the NPV of earnings within ten years of expected completion. The proportional differential is similar for 15-year earnings, although 15-year estimates are only significant at the 10% level. Notably, the premium is much smaller, less than 5%, and not statistically significant in the first one and five years after completion. Because each horizon includes a different set of cohorts, the pooled estimates cannot distinguish between the evolution of premia across experience levels and compositional differences across cohorts. I return to this in the next subsection, where I estimate the differential separately for each horizon and set of cohorts.

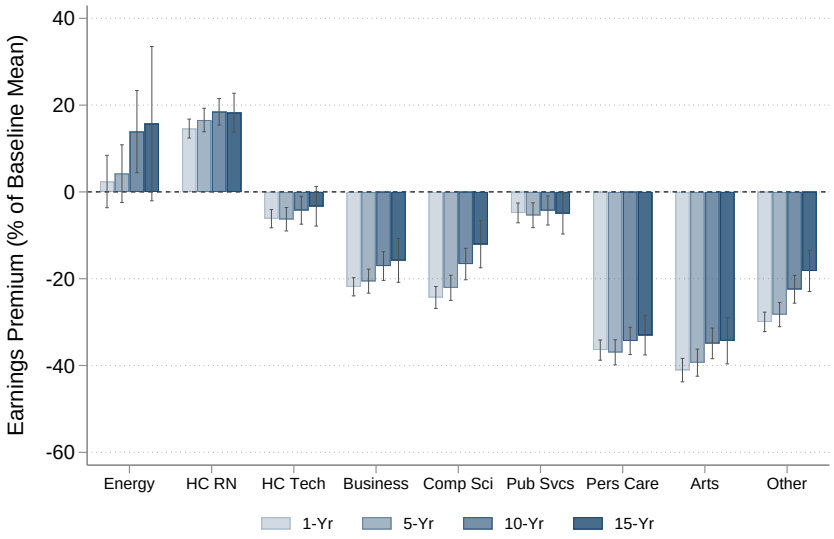
Table 13: Earnings Differentials by Earnings Horizon

Earnings NPV:	1-Year (1)	5-Year (2)	10-Year (3)	15-Year (4)
Program:				
Energy	1,431	11,943	75,115***	123,586*
	(1,842)	(9,615)	(26,142)	(71,284)
HC Nurse	8,718***	47,004***	99,819***	143,434***
	(665)	(3,931)	(8,419)	(17,963)
HC Technician	-3,689***	-17,885***	-22,845**	-25,996
	(644)	(3,904)	(8,871)	(18,328)
Business	-13,060***	-58,347***	-92,438***	-124,070***
	(641)	(4,011)	(9,118)	(20,228)
Computer Science / IT	-14,532***	-62,706***	-89,920***	-94,732***
	(768)	(4,187)	(10,017)	(21,796)
Public Service	-2,892***	-15,260***	-23,131**	-39,081**
	(696)	(4,146)	(9,236)	(18,982)
Personal Care	-21,777***	-104,891***	-185,713***	-259,501***
	(710)	(4,195)	(8,690)	(18,284)
Arts	-24,535***	-111,658***	-188,790***	-269,440***
	(822)	(4,513)	(9,764)	(21,469)
Other	-17,896***	-80,235***	-121,326***	-143,261***
	(685)	(4,034)	(8,838)	(19,016)
Cohort×Middle School FE	Y	Y	Y	Y
Cohort×CC FE	Y	Y	Y	Y
Test Score Controls	Y	Y	Y	Y
Demographic Controls	Y	Y	Y	Y
Grade 8 Cohort	1994-2011	1994-2007	1994-2002	1994-1997
Observations	179,303	129,538	75,753	31,248
R-Squared	0.367	0.345	0.335	0.322
Baseline Mean	59,749	283,810	540,972	786,142

Notes. This table reports estimates of Equation 15 where the outcome is the net present value of cumulative earnings 1, 5, 10, and 15 years after expected completion. The omitted field is other engineering and technician credentials. All specifications include the individual covariates and middle-school-by-cohort and college-by-cohort fixed effects described in the text. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are clustered at the community college by cohort level.

To visualize this comparison and put it into context, Figure 23 plots the estimated differentials for each field and experience level. The omitted field, other engineering and technicians, is among the highest-earning on its own, so energy's advantage over most other fields appears considerably larger. In fact, relative to all fields but one, energy credentials are associated with earnings premia at most horizons. The exception is nursing, which commands large and precisely estimated premia at every earnings horizon. Energy credentials are therefore among the highest-earning fields community college students complete.

Figure 23: Earnings Differentials Across Fields by Earnings Horizon



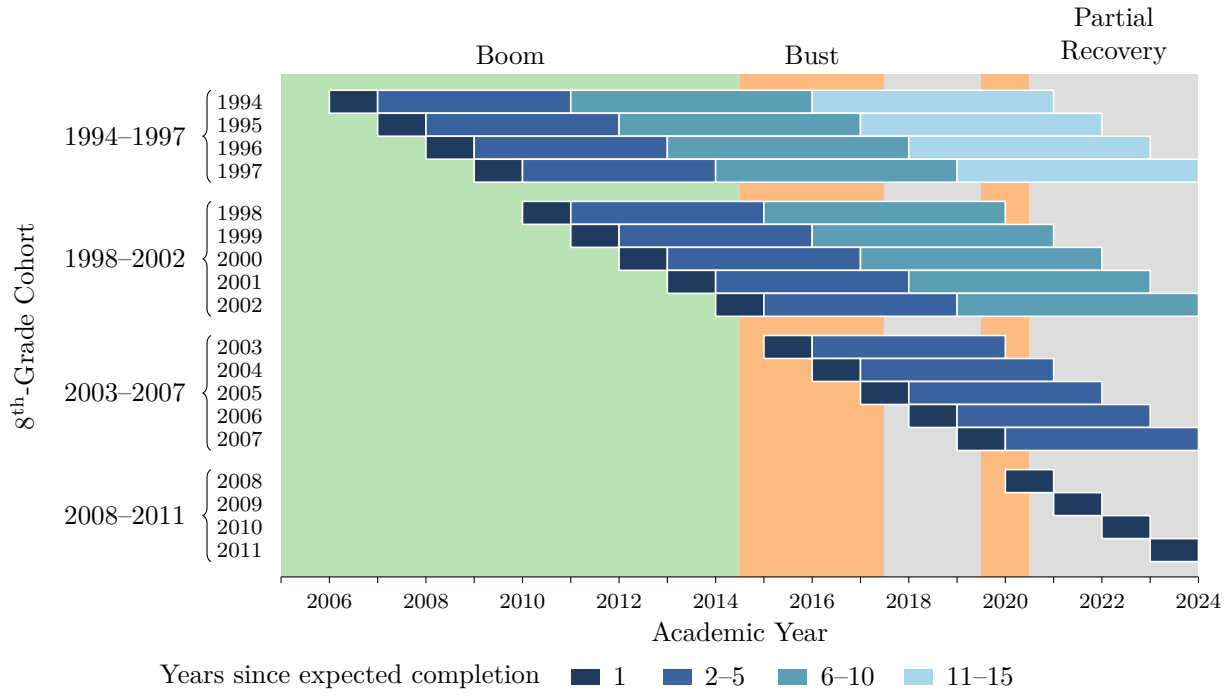
Notes. This figure plots estimates of Equation 15 for each field of study at each earnings horizon as a percent of the baseline mean for that horizon. The omitted field is other engineering and technician credentials. Capped lines indicate 95% confidence intervals. Standard errors are clustered at the community college by cohort level. Each horizon is estimated on a different set of cohorts, so differentials should be compared across fields for a given horizon.

6.3 Returns over the Boom-and-Bust Cycle

The pooled estimates above can mask variation in the energy differential across experience levels, cohort composition, and oil-and-gas boom-and-bust cycles, so I estimate Equation 15 separately for the four sets of eighth-grade cohorts. As oil and gas employment and earnings rise and fall, both the students who select into energy fields and the labor markets they enter change. The results reported in Table 14 illustrate how the earnings differential to energy fields varies across these cohorts.

The Texas oil and gas industry went through a boom and two busts over the period my earnings data cover. The shale gas boom began in the early 2000s, and in the mid-2000s, drilling shifted toward oil as horizontal wells spread across the Eagle Ford play and the Permian Basin (Gilmer et al., 2012; Federal Reserve Bank of Dallas, 2020; Thompson, 2022). Oil prices remained high until mid-2014 and then fell by more than 70% by early 2016 (Baumeister and Kilian, 2016; Slijk and Phillips, 2021), resulting in the most severe collapse since 1986 (Thompson, 2022). The boom boosted local employment and earnings and the bust eroded them (Abboud and Betz, 2021). Activity rebounded in 2018 and 2019 before the Covid pandemic brought a second collapse in 2020 (Federal Reserve Bank of Dallas, 2020; Thompson, 2022). Since 2020, oil production in the Permian Basin has climbed up

Figure 24: Calendar Timing of Earnings Horizons by Cohort

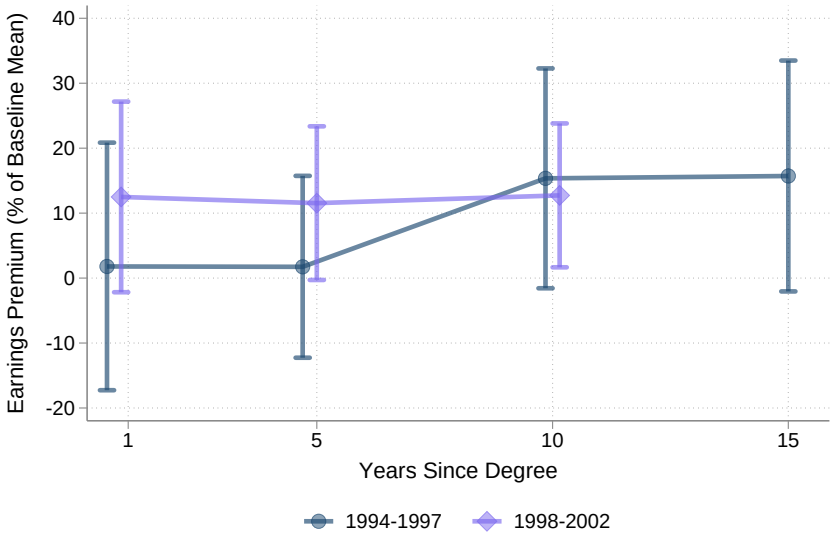


Notes. This figure shows the calendar years covered by each eighth-grade cohort's earnings horizons, measured in July–June years after expected completion. Each cohort appears only in the horizons for which its cohort group has complete earnings data. Shading marks the boom (2005–2014), bust (2015–2017 and 2020), and partial recovery (2018–2019 and 2021–2024). Earnings records end in June 2024.

to record levels (Federal Reserve Bank of Dallas, 2024), but employment has not kept pace, remaining nearly one third below 2014 levels as of 2024 (Leigh et al., 2026; BLS, 2026a). For the purposes of this analysis, I define the boom period as 2005–2014 and 2015–2017 as the first bust in line with Abboud and Betz (2021). The second bust is 2020, and 2018–2019 and 2021–2024 are partial recovery.

Because the cohorts reached expected completion at different points in this cycle, their earnings horizons overlap its phases to varying extents (see Figure 24). The 1994–1997 cohorts reached expected completion between 2006 and 2009. Their first five years of earnings fall within the boom. Their years 6 through 10 span the end of the boom and the bust, and their years 11 through 15 fall mostly in the partial recovery. The 1998–2002 cohorts reached expected completion between 2010 and 2014, during the height of the boom. Their first year falls almost entirely within the boom, their years 2 through 5 are split between the boom and the bust, and their years 6 through 10 fall entirely after the boom. The 2003–2007 cohorts reached expected completion between 2015 and 2019. Their first year of earnings falls mostly in the bust, and their years 2 through 5 fall mostly in the partial recovery. The

Figure 25: Earnings Differentials by Cohort for Energy Completers



Notes. This figure plots estimates of Equation 15 at each earnings horizon for energy completers in the 1994–1997 and 1998–2002 cohorts. Estimates for the 2003–2007 and 2008–2011 cohorts, which are observed only at short horizons, are reported in Table 14. The omitted field is other engineering and technician credentials. Bars indicate 95% confidence intervals. Standard errors are clustered at the community college by cohort level.

2008–2011 cohorts reached expected completion between 2020 and 2023, and their first year of earnings falls mostly in the partial recovery after the 2020 collapse.

To explore the three types of variation masked by the pooled results, the estimates in Table 14 can be read in three ways. First, comparing cohorts for a given earnings horizon illustrates how the differential changes across students who chose their fields and entered the labor market at different points in time. Second, since each horizon corresponds to a range of calendar years, the estimates can be interpreted in the context of the shale boom-and-bust cycle. And, third, following a single cohort across earnings horizons provides evidence on the evolution of the earnings differential with "experience" for a given set of completers.

Starting with the last comparison and tracing the 1994–1997 and 1998–2002 cohorts in Figure 25, the same completers appear at every horizon, so changes in the earnings gap over time do not reflect changes in selection. It is clear that the two cohorts follow very different trajectories. For the 1994–1997 cohorts, the differential is less than 5% and not statistically significant over the first 5 years. It rises steeply to 15% for 10-year earnings and remains near that level for 15-year earnings. Conversely, the 1998–2002 cohorts begin with a strong earnings differential of about 12% from the first year, which remains between 12-13% through year ten. These results help to explain the pooled estimates in Table 13, where

the differential is small and imprecise for 1- and 5-year earnings but large and significant for 10-year earnings. Within a cohort, horizon or "experience" can never be separated from calendar time, so these patterns could reflect returns to experience, labor market changes, or both. Because the two cohorts follow very different paths over the same earnings horizons, the results cannot be explained by experience alone, suggesting that the labor-market conditions or compositional differences in the cohorts must also play a role. I explore these mechanisms next.

Comparing across cohorts for a given earnings horizon holds "experience" fixed while allowing the composition of students to vary. However, these cohorts chose their fields at different points in time, so variation can reflect both who selects into energy and the conditions they face upon completion. Only one cohort, 1998–2002, exhibits significant earnings premia for the 1- and 5-year earnings horizons of approximately 12% relative to the omitted field. The estimates for all other cohorts are close to zero and not statistically significant (see Table 14). The earnings gap for energy credentials is largest at the 10-year horizon, around 13-15%, for the two cohorts that can be observed despite each reaching the horizon at a different calendar period.

Finally, I interpret these results through the lens of the shale boom-and-bust cycle. Table 15 re-estimates the differentials over non-overlapping earnings periods and aligns them with the phases of the shale cycle. The phase alone does not appear to be predictive of the earnings premium. Cohorts who are expected to have entered the labor market during the boom carry vastly different differentials. While the 1998–2002 cohort exhibits a 12% earnings gap for 1-year earnings, the same gap is only around 2% for the 1994–1997 cohort. Earnings during the transition from boom into bust and partial recovery carry large premia for the two early cohorts, but are essentially null for later ones. Together, these results suggest that earnings may be less sensitive to the oil and gas cycle than previously expected given the volatility of the sector. Instead, a story about overall labor market conditions, potentially exacerbated by the shale cycle, seems to fit the patterns better.

Early-career differentials line up with the conditions during each cohort's expected completion. The 1998–2002 cohort arguably entered the labor market under some of the most favorable conditions in my sample with expected completion between 2010 and 2014. Recovery from the 2008 Great Recession was underway and the 2005-2014 Texas shale boom was at its peak. Their expected community college enrollment years also had the most overlap with both the shale boom and the expansion of energy programs documented in Section 4. If the boom attracted energy completers who were less suited to the work, this cohort might

Table 14: Earnings Differentials by Cohort

Earnings NPV:	1-Year				5-Year			10-Year		15-Year
	1994-1997	1998-2002	2003-2007	2008-2011	1994-1997	1998-2002	2003-2007	1994-1997	1998-2002	1994-1997
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Program:										
Energy	661	5,775*	512	-897	3,672	30,148*	586	72,991*	74,546**	123,586*
	(3,583)	(3,460)	(2,619)	(3,727)	(15,091)	(15,768)	(13,875)	(41,045)	(33,074)	(71,284)
HC Nurse	10,592***	8,542***	7,236***	10,674***	48,832***	44,039***	50,581***	87,982***	108,153***	143,434***
	(1,029)	(979)	(1,165)	(1,556)	(5,404)	(5,422)	(7,226)	(10,822)	(11,920)	(17,963)
HC Technician	1,449	-2,077**	-5,179***	-6,034***	447	-14,678**	-29,078***	-12,902	-29,115**	-25,996
	(961)	(1,021)	(1,156)	(1,412)	(5,577)	(5,712)	(6,872)	(11,229)	(12,559)	(18,328)
Business	-7,520***	-10,938***	-14,109***	-17,466***	-43,797***	-55,275***	-69,734***	-89,713***	-94,732***	-124,070***
	(984)	(1,026)	(1,252)	(1,411)	(5,732)	(5,882)	(7,454)	(12,272)	(12,902)	(20,228)
Computer Science / IT	-6,284***	-12,273***	-16,680***	-18,926***	-36,876***	-60,848***	-80,963***	-74,817***	-104,370***	-94,732***
	(915)	(1,124)	(1,379)	(2,011)	(6,037)	(6,150)	(7,788)	(13,015)	(14,521)	(21,796)
Public Service	-1,186	-3,887***	-4,107***	-311	-11,710**	-15,311**	-16,206**	-29,679**	-19,524	-39,081**
	(916)	(1,015)	(1,269)	(1,637)	(5,443)	(6,429)	(7,386)	(11,565)	(13,281)	(18,982)
Personal Care	-12,627***	-18,081***	-22,587***	-28,809***	-74,277***	-96,467***	-126,154***	-165,441***	-198,327***	-259,501***
	(954)	(973)	(1,234)	(1,538)	(5,024)	(5,530)	(7,685)	(10,878)	(12,329)	(18,284)
Arts	-13,802***	-19,526***	-26,123***	-33,839***	-79,716***	-99,406***	-141,522***	-171,177***	-201,201***	-269,440***
	(1,235)	(1,103)	(1,325)	(1,737)	(6,068)	(6,162)	(7,948)	(12,663)	(13,838)	(21,469)
Other	-9,251***	-13,842***	-19,376***	-25,175***	-51,136***	-71,910***	-103,620***	-101,970***	-136,323***	-143,261***
	(985)	(1,019)	(1,289)	(1,584)	(5,357)	(5,695)	(7,473)	(11,227)	(12,886)	(19,016)
Cohort×Middle School FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Cohort×CC FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Test Score Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Demographic Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Grade 8 Cohort	1994-1997	1998-2002	2003-2007	2008-2011	1994-1997	1998-2002	2003-2007	1994-1997	1998-2002	1994-1997
Observations	31,248	44,505	53,785	49,765	31,248	44,505	53,785	31,248	44,505	31,248
R-Squared	0.288	0.307	0.292	0.290	0.285	0.307	0.317	0.322	0.321	0.322
Baseline Mean	36,849	46,213	57,825	81,288	211,276	261,266	336,535	475,277	585,304	786,142

Notes. This table reports estimates of Equation 15 separately for four sets of eighth-grade cohorts. Column 10 reports 15-year earnings, which are observed only for the 1994–1997 cohorts. The omitted field is other engineering and technician credentials. All specifications include the individual covariates and middle-school-by-cohort and college-by-cohort fixed effects described in the text. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are clustered at the community college by cohort level.

Table 15: Energy Earnings Differentials by Cohort, Earnings Period, and Oil and Gas Cycle

8th Grade Cohort	Expected Completion	Earnings Period (Years since Expected Completion)			
		1 Yr	2–5 Yrs	6–10 Yrs	11–15 Yrs
1994–1997	2006–2009	2%	2%	26%*	16%
1998–2002	2010–2014	12%*	11%*	14%**	
2003–2007	2015–2019	1%	0%		
2008–2011	2020–2023	−1%			

Boom Boom into Bust Bust Partial Recovery

Notes. Each cell reports the energy earnings differential within the indicated period after expected completion as a percent of the period's baseline mean. The periods do not overlap. Estimates come from Equation 15 with the NPV of earnings in each period as the outcome. Expected completion is eight years after expected high school graduation. Shading indicates the phase in which at least 60% of the period's earnings years fall, and cells without such a phase span the transition from boom to bust. Boom is 2005–2014, bust is 2015–2017 and 2020, and recovery is 2018–2019 and 2021–2024.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are clustered at the community college by cohort level.

be expected to exhibit smaller earnings differentials. Instead, theirs are the largest across all cohorts at every observable earnings horizon in dollar terms, and close to the largest in percentage terms. These estimates are consistent with positive selection into energy fields during the boom, improved training as energy programs expanded, favorable labor demand, or some combination of these mechanisms. The fact that the large premium persists through 10-year earnings suggests that a short-lived advantage at labor-market entry is unlikely to be the only factor at play. Furthermore, energy completers in the 1998–2002 cohort appear similar to those in the earlier 1994–1997 cohort on every observed characteristic in my data with the exception of high school CTE coursework, which is less common in later cohorts (see Appendix Table B3). This comparison provides no evidence of positive selection on observable characteristics, although selection on unobserved ones remains possible.

Turning to the post-boom cohorts (2003–2007 and 2008–2011), both reached expected completion during either the bust, partial recovery, or the Covid pandemic. Energy completers in these late cohorts also appear different to the earlier cohorts on observable characteristics. Appendix Table B3 indicates that a larger share are Hispanic and smaller shares are White or in the combined Black or other category. They are also more likely to be economically disadvantaged and, in the case of the last cohort, exhibit lower eighth-grade reading scores. Since I control for these observed characteristics in Equation 15, these differences only matter to the extent that they also signal differences in unobserved characteristics. However, the absence of an early differential is not yet informative about the later cohorts' longer-run

returns. The 1994–1997 cohort similarly exhibits no early differential, but develops a sizable earnings gap at longer horizons.

Finally, the 1994–1997 cohort appears to be the exception. Their early years fall within the shale boom yet their earnings differentials are small. However, they did reach expected completion between 2006 and 2009, so the first earnings years for most of these completers fell during or just after the Great Recession and the 2008 collapse in oil prices. This pattern is consistent with evidence that the labor market conditions workers face when they finish school affect their earnings for years afterward (Kahn, 2010; Oreopoulos et al., 2012) and that these effects differ across fields of study (Altonji et al., 2016). Their earnings premium emerges only in years six through ten, which is consistent with the gradual recovery from poor entry conditions documented by Oreopoulos et al. (2012) or returns to experience.

Taken together, the cohort estimates yield two main conclusions. First, selection into energy fields does not appear to explain variation in earnings differentials across cohorts. Second, I find no evidence that the shale bust lowered the returns to energy credentials[21]. Once a cohort earned a premium, it persisted after the boom subsided. Completers whose training coincided with the boom and expansion of energy programs continued to earn more than completers of other engineering and technician credentials in subsequent years. Still, the estimates are consistent with entry into a booming labor market boosting early-career returns. I also cannot rule out returns to experience contributing to later-year premia. These conclusions are limited by the small number of energy completers in the earliest cohorts and the limited timeframe over which the post-boom cohorts can be observed. One additional caveat is that my earnings measure, NPV of cumulative earnings, does not penalize volatility, so these estimates may overstate the value of energy credentials for risk-averse workers given the sector's cyclicality.

Nevertheless, one result holds across all cohorts and earnings horizons: energy completers earn at least as much as their closest alternative, completers of other engineering and technician credentials. They also out-earn completers of every other field, except nursing, at most horizons. Thus, the results provide descriptive evidence that energy credentials yield high relative returns whether they are completed during the boom, bust, or recovery. The programs colleges created in response to the boom may have channeled students into one of the highest-return vocational fields available to them. I complement these within-Texas results with a national analysis using the Census Bureau's PSEO tabulations in Appendix D, which reveals a broadly similar pattern.

[21] In contrast, Katovich et al. (2026) find that most workers hired into Brazil's oil industry late in its boom faced persistently lower earnings than matched workers in other industries.

7 Discussion and Policy Implications

The evidence here tells a consistent story. Community colleges are not the inert institutions much of the literature describes. Faced with large, local demand shocks, they reorient their portfolios, adding sections and programs. They graduate more students aligned with the booming sector, mostly from programs that did not exist prior to the boom. But, the response is slow, uneven, and not without cost. Each of these qualifications carries a lesson for how to think about the role of community colleges in the energy transformations that are currently underway.

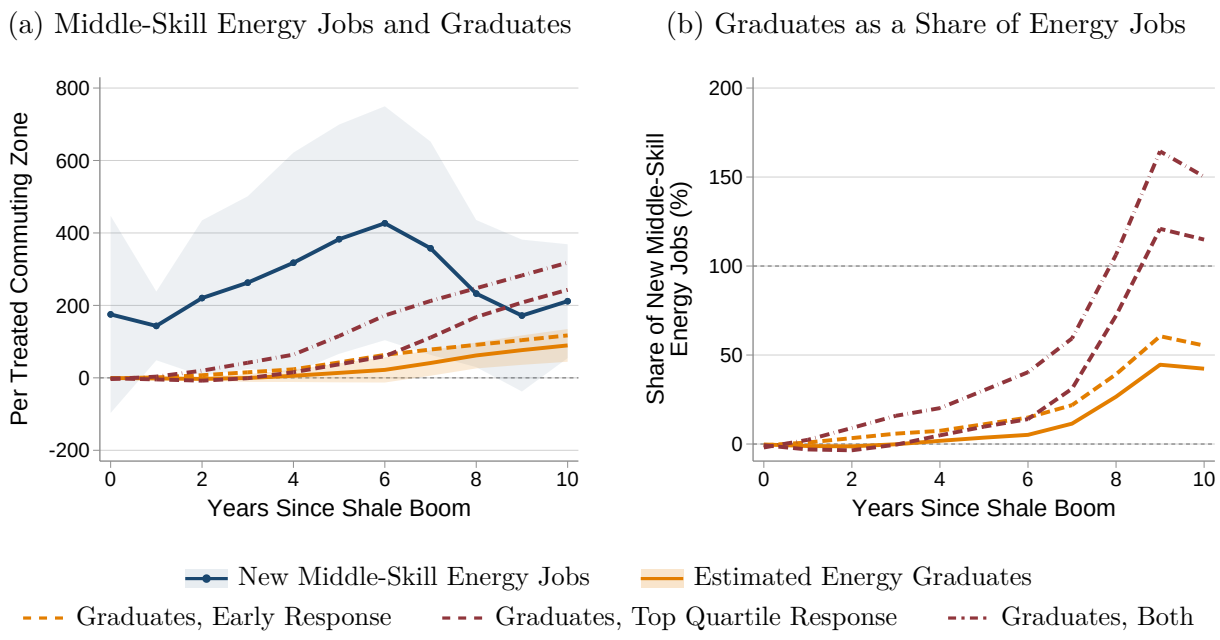
The average response is real but modest relative to the shock that provokes it. To characterize the size relative to the employment shock of the shale boom, I use the estimates from Section 4.2, measuring the magnitude of the shock on total and middle-skill energy employment. The boom raises total energy employment by 46% in treated commuting zones. Set against the roughly 1,455 additional energy jobs that appear in a treated commuting zone, the extra five energy completions each college produces annually seem like a small flow. However, of those jobs, I approximate that about 18%, or 264, are middle-skill, indicating that the boom adds about 15,300 workers with the training that community colleges provide[22]. This suggests that community colleges supply enough graduates to cover over one third of the new middle-skill employment in the 10 years after the boom[23]. Community colleges are therefore one channel of local adjustment among several, including in-house training, private trade schools, in-migration, the upskilling of incumbent workers, and hiring into roles that require no postsecondary credential.

Within 10 years, colleges cover over one third of the new middle-skill employment, but employers need trained workers at the start of the boom, not at the end. Figure 26 compares the additional boom-induced middle-skill energy jobs in a treated commuting zone in each year with the cumulative additional energy graduates from its colleges up to that year. Employment rises immediately and peaks about six years after the boom, while the graduate response takes about five years to begin and continues to grow throughout the 10-year period. By the time of the employment peak, the additional graduates can only fill about 5% of the new middle-skill jobs. It is only by year 10, after employment has fallen, that they reach about 40%. Therefore, much of the training response arrives after the period in which it may be most needed. Two hypothetical scenarios illustrate what might close this

[22]Based on 264 additional middle-skill energy workers in each of the 58 treated commuting zones in my sample.

[23]Based on 5.53 additional energy completions per treated college per year, the estimate of Equation 5 for the full post-period, and 100 treated colleges over 10 years, or about 5,500 additional completions.

Figure 26: New Middle-Skill Energy Jobs and Cumulative Additional Energy Graduates



Notes. Panel (a) plots Gardner (2021) two-stage event study estimates of additional middle-skill energy employment per treated commuting zone (Figure 7b) and cumulative additional energy completions per treated commuting zone. This is calculated as the sum of the annual energy completions estimates from the boom year through each year, scaled by the 100 treated colleges and divided by the 58 treated commuting zones. Shaded areas are 95% confidence intervals. Middle-skill energy employment is measured as in Section 4.2. The early response shifts the completions path two years earlier. The top quartile response scales completions by the completions effect of the colleges in the top quartile of energy program response (about 2.7 times the average college after empirical Bayes shrinkage). Panel (b) plots cumulative additional graduates as a share of the new middle-skill jobs in each year. Standard errors are clustered at the play level.

gap. The first is on timing. If colleges responded earlier, graduates could respond in turn. Accelerating graduate responses by two years shows that additional energy graduates could fill about 15% of the new jobs at the peak. The second is on scale. If every college responded like the most responsive colleges, those in the top quartile of energy program adjustment, they would add approximately 2.7 times as many energy graduates as the average college, filling about 14% of peak employment. Since the two effects compound, together they could cover about 40% of peak employment. Panel (b) suggests that helping colleges adjust more quickly to labor demand may be a more appropriate lever. On its own, a stronger response would produce more graduates than new jobs once the boom fades (120% by year nine), while an earlier one peaks at about 60%. These shares are upper bounds. They assume that every additional graduate stays in the commuting zone and works in energy. Net job growth also understates hiring demand.

As another way to benchmark the adjustment, I measure the elasticity of overall program and completion response to approximated middle-skill energy employment growth. Table 16 scales the program and completion effects against the local middle-skill energy-employment

Table 16: Back-of-Envelope Elasticities

Outcome	DiD Estimate	Baseline Mean	% Change	Elasticity
Middle-Skill Energy Employment	264	695	38.0%	—
Energy CIPs	0.128	0.52	24.5%	0.65
Energy Completions	4.61	4.71	97.9%	2.58

Notes. Column 1 reports the DiD point estimate; column 2 the pre-period treated mean; column 3 the implied percent change; column 4 the elasticity with respect to the percent change in middle-skill energy employment. Row 1 is middle-skill energy employment; row 2 is the count of energy CIPs; row 3 is energy completions. College outcomes are estimated over a 10-year post-period to match the employment estimates.

shock. A boom that raises middle-skill energy employment in a treated area by about 38% raises the number of energy programs at a nearby college by 25% and energy completions by 98%[24]. Expressed as elasticities with respect to that employment shock, the program response is inelastic (0.65), while completions are elastic (2.58), reflecting both the student and institutional responses. The gap between the two is informative and consistent with the overall story. Community colleges add a handful of new programs, but graduations in those programs rise far more, suggesting that the program count results may understate institutional influence.

These magnitudes help reconcile the paper with the prevailing view that postsecondary training only tracks labor demand loosely and the effect mainly operates through graduations, where students respond to local signals (Grosz, 2022; Conzelmann et al., 2023). My contribution is to show that institutional supply moves in the same direction, and as colleges create new programs, graduations rise alongside them. The earnings evidence ties the supply response to the human capital literature on resource booms. Where Kovalenko (2023) and Cascio and Narayan (2022) document the costs booms impose on schooling at the K–12 and enrollment margins, the durable energy earnings premium for boom-trained graduates suggests the institutional response built skills that may have retained their value beyond the boom.

Although, to my knowledge, no prior work has directly studied program creation or calculated the same elasticities, Conzelmann et al. (2023) find that the elasticity of 4-year degrees with respect to major-specific labor demand is 1.3, where the variation in demand is driven by national shifts in industry employment. This is from a different college market and shock intensity, so it is not surprising that it comes out at about half my 2.6 estimate.

[24]Energy program response is 30% relative to the baseline mean for the full post-period, but 25% when restricted to a 10-year post-period. Energy completion response is 117% as a percentage of the baseline mean for the full post-period, but 98% for a 10-year post-period.

Using a similar instrument at the 2-year level, Grosz (2022) finds that a one-percentage-point increase in an occupation's share of employment raises its share of community college completions by about 0.53 percentage points. Here the comparable quantity is the ratio of the boom's effect on energy's share of completions (+0.57 percentage points) to its effect on energy's share of local employment (0.99 percentage points), which is roughly 0.58. This is similar to Grosz's estimate even though the shale boom was a concentrated and sustained shock rather than a gradual shift in occupational demand. The two benchmarks differ partly because energy completions start from a small base, so an increase that is large relative to the baseline mean may be a modest shift in the composition of graduates.

The heterogeneity carries another policy lesson. The response is concentrated in colleges with higher administrative salaries, and it is weaker where faculty costs are high. The faculty pay dimension suggests that budget constraints may play a role in a college's ability to respond to demand, with larger post-boom revenue declines relative to non-exposed colleges at institutions with higher pre-boom faculty pay. On the other hand, the administrative dimension does not appear to run through funding, pointing to something more complex. Higher average administrative pay may be associated with aspects of "quality" that are more difficult to quantify, for example, curriculum-development support, hiring pipelines, or integration with local employment. Understanding which of these mechanisms matters most would help identify the colleges that are constrained in their responses.

More work is needed to weigh the relative costs that reallocation imposes. The energy expansion came alongside a material relative contraction in non-energy completions and programs. This likely reflects students switching into energy, leaving college, or delaying enrollment, along with possible shifts in instructional resources toward the booming sector. Encouraging colleges to pursue a booming field is therefore not a free lever since it may draw capacity and student attention away from other vocational tracks. Energy appears to be one of the highest-return fields, mitigating some of these concerns; however, the non-energy decomposition shows that nursing, the highest-return field, is among those that gave way. Whether the tradeoff is worth making depends on what energy students would have otherwise studied, how sustainable the new demand is, and whether the skills acquired in those credentials are transferable beyond energy. These questions are beyond the scope of this paper but provide opportunities for future work.

Finally, the renewables results suggest the same institutional margin may be available for the clean energy transition, but they come with similar caveats. Colleges expand solar training where solar investment grows, yet that response may carry the same opportunity

cost. Community colleges might be a real and underused lever for regional adaptation, but we first need to better understand whether there are frictions dampening their adjustment before they can provide a meaningful reprieve from the energy transition's labor bottlenecks.

8 Conclusion

This paper asks what role community colleges play in labor market adjustment to energy-driven demand shifts. Using the shale boom as a natural experiment, I find that community colleges respond, expanding energy programs and more than doubling the energy credentials they award, while non-energy offerings contract. Responses to solar power expansion show the same institutional margin operating in the clean energy transition, though on a smaller scale for graduates.

At the same time, the response is partial in both contexts, covering only a fraction of the new local employment generated by the energy transformations. In the case of shale, it also appears to be concentrated in colleges with the administrative capacity to act, leaving others unmoved despite similar exposure. As Section 7 sets out, community colleges may be an underused lever for regional adjustment, but not one that can currently clear the market on its own.

This paper offers causal evidence that the U.S. community college system adapts its supply of training to major energy-driven labor demand shocks. As the country confronts the workforce demands of unprecedented electricity demand and a growing and decarbonizing energy sector, understanding and strengthening the institutions that can train those workers is essential. Whether community colleges can mediate the renewable response as fully as they did the shale boom remains an open question for future work.

References

Abboud, A. and Betz, M. R. (2021). The Local Economic Impacts of the Oil and Gas Industry: Boom, Bust and Resilience to Shocks. *Energy Economics*, 99, 105285. https://doi.org/10.1016/j.eneco.2021.105285

Acemoglu, D. and Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, Vol. 4B, 1043–1171. Elsevier.

Acton, R. (2021). Community College Program Choices in the Wake of Local Job Losses. *Journal of Labor Economics*, 39(4).

Acton, R. K., Morales, C., Cortes, K., Turner, J., and Miller, L. (2026). Community College Bachelor's Degrees: How CCB Graduates' Earnings Compare to AAs and BAs. EdWorkingPaper No. 26-1397, Annenberg Institute at Brown University.

Allcott, H. and Keniston, D. (2018). Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America. *Review of Economic Studies*, 85(2), 695–731.

Altonji, J. G., Kahn, L. B., and Speer, J. D. (2016). Cashier or Consultant? Entry Labor Market Conditions, Field of Study, and Career Success. *Journal of Labor Economics*, 34(S1), S361–S401.

American Association of Colleges of Nursing (AACN) (2026). Nursing Faculty Shortage Fact Sheet. Updated August 2026. https://www.aacnnursing.org/news-data/fact-sheets/nursing-faculty-shortage.

Anderson, T. W., and Rubin, H. (1949). Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations. *Annals of Mathematical Statistics*, 20(1), 46–63.

Andrews, R. J., Imberman, S. A., Lovenheim, M. F., and Stange, K. M. (2024). The Returns to College Major Choice: Average and Distributional Effects, Career Trajectories, and Earnings Variability. *Review of Economics and Statistics*, 1–45. https://doi.org/10.1162/rest_a_01503

Angrist, J. D., Graddy, K., and Imbens, G. W. (2000). The Interpretation of Instrumental Variables Estimators in Simultaneous Equations Models with an Application to the Demand for Fish. *Review of Economic Studies*, 67(3), 499–527.

Angrist, J. D., and Imbens, G. W. (1995). Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity. *Journal of the American Statistical Association*, 90(430), 431–442.

Arkolakis, C. and Walsh, C. (2024). The Economic Impacts of Clean Energy. *Brookings Papers on Economic Activity*, Fall 2024, 161–238.

Atkin, D. (2016). Endogenous Skill Acquisition and Export Manufacturing in Mexico. *American Economic Review*, 106(8), 2046–2085.

Autor, D. H., and Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5), 1553–1597.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, 103(6), 2121–2168.

Autor, D. H., Dorn, D., and Hanson, G. H. (2016). The China Shock: Learning from Labor Market Adjustment to Large Changes in Trade. *Annual Review of Economics*, 8, 205–240.

Autor, D., Dorn, D., and Hanson, G. (2019). When Work Disappears: Manufacturing Decline and the Falling Marriage Market Value of Young Men. *American Economic Review: Insights*, 1(2), 161–178.

Autor, D. H., Dorn, D., Hanson, G. H., Jones, M. R., and Setzler, B. (2025). Places versus People: The Ins and Outs of Labor Market Adjustment to Globalization. NBER Working Paper No. 33424.

Barr, A., & Turner, S. E. (2013). Expanding Enrollments and Contracting State Budgets: The Effect of the Great Recession on Higher Education. The ANNALS of the American Academy of Political and Social Science, 650(1), 168–193. https://doi.org/10.1177/0002716213500035

Bartik, T. J. (1991). Who benefits from state and local economic development policies? W.E. Upjohn Institute for Employment Research.

Bartik, A. W., Currie, J., Greenstone, M., & Knittel, C. R. (2019). The Local Economic and Welfare Consequences of Hydraulic Fracturing. American Economic Journal: Applied Economics, 11(4), 105–155. https://doi.org/10.1257/app.20170487

Baumeister, C. and Kilian, L. (2016). Understanding the Decline in the Price of Oil since June 2014. *Journal of the Association of Environmental and Resource Economists*, 3(1), 131–158.

Bertrand, M. and Schoar, A. (2003). Managing with Style: The Effect of Managers on Firm Policies. *Quarterly Journal of Economics*, 118(4), 1169–1208.

Betts, J. R. and McFarland, L. L. (1995). Safe Port in a Storm: The Impact of Labor Market Conditions on Community College Enrollments. *Journal of Human Resources*, 30(4), 741–765.

Black, D., McKinnish, T., and Sanders, S. (2005). Tight Labor Markets and the Demand for Education: Evidence from the Coal Boom and Bust. *Industrial and Labor Relations Review*, 59(1), 3–16.

Blanchard, O. J., & Katz, L. F. (1992). Regional Evolutions. Brookings Papers on Economic Activity, 1992(1), 1. https://doi.org/10.2307/2534556

Bloom, N., Handley, K., Kurmann, A., and Luck, P. A. (2024). The China Shock Revisited: Job Reallocation and Industry Switching in U.S. Labor Markets. NBER Working Paper No. 33098.

Borusyak, K., and Hull, P. (2023). Non-Random Exposure to Exogenous Shocks. *Econometrica*, 91(6), 2155–2185.

Borusyak, K., Hull, P., and Jaravel, X. (2022). Quasi-Experimental Shift-Share Research Designs. *Review of Economic Studies*, 89(1), 181–213.

Bound, J., & Turner, S. (2007). Cohort crowding: How resources affect collegiate attainment. Journal of Public Economics, 91(5–6), 877–899. https://doi.org/10.1016/j.jpubeco.2006.07.006

Bound, J., Lovenheim, M. F., & Turner, S. (2010). Why Have College Completion Rates Declined? An Analysis of Changing Student Preparation and Collegiate Resources. American Economic Journal: Applied Economics, 2(3), 129–157. https://doi.org/10.1257/app.2.3.129

California Community Colleges Chancellor's Office (2026). Management Information Systems Data Mart [dataset]. Sacramento, CA. https://datamart.cccco.edu/.

Cascio, E. U. and Narayan, A. (2022). Who Needs a Fracking Education? The Educational Response to Low-Skill-Biased Technological Change. *ILR Review*, 75(1), 56–89.

Cellini, S. R. (2009). Crowded Colleges and College Crowd-Out: The Impact of Public Subsidies on the Two-Year College Market. *American Economic Journal: Economic Policy*, 1(2), 1–30.

Chakrabarti, R., Gorton, N., and Lovenheim, M. F. (2020). State Investment in Higher Education: Effects on Human Capital Formation, Student Debt, and Long-Term Financial Outcomes of Students. NBER Working Paper No. 27885.

Chan, H. R. and Zhou, Y. C. (2025). Charged Up: Impacts of Green Energy Transition on Local Labor Markets. Working paper, SSRN No. 5954499.

Charles, K. K., Hurst, E., & Notowidigdo, M. J. (2018). Housing Booms and Busts, Labor Market Opportunities, and College Attendance. American Economic Review, 108(10), 2947–2994. https://doi.org/10.1257/aer.20151604

Cohen, A. M. and Brawer, F. B. (2003). *The American Community College*. 4th ed. Jossey-Bass.

Cohodes, S. R., and Goodman, J. S. (2014). Merit Aid, College Quality, and College Completion: Massachusetts' Adams Scholarship as an In-Kind Subsidy. *American Economic Journal: Applied Economics*, 6(4), 251–285.

Conzelmann, J. G., Hemelt, S. W., Hershbein, B., Martin, S. M., Simon, A., and Stange, K. M. (2023). Skills, Majors, and Jobs: Does Higher Education Respond? NBER Working Paper No. 31572.

Corden, W. M. and Neary, J. P. (1982). Booming Sector and De-Industrialisation in a Small Open Economy. *Economic Journal*, 92(368), 825–848.

Crago, C. L. and Chernyakhovskiy, I. (2017). Are Policy Incentives for Solar Power Effective? Evidence from Residential Installations in the Northeast. *Journal of Environmental Economics and Management*, 81, 132–151.

Curtis, E. M. and 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(1), 202–237.

Curtis, E. M., O'Kane, L., and Park, R. J. (2024). Workers and the Green-Energy Transition: Evidence from 300 Million Job Transitions. *Environmental and Energy Policy and the Economy*, 5, 127–161.

Dal Bó, E., Finan, F., and Rossi, M. A. (2013). Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service. *Quarterly Journal of Economics*, 128(3), 1169–1218.

De Chaisemartin, C., & D'Haultfœuille, X. (2020). Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects. American Economic Review, 110(9), 2964–2996. https://doi.org/10.1257/aer.20181169

Deming, D. J. and Walters, C. R. (2017). The Impact of Price Caps and Spending Cuts on U.S. Postsecondary Attainment. NBER Working Paper No. 23736. https://www.nber.org/papers/w23736.

Eckert, F., Fort, T. C., Schott, P. K., & Yang, N. J. (2021). Imputing Missing Values in the US Census Bureau's County Business Patterns. NBER Working Paper No. 26632.

U.S. Energy Information Administration. (2013). *Technically Recoverable Shale Oil and Shale Gas Resources: An Assessment of 137 Shale Formations in 41 Countries Outside the United States*. Washington, DC: U.S. Department of Energy. https://www.eia.gov/analysis/studies/worldshalegas/.

Enverus (2025). U.S. Horizontal Well Data, Wells Producing from 2000 to July 2025 [proprietary dataset]. Retrieved from Enverus Prism (formerly DrillingInfo). https://prism.enverus.com.

Fabra, N., Gutiérrez, E., Lacuesta, A., and Ramos, R. (2024). Do Renewable Energy Investments Create Local Jobs? *Journal of Public Economics*, 239, 105212.

Falch, T. (2010). The Elasticity of Labor Supply at the Establishment Level. *Journal of Labor Economics*, 28(2), 237–266.

Federal Reserve Bank of Dallas (2020). Collapse in Oil and Gas Activity, Employment Deepens. *Dallas Fed Energy Survey*, Second Quarter 2020. https://www.dallasfed.org/research/surveys/des/2020/2002.

Federal Reserve Bank of Dallas (2024). *Permian Basin Economic Indicators*, Third Quarter 2024. https://www.dallasfed.org/research/indicators/pb/2024/pb2403.

Feyrer, J., Mansur, E. T., and Sacerdote, B. (2017). Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution. *American Economic Review*, 107(4), 1313–1334.

Flood, S., King, M., Rodgers, R., Ruggles, S., Warren, J. R., Backman, D., Breton, E., Cooper, G., Rivera Drew, J. A., Richards, S., and Van Riper, D. (2025). Integrated Public Use Microdata Series, Current Population Survey: Version 13.0 [dataset]. Minneapolis, MN: IPUMS. https://doi.org/10.18128/D030.V13.0.

Foote, A. and Grosz, M. (2020). The Effect of Local Labor Market Downturns on Postsecondary Enrollment and Program Choice. *Education Finance and Policy*, 15(4), 593–622.

Foote, A. and Stange, K. M. (2022). Attrition from Administrative Data: Problems and Solutions with an Application to Postsecondary Education. NBER Working Paper No. 30232.

Gagliardi, L., Moretti, E., and Serafinelli, M. (2023). The World's Rust Belts: The Heterogeneous Effects of Deindustrialization on 1,993 Cities in Six Countries. NBER Working Paper No. 31948.

Gagliarducci, S. and Nannicini, T. (2013). Do Better Paid Politicians Perform Better? Disentangling Incentives from Selection. *Journal of the European Economic Association*, 11(2), 369–398.

Gardner, J. (2021). Two-Stage Differences in Differences. Working Paper.

Gilmer, R. W., Hernandez, R., and Phillips, K. R. (2012). Oil Boom in Eagle Ford Shale Brings New Wealth to South Texas. *Southwest Economy*, Second Quarter 2012. Federal Reserve Bank of Dallas.

Goldin, C. and Katz, L. F. (2008). *The Race between Education and Technology*. Harvard University Press.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik Instruments: What, When, Why, and How. *American Economic Review*, 110(8), 2586–2624.

Goodman, J. and Winkelmann, J. (2025). Labor Market Strength and Declining Community College Enrollment. NBER Working Paper No. 34498.

Goodman-Bacon, A. (2021). Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics*, 225(2), 254–277.

Greenspon, J. and Raimi, D. (2025). Finding the Right Fit: What Jobs Offer a Good Match for Fossil Fuel Workers' Skills? Working Paper 25-06, Resources for the Future.

Grosz, M. (2020). The Returns to a Large Community College Program: Evidence from Admissions Lotteries. *American Economic Journal: Economic Policy*, 12(1).

Grosz, M. (2022). Do Postsecondary Training Programs Respond to Changes in the Labor Market? *Journal of Human Capital*, 16(4).

Gylfason, T. (2001). Natural resources, education, and economic development. *European Economic Review*, 45(4–6), 847–859.

Hanson, G. H. (2023). Local Labor Market Impacts of the Energy Transition. NBER Working Paper No. 30871.

Hemelt, S. W., Stange, K. M., Furquim, F., Simon, A., and Sawyer, J. E. (2021). Why Is Math Cheaper than English? Understanding Cost Differences in Higher Education. *Journal of Labor Economics*, 39(2), 397–435. https://doi.org/10.1086/709535

Hillman, N. and Weichman, T. (2016). Geography of College Opportunity. *American Educational Research Journal*, 53(4), 987–1021.

Holzer, H., Feygin, A., Bailey, P., and Rutschow, E. Z. (2026). Creating Credentials of Value in Community Colleges: Correlates of High Performance. American Institutes for Research.

Hoxby, C. M. and Leigh, A. (2004). Pulled Away or Pushed Out? Explaining the Decline of Teacher Aptitude in the United States. *American Economic Review*, 94(2), 236–240.

Huang, R. and Kahn, M. E. (2024). Do Red States Have a Comparative Advantage in Generating Green Power? *Environmental and Energy Policy and the Economy*, 5, 200–238.

Imbens, G. W., and Angrist, J. D. (1994). Identification and Estimation of Local Average Treatment Effects. *Econometrica*, 62(2), 467–475.

InformEd States (n.d.). State Higher Education Policy Database. https://informedstates.org/data.

Internal Revenue Service (IRS) (2024). *Inflation Reduction Act Prevailing Wage and Registered Apprenticeship Overview*. Publication 5855, revised June 2024. https://www.irs.gov/pub/irs-pdf/p5855.pdf.

International Energy Agency (2025). *World Energy Employment 2025*. IEA, Paris.

Interstate Renewable Energy Council (IREC) (2024). *National Solar Jobs Census 2023*. https://irecusa.org/programs/solar-jobs-census-2023/.

Interstate Renewable Energy Council (IREC) (n.d.). Solar O&M Technician (Commercial/Utility). *Solar Career Map*. https://www.irecsolarcareermap.org/jobs/solar-o-and-m-technician-commercial-utility.

Jacobson, L., LaLonde, R., and Sullivan, D. G. (2005). Estimating the Returns to Community College Schooling for Displaced Workers. *Journal of Econometrics*, 125(1–2), 271–304.

Kahn, L. B. (2010). The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy. *Labour Economics*, 17(2), 303–316.

Kane, T. J., & Rouse, C. E. (1999). The Community College: Educating Students at the Margin Between College and Work. Journal of Economic Perspectives, 13(1), 63–84. https://doi.org/10.1257/jep.13.1.63

Katovich, E., Parker, D., and Poelhekke, S. (2026). Timing Is Everything: Labor Market Winners and Losers During Boom-Bust Cycles. *Journal of the Association of Environmental and Resource Economists*, 13(1), 229–260.

Katz, L. F. and Murphy, K. M. (1992). Changes in Relative Wages, 1963–1987: Supply and Demand Factors. *Quarterly Journal of Economics*, 107(1), 35–78.

Kovalenko, A. (2023). Natural Resource Booms, Human Capital, and Earnings: Evidence from Linked Education and Employment Records. *American Economic Journal: Applied Economics*, 15(2), 184–217.

Krause, E. (2025a). Selective Migration and Regional Decline: Evidence from Coal Country. Working Paper.

Krause, E. (2025b). Adjusting to the Energy Transition: Training and Transfers in Coal Country. *Journal of Environmental Economics and Management*.

Kurlaender, M., Carrell, S., and Jackson, J. (2016). The Promises and Pitfalls of Measuring Community College Quality. *RSF: The Russell Sage Foundation Journal of the Social Sciences*, 2(1), 174.

Lawrence Berkeley National Laboratory (2025a). RPS & CES Targets and Demand [data file], August 2025 version. Berkeley, CA. https://emp.lbl.gov/projects/renewables-portfolio.

Lawrence Berkeley National Laboratory (2025b). *Tracking the Sun* Public Data File, 2025 Edition. Released September 29, 2025. Berkeley, CA. https://emp.lbl.gov/tracking-the-sun.

Leigh, R., Thompson, J., and Torres, L. (2026). Oil Prices Are Up; Whither the Texas Boom? *Southwest Economy*. Federal Reserve Bank of Dallas. https://www.dallasfed.org/research/swe/2026/swe2614.

Leung, P. and Pei, Z. (2020). Further Education during Unemployment. Working Paper 642, Industrial Relations Section, Princeton University.

Lim, J., Aklin, M., and Frank, M. R. (2023). Location Is a Major Barrier for Transferring US Fossil Fuel Employment to Green Jobs. *Nature Communications*, 14, 5711.

Long, M. C., Goldhaber, D., and Huntington-Klein, N. (2015). Do Completed College Majors Respond to Changes in Wages? *Economics of Education Review*, 49(C), 1–14.

Maniloff, P. and Mastromonaco, R. (2014). The Local Economic Impacts of Hydraulic Fracturing and Determinants of Dutch Disease. Working Paper.

Marchand, J. and Weber, J. G. (2020). How Local Economic Conditions Affect School Finances, Teacher Quality, and Student Achievement: Evidence from the Texas Shale Boom. *Journal of Policy Analysis and Management*, 39(1), 36–63.

McCormack, J., Propper, C., and Smith, S. (2014). Herding Cats? Management and University Performance. *Economic Journal*, 124(578), F534–F564.

Michaels, G. (2011). The Long Term Consequences of Resource-Based Specialisation. *Economic Journal*, 121(551), 31–57.

Montiel Olea, J. L., and Pflueger, C. (2013). A Robust Test for Weak Instruments. *Journal of Business & Economic Statistics*, 31(3), 358–369.

Morissette, R., Chan, P. C. W., and Lu, Y. (2015). Wages, Youth Employment, and School Enrollment: Recent Evidence from Increases in World Oil Prices. *Journal of Human Resources*, 50(1), 222–253.

Murnane, R. J., Willett, J. B., McDonough, A., Papay, J. P., and Mantil, A. (2024). Are Community-College Students Increasingly Choosing High-Paying Fields of Study? Evidence from Massachusetts. *AERA Open*, 10, 23328584241304696. Also NBER Working Paper No. 33073. https://doi.org/10.1177/23328584241304696.

National Center for Education Statistics (2023). Digest of Education Statistics, Table 316.80. Percentage of degree-granting postsecondary institutions with a tenure system and percentage of full-time faculty with tenure at these institutions, by control and level of institution and selected characteristics of faculty: Selected academic years, 1993–94 through 2021–22. U.S. Department of Education, Institute of Education Sciences.

National Center for Education Statistics (2024). Integrated Postsecondary Education Data System (IPEDS): Finance Survey Component. U.S. Department of Education, Institute of Education Sciences.

National Center for Education Statistics (2025a). Digest of Education Statistics, Table 303.70. Total undergraduate fall enrollment in degree-granting postsecondary institutions. U.S. Department of Education, Institute of Education Sciences.

National Center for Education Statistics (2025b). Integrated Postsecondary Education Data System (IPEDS), 1990–2024 [dataset]. U.S. Department of Education, Institute of Education Sciences. https://nces.ed.gov/ipeds/.

National Science Foundation (2026). NSF Award Search, Awards by Year, 1990–1995 [dataset]. https://www.nsf.gov/awardsearch/download.jsp.

Notowidigdo, M. J. (2020). The Incidence of Local Labor Demand Shocks. Journal of Labor Economics, 38(3), 687–725. https://doi.org/10.1086/706048

Oreopoulos, P., von Wachter, T., and Heisz, A. (2012). The Short- and Long-Term Career Effects of Graduating in a Recession. *American Economic Journal: Applied Economics*, 4(1), 1–29.

Our World in Data (2026). Solar Photovoltaic Module Price [dataset]. Data from IRENA, *Renewable Power Generation Costs in 2024* (2025); Nemet (2009); and Farmer and Lafond (2016), with major processing by Our World in Data. Retrieved August 17, 2026, from https://ourworldindata.org/grapher/solar-pv-prices.

Pless, J. and van Benthem, A. A. (2019). Pass-Through as a Test for Market Power: An Application to Solar Subsidies. *American Economic Journal: Applied Economics*, 11(4), 367–401.

Popp, D., Vona, F., Marin, G., and Chen, Z. (2021). The Employment Impact of a Green Fiscal Push: Evidence from the American Recovery and Reinvestment Act. *Brookings Papers on Economic Activity*, Fall 2021, 1–69.

Ruggles, S., Flood, S., Sobek, M., Backman, D., Cooper, G., Rivera Drew, J. A., Richards, S., Rodgers, R., Schroeder, J., and Williams, K. (2025). IPUMS USA: Version 16.0 [dataset]. Minneapolis, MN: IPUMS. https://doi.org/10.18128/D010.V16.0.

Sachs, J. D. and Warner, A. M. (2001). The Curse of Natural Resources. *European Economic Review*, 45(4–6), 827–838.

Schanzenbach, D. W. and Turner, S. (2022). Limited Supply and Lagging Enrollment: Production Technologies and Enrollment Changes at Community Colleges during the Pandemic. *Journal of Public Economics*, 212.

Sengupta, M., Xie, Y., Lopez, A., Habte, A., Maclaurin, G., and Shelby, J. (2018). The National Solar Radiation Data Base (NSRDB). *Renewable and Sustainable Energy Reviews*, 89, 51–60.

Sentz, R., Metsker-Galarza, N., and Leininger, P. (2018). Where Value Meets Values: The Economic Impact of Community Colleges. Emsi.

Slijk, C. and Phillips, K. R. (2021). Once-Oil-Dependent Texas Economy to Keep Growing as Renewable Energy Expands. *Southwest Economy*, Third Quarter 2021. Federal Reserve Bank of Dallas.

Stange, K. (2012). Ability Sorting and the Importance of College Quality to Student Achievement: Evidence from Community Colleges. *Education Finance and Policy*, 7(1), 74–105.

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

Surveillance, Epidemiology, and End Results (SEER) Program (2026). U.S. County Population Data, 1990–2024 [dataset]. National Cancer Institute, DCCPS, Surveillance Research Program. Released February 2026. https://seer.cancer.gov/popdata/.

Thompson, J. (2022). Shale Oil Boom Gave Permian Basin a Second Life. *Southwest Economy*, First Quarter 2022. Federal Reserve Bank of Dallas.

U.S. Bureau of Economic Analysis (2023). Table CAEMP25N, Total Full-Time and Part-Time Employment by NAICS Industry, County Level, 2001–2022 [dataset]. Regional Economic Accounts. Last updated November 16, 2023. https://apps.bea.gov/regional/downloadzip.htm.

U.S. Bureau of Labor Statistics (2026a). Current Employment Statistics, State and Area: All Employees, Oil and Gas Extraction and Support Activities for Mining, Texas. Retrieved from FRED, Federal Reserve Bank of St. Louis, series SMU48000001021100001 and SMU48000001021300001.

U.S. Bureau of Labor Statistics (2026b). Electricians. *Occupational Outlook Handbook*. Last modified August 27, 2026. https://www.bls.gov/ooh/construction-and-extraction/electricians.htm.

U.S. Bureau of Labor Statistics (2026c). Solar Photovoltaic Installers. *Occupational Outlook Handbook*. Last modified August 27, 2026. https://www.bls.gov/ooh/construction-and-extraction/solar-photovoltaic-installers.htm.

U.S. Department of Energy (2025). *2025 U.S. Energy and Employment Report*. Washington, DC.

U.S. Energy Information Administration (2011). Maps: Exploration, Resources, Reserves, and Production. https://www.eia.gov/maps/maps.htm.

U.S. Energy Information Administration (2025a). *Annual Energy Outlook 2025*. U.S. Department of Energy.

U.S. Energy Information Administration (2025b). *Annual Electric Generator Report (Form EIA-860)*, 2024 data. Washington, DC. https://www.eia.gov/electricity/data/eia860/.

Van der Ploeg, F. (2011). Natural Resources: Curse or Blessing? *Journal of Economic Literature*, 49(2), 366–420.

Vona, F., Marin, G., and Consoli, D. (2019). Measures, Drivers and Effects of Green Employment: Evidence from US Local Labor Markets, 2006–2014. *Journal of Economic Geography*, 19(5), 1021–1048.

Weber, J. G. (2014). A Decade of Natural Gas Development: The Makings of a Resource Curse? *Resource and Energy Economics*, 37, 168–183.

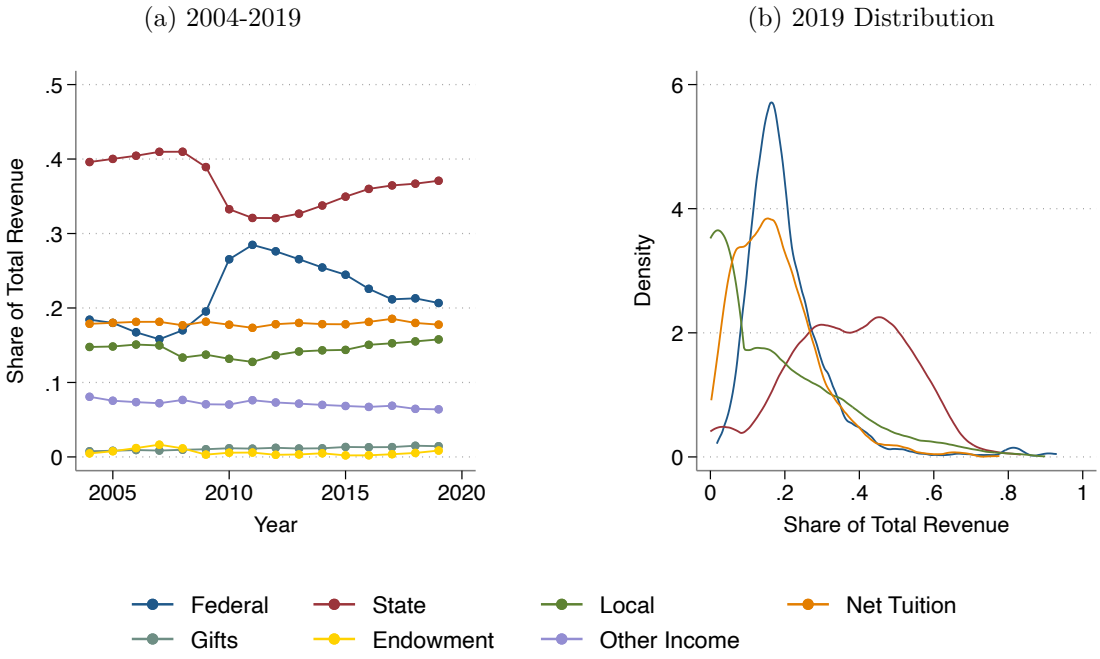
Wing, C., Freedman, S. M., and Hollingsworth, A. (2024). Stacked Difference-in-Differences. NBER Working Paper No. 32054.

Weinstein, R. (2022). Local Labor Markets and Human Capital Investments. *Journal of Human Resources*, 57(5), 1498–1525.

Zuo, N., Schieffer, J., and Buck, S. (2019). The Effect of the Oil and Gas Boom on Schooling Decisions in the U.S. *Resource and Energy Economics*, 55, 1–23.

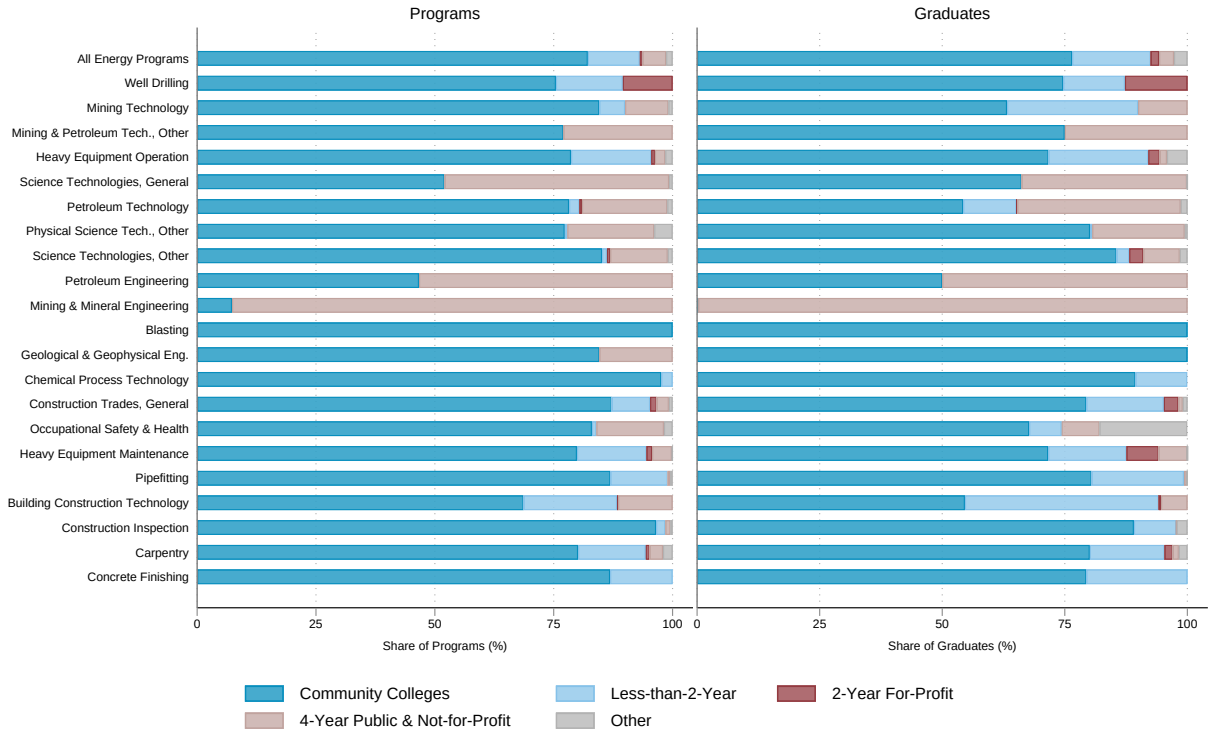
A Appendix Figures

Figure A1: Revenue Composition for Two-Year Colleges



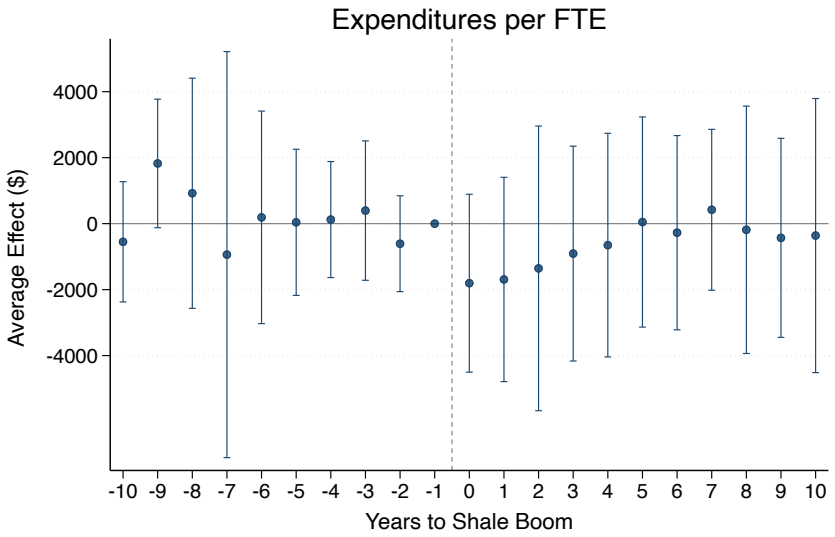
Notes. This figure shows the composition of revenue for two-year colleges in the IPEDS Finance Survey. Panel (a) plots the time series of average revenue shares from 2004 to 2019. Panel (b) plots the distribution of revenue shares across colleges in 2019. Revenue sources include tuition and fees, state appropriations, local appropriations, federal appropriations, net tuition, gifts, endowments, and other sources.

Figure A2: Energy Programs and Graduates by Institution Type and CIP Code



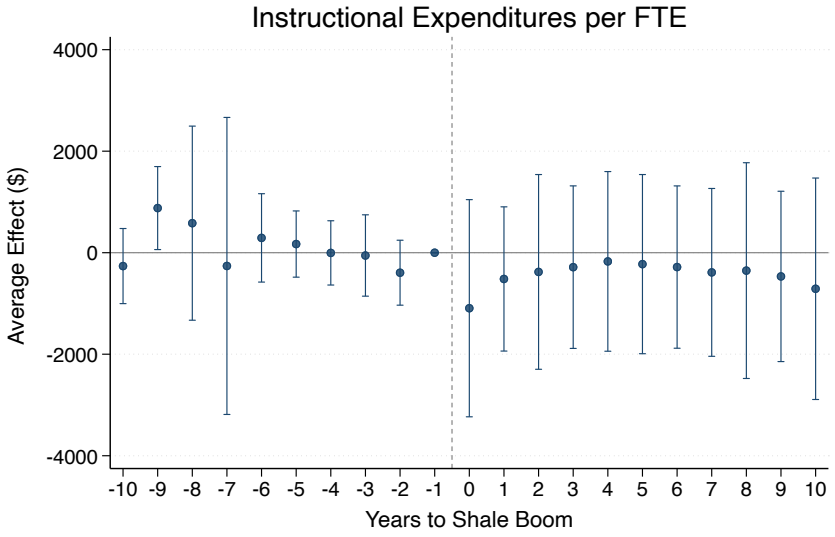
Notes. This figure shows the share of programs (left) and graduates (right) by institution type for certificates and associate degrees in each energy CIP code listed in Table 2, ordered by the share of graduates expected to enter the energy sector. The top row pools all energy programs, as in Figure 5. A program is a CIP code and award level reported by an institution in a year. Shares are computed in each year and averaged over the years a CIP code is offered between 1990 and 2024. Community colleges are public two-year institutions, and Other includes four-year for-profit and two-year not-for-profit institutions. Data are from NCES IPEDS Completions.

Figure A3: Shale Effects on Total Expenditures Over Time



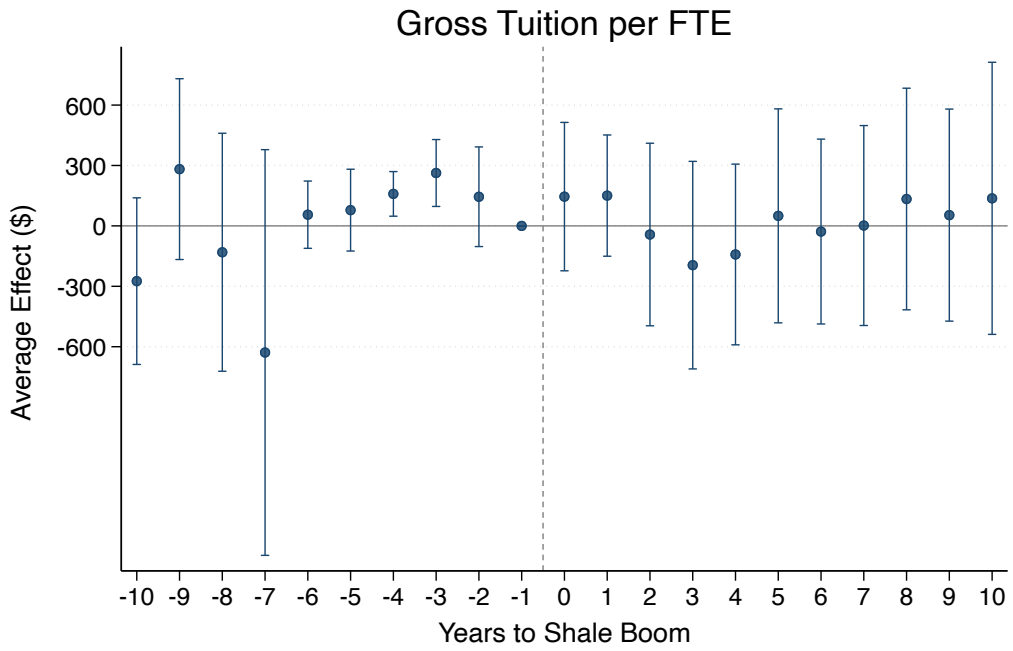
Notes. This figure plots the Gardner (2021) two-stage event study estimates for total expenditures per FTE. Standard errors are bootstrapped and clustered at the play level.

Figure A4: Shale Effects on Instructional Expenditures Over Time



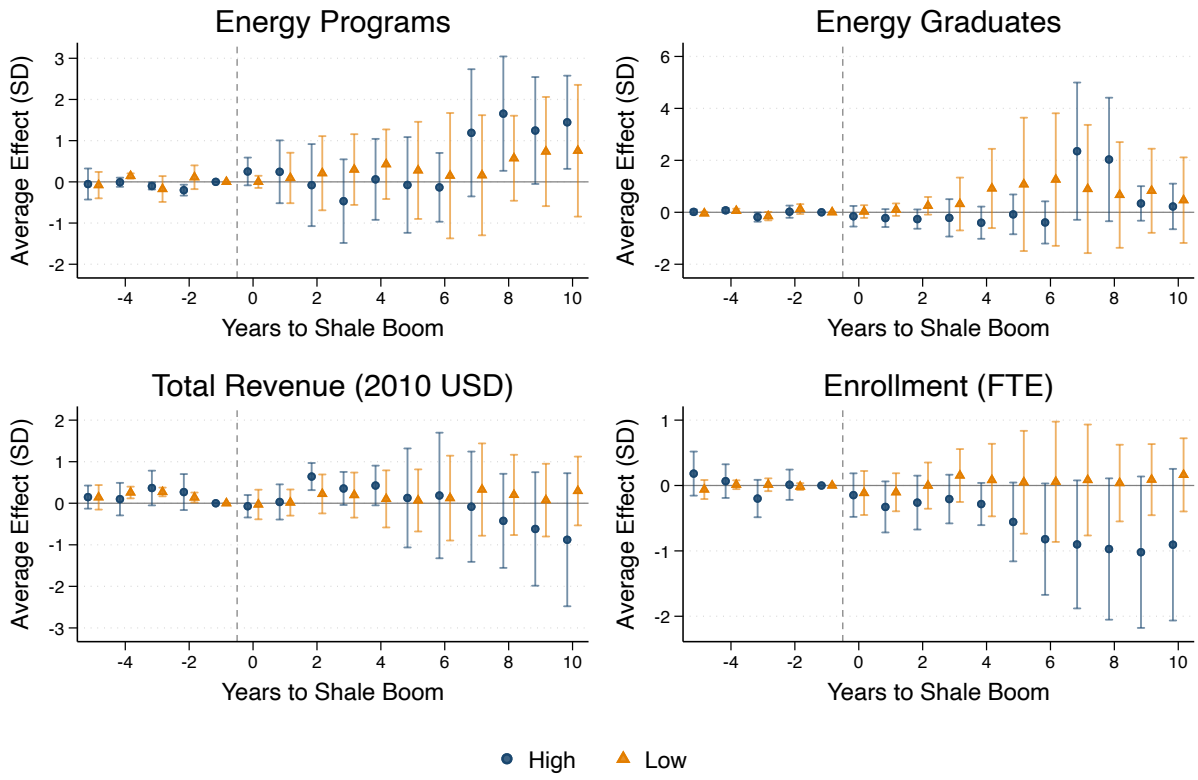
Notes. This figure plots the Gardner (2021) two-stage event study estimates for instructional expenditures per FTE. Standard errors are bootstrapped and clustered at the play level.

Figure A5: Shale Effects on Gross Tuition Over Time



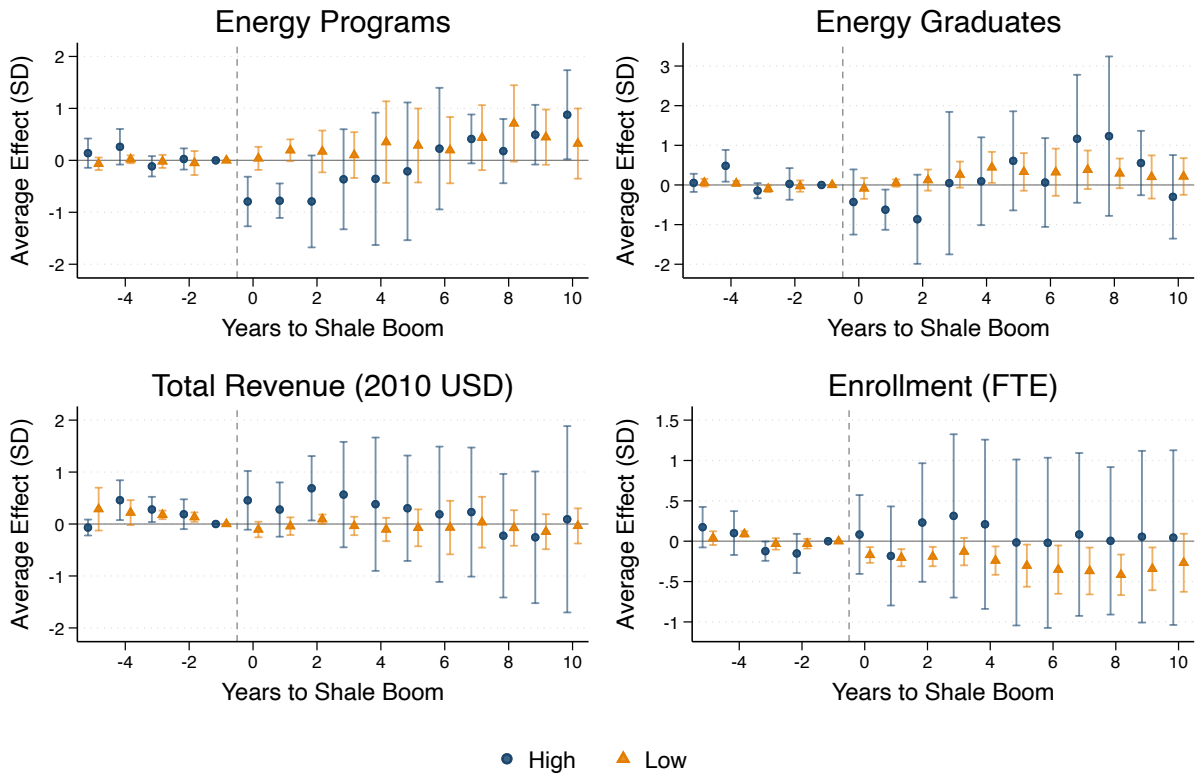
Notes. This figure plots the Gardner (2021) two-stage event study estimates for gross tuition per FTE. Standard errors are bootstrapped and clustered at the play level.

Figure A6: Shale Effects Over Time by Pre-Boom Instructional Spending per FTE



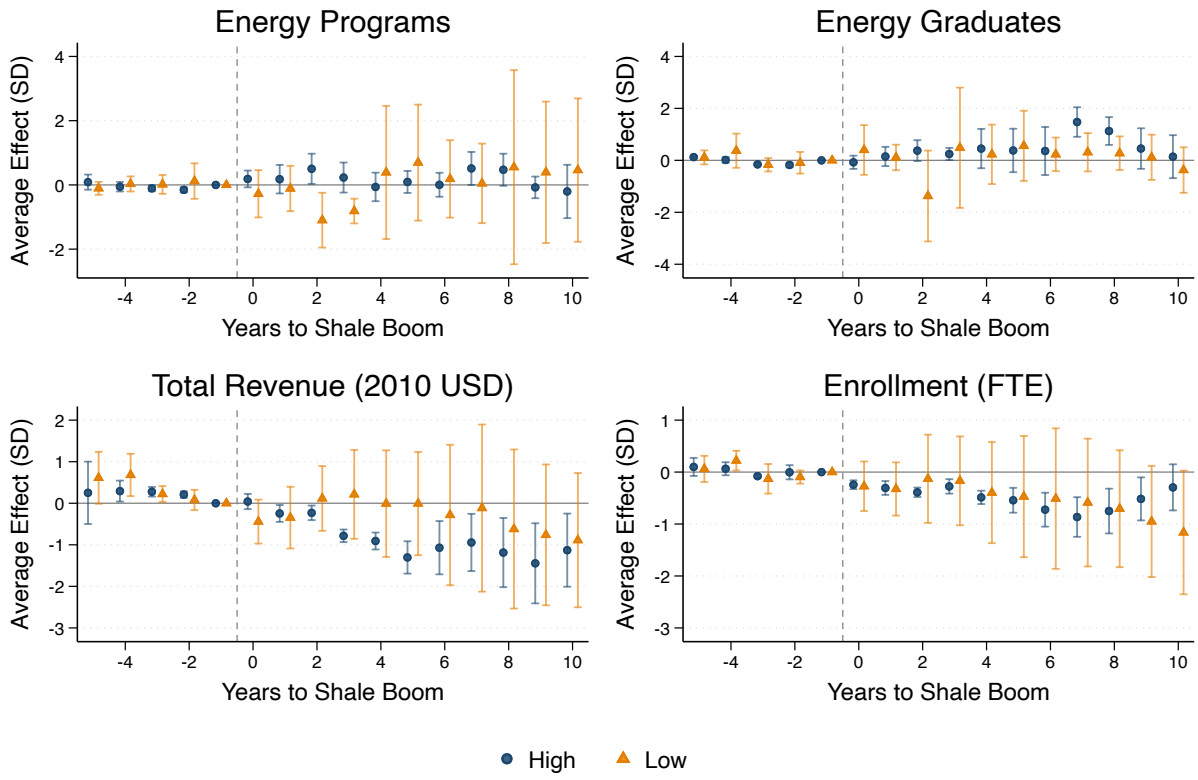
Notes. This figure plots Gardner (2021) two-stage event study estimates separately for colleges in the top quartile (High) and bottom quartile (Low) of pre-boom instructional spending per FTE student, for energy programs, energy graduates, total revenue, and FTE enrollment. Colleges in the middle two quartiles are excluded. Year effects vary with pre-boom enrollment. Outcomes are expressed in standard deviations of their annual change. Standard errors are clustered at the play level.

Figure A7: Shale Effects Over Time by Prior NSF Grants



Notes. This figure plots Gardner (2021) two-stage event study estimates separately for colleges that received any NSF grant in 1990–1995 (High) and those that did not (Low), for energy programs, energy graduates, total revenue, and FTE enrollment. Year effects vary with pre-boom enrollment. Outcomes are expressed in standard deviations of their annual change. Standard errors are clustered at the play level.

Figure A8: Shale Effects Over Time by Pre-Boom Faculty Pay



Notes. This figure plots Gardner (2021) two-stage event study estimates separately for colleges in the top quartile (High) and bottom quartile (Low) of pre-boom faculty pay relative to county per capita income, for energy programs, energy graduates, total revenue, and FTE enrollment. Colleges in the middle two quartiles are excluded. Year effects vary with pre-boom enrollment. Outcomes are expressed in standard deviations of their annual change. Standard errors are clustered at the play level.

B Appendix Tables

Table B1: Shale Treatment Breakdown

Basin	Play	Boom Year	10 Comp Wells Year	Treated Colleges	Control Colleges
(1)	(2)	(3)	(4)	(5)	(6)
Ft. Worth	Barnett	2001	2005	7	60
TX-LA-MS Salt	Tuscaloosa*	2001	2001	8	100
Arkoma	Fayetteville	2005	2006	2	12
Permian	Delaware	2005	2001	1	71
Permian	Wolfcamp-Midland	2005	2010	3	55
Arkoma	Woodford-Caney	2006	2006	1	32
Appalachian	Devonian*	2007	2007	37	104
Ardmore	Woodford	2007	2008	1	86
Williston	Bakken	2007	2001	3	16
Anadarko	Woodford	2008	2008	3	28
Appalachian	Marcellus	2008	2008	10	109
TX-LA-MS Salt	Haynesville-Bossier	2008	2000	8	73
Western Gulf	Eagle Ford	2009	2001	1	50
Denver	Niobrara	2010	2010	8	20
Powder River	Niobrara	2010	2011	1	13
Appalachian	Utica	2012	2011	5	73
Greater Green River	Hilliard-Baxter-Mancos-Niobrara	2012	2007	1	17
N				100	322

Notes. This table reports the details of the shale boom treatment. Columns 1 and 2 report the shale basin and play. Column 3 reports the treatment year used for the main analysis - the first year shale became public in line with Bartik et al. (2019) or the first year with 10 completed wells. * denotes plays that use 10 completed wells as the treatment date. Column 4 reports the first year a play reached 10 completed wells for comparison. Columns 5 and 6 report the number of treated and control community colleges for each state.

Table B2: Pre-Boom Summary Statistics and Balance

	Treated (1)	Control (2)	Norm. Diff. (3)
Panel A. Outcomes			
Energy Programs	0.53	0.47	0.08
Non-Energy Programs	26.3	30.2	-0.19*
Energy Completions	4.69	3.21	0.15
Non-Energy Completions	316	371	-0.17
Panel B. Institutional Characteristics			
FTE Enrollment	2,167	2,827	-0.24*
% Part-time	49.5	53.8	-0.27**
% Black	9.69	12.3	-0.19
% Hispanic	5.27	8.96	-0.27**
% Female	58.0	58.2	-0.02
Metro Indicator	0.58	0.55	0.06
N	100	322	

Notes. This table reports summary statistics for pre-boom community colleges. Columns 1 and 2 report pre-boom means for treated and control community colleges, respectively. Column 3 reports the normalized differences between the two. $^{***}p < 0.01,^{**}p < 0.05,^{*}p < 0.10$

Table B3: Summary Statistics for Energy Completers by Cohort

	Mean				Normalized Difference (vs. 1994–1997)		
Cohorts:	1994–1997	1998–2002	2003–2007	2008–2011	1998–2002	2003–2007	2008–2011
Test Scores:							
Grade 8 Math	0.159	0.098	0.008	0.071	-0.077	-0.194	-0.113
Grade 8 Reading	0.209	0.186	0.099	0.003	-0.031	-0.142	-0.262**
Demographics:							
Female	0.244	0.262	0.256	0.199	0.042	0.029	-0.106
Black or Other	0.115	0.070	0.080	0.058	-0.158	-0.120	-0.204*
Hispanic	0.244	0.310	0.383	0.530	0.149	0.303**	0.614***
White	0.641	0.620	0.537	0.412	-0.043	-0.211*	-0.470***
Language, Disability, or Immigrant Status	0.103	0.059	0.050	0.131	-0.160	-0.199*	0.089
Economically Disadvantaged	0.282	0.348	0.408	0.424	0.141	0.266**	0.300**
Gifted	0.090	0.102	0.097	0.083	0.040	0.025	-0.023
Vocational Education	0.410	0.203	0.264	0.283	-0.458***	-0.313***	-0.269**
N	78	187	402	396			

Notes. This table reports summary statistics for energy completers in the Texas ERC data by eighth-grade cohort. Columns 1 through 4 report means for observable characteristics by cohort. Columns 5-7 report the normalized differences between the means of the reported cohort and the 1994–1997 cohort. $^{***}p < 0.01,^{**}p < 0.05,^{*}p < 0.10$

Table B4: First-Stage and Reduced Form Estimates for National Solar Investment

	First Stage	Reduced Form	
Outcome:	Solar Investment$_{t-2}$	Δ Solar CIPs	Δ Solar Completions
Raw Outcome:			
Solar Potential ×State RPS$_{t-3}$ (SD)	108***	0.002*	0.013**
	(6.49)	(0.001)	(0.007)
Solar Potential ×Solar Module Price$_{t-3}$ (SD)	-39***	0.000	-0.008
	(9.49)	(0.002)	(0.011)
Effect per SD:			
Solar Potential ×State RPS$_{t-3}$ (SD)	0.74***	0.010*	0.005**
	(0.04)	(0.006)	(0.002)
Solar Potential ×Solar Module Price$_{t-3}$ (SD)	-0.27***	0.000	-0.003
	(0.07)	(0.009)	(0.004)
Year FE	Y	Y	Y
State FE	Y	Y	Y
Population Control	Y	Y	Y
Joint F-Stat of Instruments	1144	5.630	7.160
Effective F	173		
Anderson-Rubin p-value (Joint Test)		0.004	0.001
Observations	24,268	24,268	24,268
Sample Mean	34	0.184	0.653
SD of Dep. Variable	146	0.240	2.720

Notes. Column 1 reports the first stage, where the outcome is annual megawatts of solar capacity installed in the commuting zone. Columns 2 and 3 report the reduced form, where the outcomes are the annual changes in solar programs and solar completions. The instruments interact local solar potential (irradiance × land area) with state RPS targets and with the global solar module price, and are expressed in standard deviations. The top panel reports estimates in the units of each outcome, and the bottom panel reports them in standard deviations of the dependent variable. The specification includes population controls and year and state fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the commuting-zone level in parentheses.

Table B5: First-Stage and Reduced Form Estimates for California Solar Investment

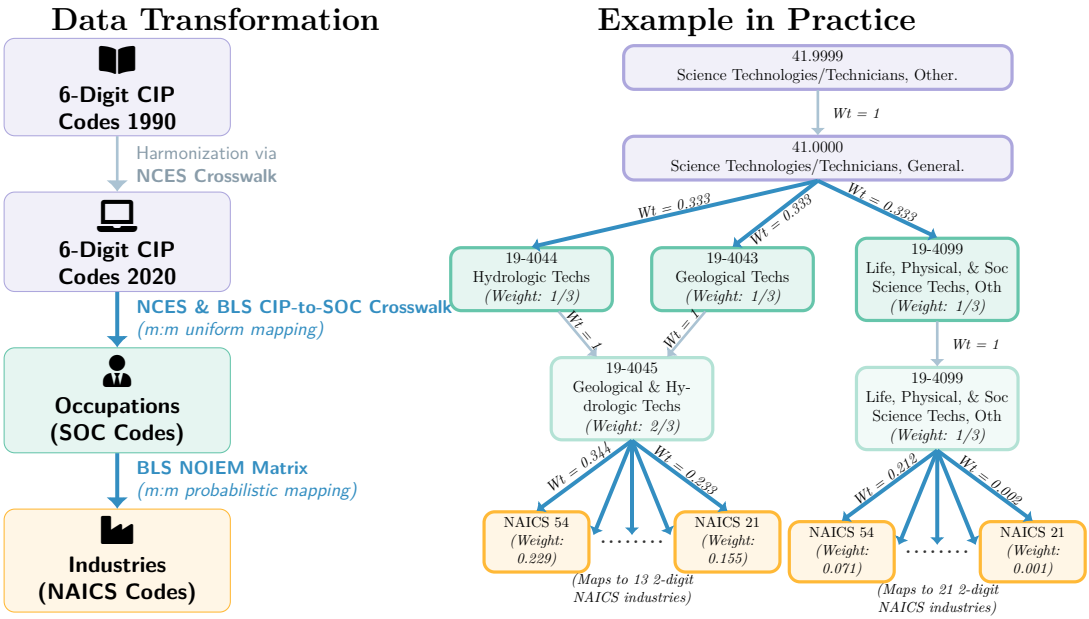
	First Stage	Reduced Form		
Outcome:	Solar Investment$_{t-1}$	Δ Share Solar Sections	Δ Solar Sections	Δ Non-Solar Sections
Raw Outcome:				
Solar Potential ×State RPS$_{t-2}$ (SD)	42***	0.002***	0.017***	0.465
	(14)	(0.001)	(0.004)	(1.501)
Effect per SD:				
Solar Potential ×State RPS$_{t-2}$ (SD)	0.32***	0.019***	0.025***	0.005
	(0.11)	(0.005)	(0.006)	(0.015)
Year FE	Y	Y	Y	Y
CZ FE	Y	Y	Y	Y
Population Control	Y	Y	Y	Y
Joint F-Stat of Instruments	8.67	13	16	0.096
Anderson-Rubin p-value (Joint Test)		0.001	0.000	0.759
Observations	2,567	2,567	2,567	2,567
Sample Mean	65	0.062	0.427	679
SD of Dep. Variable	131	0.120	0.670	98

Notes. Column 1 reports the first stage, where the outcome is megawatts of solar capacity installed in the California county. Columns 2 through 4 report the reduced form, where the outcomes are the annual changes in the share of solar sections, solar sections, and non-solar sections. The instrument is county solar potential interacted with California's RPS targets, expressed in standard deviations. The top panel reports estimates in the units of each outcome, and the bottom panel reports them in standard deviations of the dependent variable. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the county level in parentheses.

C Crosswalk Construction

Figure C1 summarizes the construction of the CIP-to-industry mapping. The left panel shows the data-transformation pipeline: 1990-vintage six-digit CIP codes are harmonized to the 2020 vintage using the NCES CIP crosswalk; the 2020 CIPs are then mapped to Standard Occupational Classification (SOC) codes via the NCES/BLS CIP-to-SOC crosswalk (a many-to-many uniform mapping in which a CIP is split equally across the SOCs it lists); and SOCs are finally mapped to two-digit NAICS industries using the BLS National Occupation-Industry Employment Matrix (NOIEM, a many-to-many mapping with weights equal to the employment share of each occupation in each industry). The right panel illustrates the procedure for two example CIPs in the 41 (Science Technologies) family, showing the equal-weight expansion across SOCs and the probabilistic expansion across NAICS industries that follows.

Figure C1: CIP-to-Industry Mapping



D National Earnings Evidence

This appendix complements the earnings analysis in Section 6 with national evidence from the PSEO on graduate earnings. The PSEO tabulations cover community colleges across participating states, so they provide a check on whether the earnings patterns documented in Texas extend beyond a single state. The data are at the level of the institution rather than individual records, so they represent a coarser measure. One issue in particular is that the composition of energy graduates likely changes as the boom draws new students into the field, so selection differences across cohorts confound changes in the returns to energy training with changes in who receives it. I therefore interpret this analysis with caution as a descriptive complement to my findings on Texas rather than as independent estimates of returns.

D.1 Empirical Design

I use the Census Bureau's PSEO tabulations (as in Acton et al., 2026). Produced by the Longitudinal Employer–Household Dynamics (LEHD) program, the PSEO links transcript records from participating institutions and state systems to the national database of unemployment-insurance (UI) earnings records. The tabulations report median earnings at one, five, and ten years after graduation, along with employment, by institution, degree level, field of study (4-digit CIP code), and graduation cohort. Since the tabulations are

disclosure protected, small cells are suppressed. Because the underlying transcript data are contributed voluntarily, coverage is partial and uneven. Three features of the data govern what can be learned from it. First, the tabulations are thin at the community-college level. Of the colleges in my sample, 167 (22 treated) across 10 plays appear in PSEO with energy and non-energy graduates observed in both pre- and post-boom cohorts. Second, subbaccalaureate cohorts are pooled into five-year bands, which coarsens the timing of the boom relative to a graduate's cohort and limits how finely the dynamics can be measured. Third, the earnings medians are conditional on appearing in covered employment, so a level comparison of earnings could reflect differences in who is found working rather than differences in pay. I address this concern partially by examining employment rates alongside earnings.

My design holds institution quality fixed by comparing fields within the same college and cohort across treated institutions. For graduates of college i in cohort t at degree level m, I estimate

$$Y_{ictm} = \alpha_i + \lambda_t + \sum_e \gamma_e \, \mathbf{1}[\text{event} = e] \cdot \text{energy}_c + \varepsilon_{ictm}, \tag{16}$$

where Y_{ictm} is median earnings for field c at one of one, five or ten years post graduation, energy_c indicates an energy program, and event time is the years between a cohort's graduation and the play's boom onset. The institution and cohort fixed effects difference out the college's overall earnings level and any common cohort shock, so each γ_e is the energy–non-energy earnings gap among graduates of the same college who finished at the same time. Because this is a within-college comparison of two fields, rather than a comparison of treated and control colleges, γ_e reports how much more energy graduates earn than their non-energy classmates, not a deviation from a control group or omitted time period.

I visualize this gap in two ways, which answer two different questions. The first plots the gap against event time to assess whether the energy premium changed with the boom and how it differs across cohorts who trained at different points in the cycle. The second plots the gap by years since graduation, separately for cohorts who graduated before and after the boom. This addresses the question of durability more directly, tracing the energy premium through pre- and post-boom cohorts' careers.

D.2 Results

I report the within-college gap by event time, separately for certificates of less than one year and associate's degrees and at each earnings horizon[25]. Across all earnings years and

[25]The data for certificates of 1-2 years was too sparse to be included.

award levels, the data is noisy and relatively imprecise, but the gap seems to be positive, suggesting that energy graduates out-earn their non-energy classmates at the same college and cohort. Furthermore, the gap does not seem to dissipate over time; it is, if anything, as large for cohorts who trained well after the boom as for those who trained before it, giving no indication that the premium is an artifact of earnings at the boom's peak. Individual event-time coefficients are imprecise and confidence intervals often include zero, so I do not analyze the year-to-year movements.

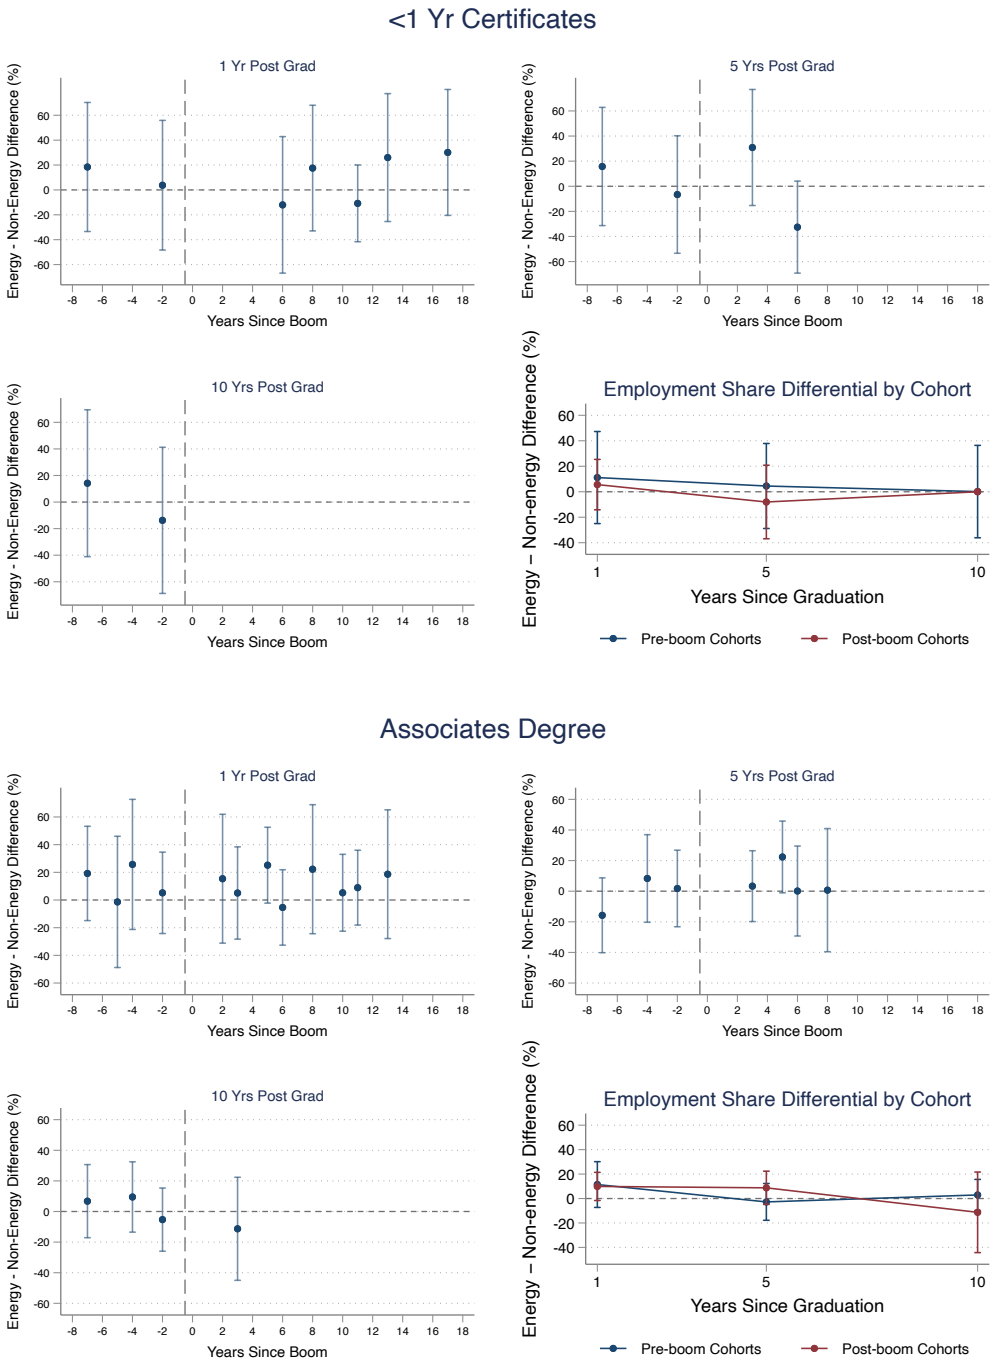
Another way to visualize the results is to look at the development of earnings over time by cohort (pre and post boom). Looking at earnings by event time within cohorts addresses this by tracing the energy–non-energy gap over the career of graduates separately for pre-boom and post-boom cohorts. For pre-boom cohorts, the premium is flat or modestly declining as the career progresses. For post-boom cohorts who were trained during and after the boom and are disproportionately graduates of the new programs documented in Section 4.5, the premium does not collapse and potentially increases. The within-college energy premium for boom-trained graduates seems to strengthen rather than erode as their careers carry them past the bust, which is the opposite of the stranding pattern one would expect if the response had equipped students with narrowly boom-specific skills.

Because the earnings medians condition on being found in covered employment, I re-estimate equation (16) with the share of graduates found employed as the outcome (Figure D1). The energy–non-energy employment gap is close to zero and flat across cohorts and horizons, indicating that energy and non-energy graduates are about equally likely to be working and that the earnings premium does not seem to be an artifact of a more selected energy workforce. As a robustness check, I also use graduates of healthcare programs at control colleges in the same play as the comparison non-energy group. Although these estimates combine differences in institution quality and student composition that the within-college design holds fixed, the results are remarkably similar to the within-college design (Figure D2).

Overall, I find no evidence that the institutional response resulted in program offerings with lower long-term earnings. But the limits of the design remain a significant constraint. The within-school gap recovers the causal return to energy training only if, among graduates of the same college and cohort, a student's field of study is independent of what he would have earned in either field. Thus if more able or more motivated students sort into energy programs, the gap overstates the return to the training. The boom also changes who graduates from energy programs, so the pre-to-post movement in the gap conflates any change in the return with a change in the composition of energy graduates. Finally, one additional

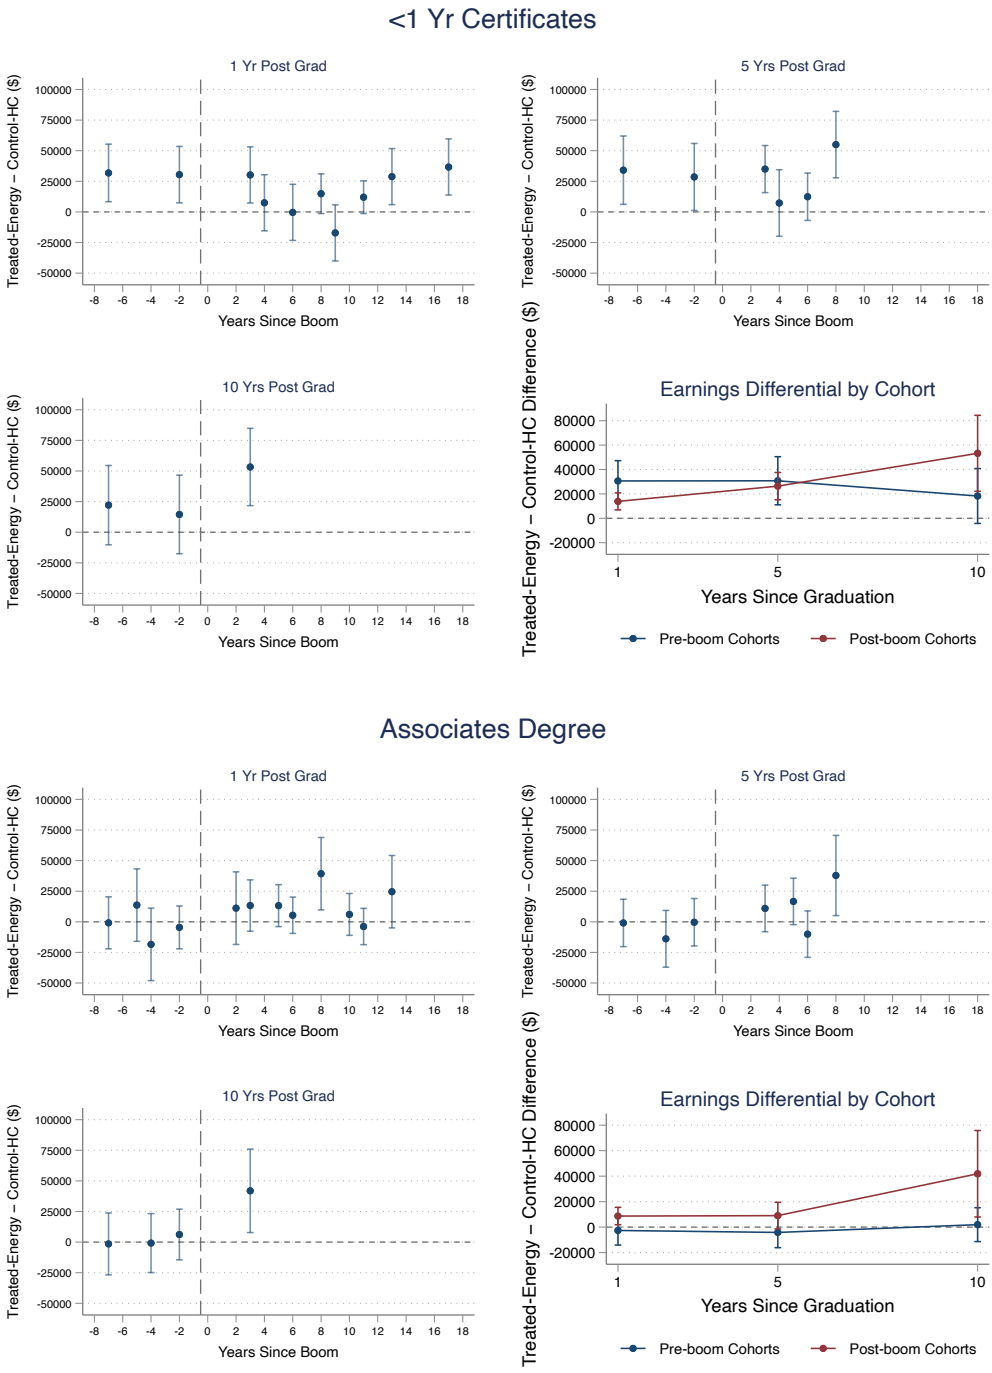
concern is that the five-year cohort bands straddle the boom onset, blurring the pre/post split. I therefore see these results as a credibility check on the earnings impacts of the supply response rather than as a precise estimate of its returns. The linked administrative design in Section 6 is better suited to establish the latter.

Figure D1: Within-College Energy–Non-Energy Employment Gap Over Time



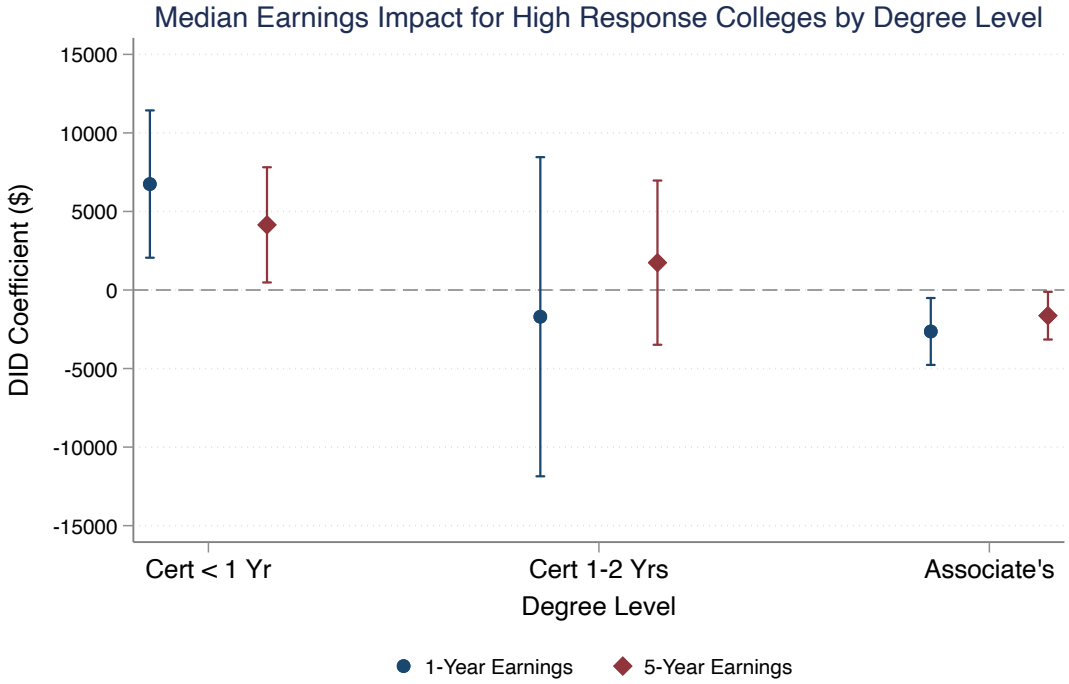
Notes. This figure plots the within-college, within-cohort difference in the share of graduates found in covered employment, between energy and non-energy graduates.

Figure D2: Cross-School Earnings Gap: Energy Graduates vs. Healthcare Graduates



Notes. This figure plots the cross-college earnings gap between energy graduates at treated colleges and healthcare graduates at control colleges in the same shale play, by event time and cohort.

Figure D3: Earnings DiD: High- vs. Low-Response Colleges



Notes. This figure plots event-study coefficients from a DiD comparing graduate earnings at high-response colleges (those that expanded energy programs) to low-response colleges within the same shale play, by event time.

E Robustness Checks

E.1 Weighted-CIP Measure

As an alternative to the binary classification used in the main text, I use the crosswalk weights themselves as the outcome measure. Rather than counting a program as energy when its energy weight exceeds the 2.5% threshold, each program contributes its energy share to the energy outcome and the remainder to the non-energy outcome. The institution-level outcome is the sum of these shares across the programs offered in a given year,

$$M_{i,t}^{J} = \sum_{m \in \mathcal{M}_{i,t}} \omega_{m,J}, \qquad J \in \{\text{Energy},\ \text{Non-Energy}\}, \tag{17}$$

where $\mathcal{M}_{i,t}$ is the set of programs at institution i in year t and $\omega_{m,J}$ is the share of program m mapped to industry group J. Energy's share is $M_{i,t}^{\text{Energy}}/(M_{i,t}^{\text{Energy}} + M_{i,t}^{\text{Non-Energy}}) \times 100$. I construct the analogous completion measures by weighting each program's share by the

awards it confers. All other features of the design follow Equation (5).

Tables E1 and E2 report the results, and Figures E1 and E2 plot the event studies. The shale boom increases energy-weighted programs by 0.037 (s.e. 0.014), 28% of the baseline mean of 0.13, and reduces non-energy weighted programs by 2.2 (s.e. 0.93). Energy-weighted completions rise by 0.59 (s.e. 0.28), 57% of their baseline mean, while non-energy completions fall by 84 (s.e. 39) and energy's share of completions rises by 0.10 percentage points (s.e. 0.05). As under the binary classification, energy's share of programs does not change significantly. In standard deviation units, the effects on energy programs (0.57) and energy completions (0.74) are, if anything, larger than under the binary classification (0.33 and 0.66), and the non-energy effects are similar. The expansion of energy offerings alongside the contraction of non-energy offerings therefore does not depend on the classification rule.

Table E1: Effect of Shale Boom on Programs Offered, Weighted-CIP Measure

Outcome:	Energy CIPs	Non-Energy CIPs	Energy's Share of CIPs
Raw Outcome:			
Treat × Post	0.037***	-2.248**	0.031
	(0.014)	(0.933)	(0.070)
Effect per SD:			
Treat × Post	0.565***	-0.230**	0.105
	(0.211)	(0.096)	(0.234)
CC FE	Y	Y	Y
Play × Year FE	Y	Y	Y
Observations	32,566	32,566	32,566
Plays	17	17	17
Baseline Mean	.13	25	.56
SD of ΔOutcome	.066	9.8	.3

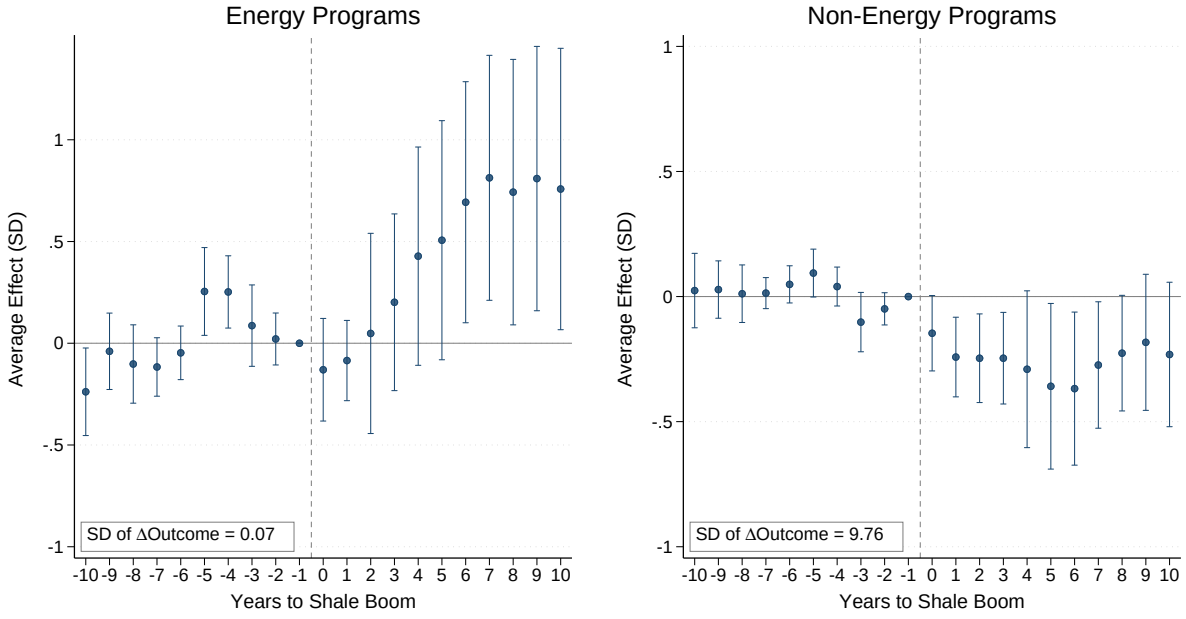
Notes. This table repeats Table 4 using the weighted-CIP measure, in which each program counts as its share of graduates' occupational employment in the energy industry rather than as a whole energy or non-energy program. Columns 1 and 2 report the effect on energy and non-energy weighted programs, respectively, and Column 3 reports energy's share of weighted programs in percentage points. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the play level in parentheses.

Table E2: Effect of Shale Boom on Completions, Weighted-CIP Measure

Outcome:	Energy Completions	Non-Energy Completions	Energy's Share Completions
Raw Outcome:			
Treat × Post	0.592**	-84.207**	0.101**
	(0.280)	(38.971)	(0.050)
Effect per SD:			
Treat × Post	0.740**	-0.377**	0.353**
	(0.350)	(0.174)	(0.176)
CC FE	Y	Y	Y
Play × Year FE	Y	Y	Y
Observations	32,566	32,566	32,551
Plays	17	17	17
Baseline Mean	1.04	304	.35
SD of ΔOutcome	.8	223	.286

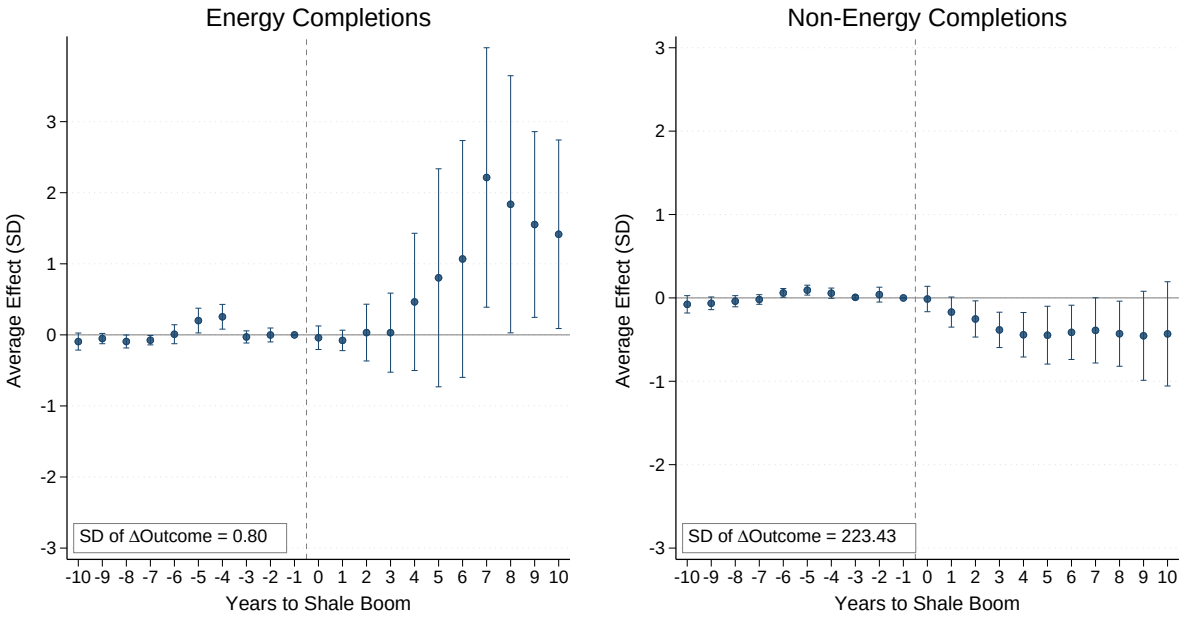
Notes. This table repeats Table 5 using the weighted-CIP measure, in which each program's completions are weighted by its share of graduates' occupational employment in the energy industry. The outcome in column 1 is energy-weighted completions, in column 2 non-energy weighted completions, and in column 3 energy's share of weighted completions (percentage points). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the play level in parentheses.

Figure E1: Event Study Estimates for Programs Offered, Weighted-CIP Measure



Notes. This figure plots the Gardner (2021) two-stage event study estimates for energy (left) and non-energy (right) weighted programs. Estimates are in standard deviations of the annual change in each outcome. Standard errors are bootstrapped and clustered at the play level.

Figure E2: Event Study Estimates for Completions, Weighted-CIP Measure



Notes. This figure plots the Gardner (2021) two-stage event study estimates for energy (left) and non-energy (right) weighted completions. Estimates are in standard deviations of the annual change in each outcome. Standard errors are bootstrapped and clustered at the play level.

E.2 Donut-Hole Design with Buffer Zone

Table E3: Donut-Hole Robustness by Buffer Width

Outcome:	Baseline (No Buffer)	Buffer Zone			
		10 miles	25 miles	50 miles	75 miles
	(1)	(2)	(3)	(4)	(5)
Energy CIPs	0.156**	0.145**	0.180***	0.150**	0.162**
	(0.063)	(0.065)	(0.060)	(0.059)	(0.077)
Non-Energy CIPs	-2.93***	-2.38**	-1.86	-2.09	-0.60
	(1.05)	(1.12)	(1.14)	(1.35)	(1.98)
Energy's Share of CIPs	0.247	0.253	0.224	0.128	0.146
	(0.241)	(0.243)	(0.265)	(0.244)	(0.311)
Energy Completions	5.53***	5.33***	5.26***	4.44***	4.73***
	(1.29)	(1.31)	(1.39)	(1.45)	(1.27)
Non-Energy Completions	-105**	-99**	-143***	-197***	-183***
	(46)	(44)	(46)	(55)	(64)
Energy's Share of Completions	0.575**	0.581**	0.474**	0.416**	0.554**
	(0.228)	(0.226)	(0.210)	(0.205)	(0.228)
CC FE	Y	Y	Y	Y	Y
Play × Year FE	Y	Y	Y	Y	Y
Observations	32,566	30,046	24,477	19,335	14,681
Plays	17	17	17	17	17

Notes. This table reports the Gardner (2021) two-stage difference-in-differences results of Equation 5 excluding control colleges within successively wider buffer zones around play boundaries. Column 1 reports the baseline specification with no buffer. Columns 2 through 5 report estimates after removing controls within the indicated distance of a shale play at 10, 25, 50, and 75 miles, respectively. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the play level in parentheses.

E.3 Stacked Regression with Alternative Standard Errors

Table E4 re-estimates Equation 5 as a single stacked regression rather than with the two-stage estimator, weighting the stacks following Wing et al. (2024). Because the sample has only 17 play clusters, the table reports three sets of standard errors. The estimates are similar to the headline results, although the energy completion effect is somewhat smaller. The energy program and completion effects remain statistically significant under all three sets of standard errors, while the non-energy completion effect is not significant with two-way clustering.

Table E4: Stacked Regression Estimates with Alternative Standard Errors

Outcome:	Clustered by Play	Two-Way Clustered by Play & College	Bootstrapped by Play
	(1)	(2)	(3)
Energy CIPs	0.342**	0.342*	0.342**
	(0.148)	(0.167)	(0.153)
Non-Energy CIPs	-0.258**	-0.258*	-0.258**
	(0.103)	(0.126)	(0.111)
Energy's Share of CIPs	0.150	0.150	0.150
	(0.098)	(0.109)	(0.116)
Energy Completions	0.542***	0.542***	0.542***
	(0.160)	(0.167)	(0.180)
Non-Energy Completions	-0.409*	-0.409	-0.409*
	(0.213)	(0.242)	(0.218)
Energy's Share of Completions	0.205**	0.205*	0.205**
	(0.088)	(0.100)	(0.102)
CC FE	Y	Y	Y
Play × Year FE	Y	Y	Y
Stacking Weights	Y	Y	Y
Observations	32,566	32,566	32,566
Plays	17	17	17
SD of Δ*Outcome:*			
Energy CIPs	0.467	0.467	0.467
Non-Energy CIPs	11.4	11.4	11.4
Energy's Share of CIPs	2.41	2.41	2.41
Energy Completions	8.42	8.42	8.42
Non-Energy Completions	247	247	247
Energy's Share of Completions	2.51	2.51	2.51

Notes. This table reports stacked difference-in-differences estimates of Equation 5 from a single regression with college and play-by-year fixed effects, weighted following Wing et al. (2024). Coefficients are in standard deviations of the annual change in each outcome, reported in the bottom panel. Column 1 clusters standard errors by play. Column 2 clusters them by play and college. Column 3 reports standard errors from a pairs cluster bootstrap over plays with 500 replications, rebuilding the stacking weights in each replication.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

E.4 Boom Definition: Ten Completed Wells

The main specification dates each boom to the year a play's shale potential became public knowledge, following Bartik et al. (2019), for 15 of the 17 plays. Public knowledge may not coincide with the start of drilling and hiring, so I re-estimate Equation 5 dating every boom to the first year a play reached ten completed horizontal wells. This moves several onsets by five years or more. The Bakken, Haynesville-Bossier, and Eagle Ford booms move earlier to 2000 or 2001, while the Barnett and Wolfcamp-Midland booms move later to 2005 and 2010 (Column 4 of Table B1). The estimates are similar to the main results (Tables E5 and E6). Energy programs rise by 26% of the baseline mean, compared with 30% in the main specification. Energy completions rise by 108%, compared with 117%. Non-energy programs and completions fall by 12% and 29%, compared with 10% and 31%. The event studies in Figures E3 and E4 show small pre-period coefficients and the same timing as the main results, with non-energy outcomes falling soon after onset and energy completions rising after several years. The results therefore do not depend on how the start of each boom is dated.

Table E5: Effect of Shale Boom on Programs Offered, Ten Completed Wells

Outcome:	Energy CIPs	Non-Energy CIPs	Energy's Share of CIPs
Raw Outcome:			
Treat × Post	0.140**	-3.423***	0.283
	(0.069)	(1.134)	(0.244)
Effect per SD:			
Treat × Post	0.290**	-0.316***	0.117
	(0.141)	(0.105)	(0.101)
CC FE	Y	Y	Y
Play × Year FE	Y	Y	Y
Observations	32,883	32,883	32,883
Plays	17	17	17
Baseline Mean	.53	29	1.7
SD of ΔOutcome	.48	11	2.4

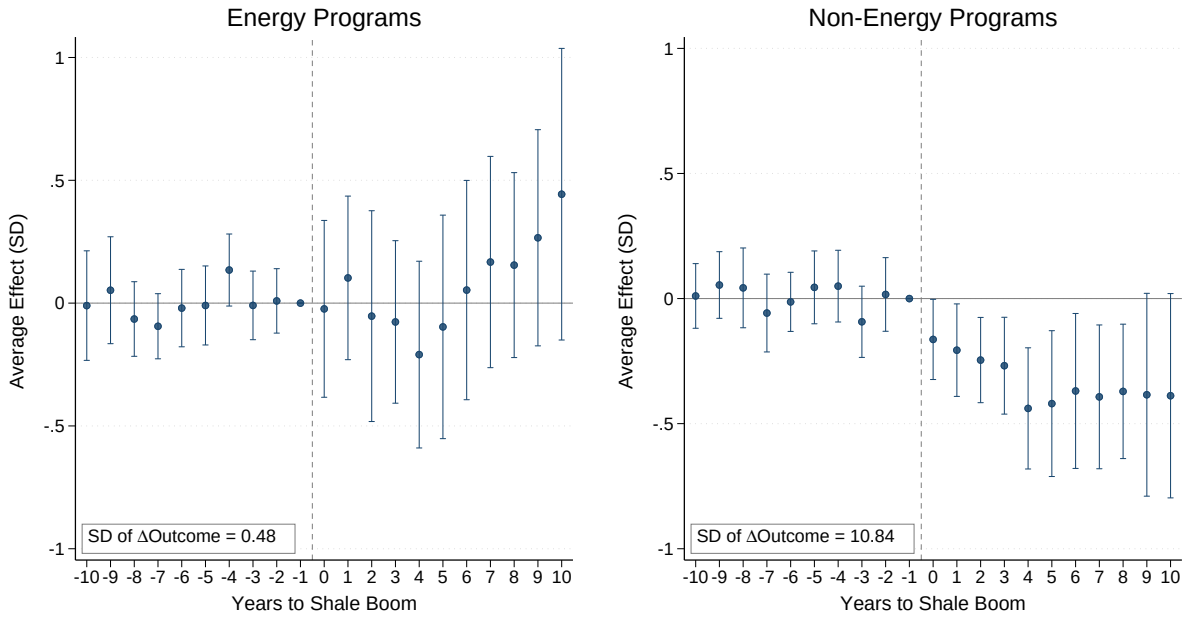
Notes. This table repeats Table 4 with each play's boom dated to the first year it reached ten completed horizontal wells. Columns 1 and 2 report the effect on the number of energy and non-energy programs (CIPs), respectively, and Column 3 reports energy's share of programs (CIPs) in percentage points. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the play level in parentheses.

Table E6: Effect of Shale Boom on Completions, Ten Completed Wells

Outcome:	Energy Completions	Non-Energy Completions	Energy's Share Completions
Raw Outcome:			
Treat × Post	5.221***	-100.250**	0.568***
	(1.313)	(43.995)	(0.219)
Effect per SD:			
Treat × Post	0.609***	-0.400**	0.239***
	(0.153)	(0.176)	(0.092)
CC FE	Y	Y	Y
Play × Year FE	Y	Y	Y
Observations	32,883	32,883	32,872
Plays	17	17	17
Baseline Mean	4.84	347	1.07
SD of ΔOutcome	8.57	250	2.37

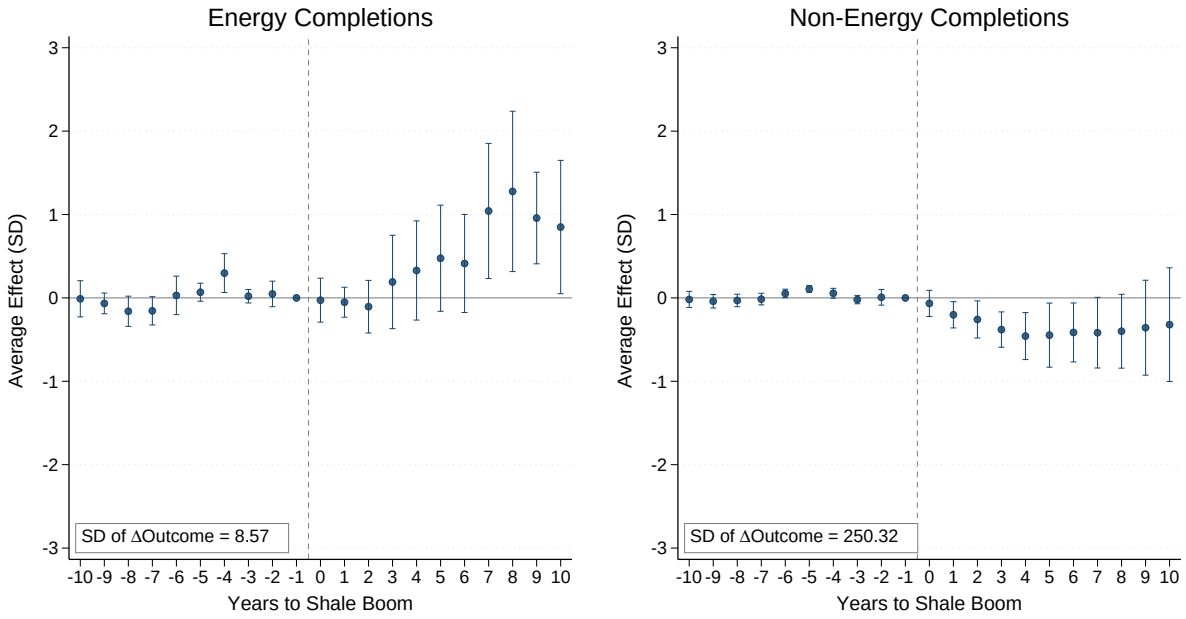
Notes. This table repeats Table 5 with each play's boom dated to the first year it reached ten completed horizontal wells. The outcome in column 1 is the count of energy completions, in column 2 the count of non-energy completions, and in column 3 energy's share of completions (percentage points). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the play level in parentheses.

Figure E3: Event Study Estimates for Programs Offered, Ten Completed Wells



Notes. This figure plots the Gardner (2021) two-stage event study estimates for energy (left) and non-energy (right) programs, with each play's boom dated to the first year it reached ten completed horizontal wells. Estimates are in standard deviations of the annual change in each outcome. Standard errors are bootstrapped and clustered at the play level.

Figure E4: Event Study Estimates for Completions, Ten Completed Wells



Notes. This figure plots the Gardner (2021) two-stage event study estimates for energy (left) and non-energy (right) completions, with each play's boom dated to the first year it reached ten completed horizontal wells. Estimates are in standard deviations of the annual change in each outcome. Standard errors are bootstrapped and clustered at the play level.

E.5 Alternative Solar Instruments

The main solar instruments interact local solar potential with state RPS targets and with the global solar module price. The federal investment tax credit (ITC) is a third potential shifter of solar profitability at the national level (Section 2.3). I interact solar potential with the ITC rate, lagged consistently with the other instruments, and re-estimate Equations 12 and 13.

The ITC predicts solar investment when it replaces RPS in the instrument, but adds no predictive power once the RPS and price instruments are included (Table E7). The first stage is weaker with the ITC instrument, with an effective F-statistic of 21 compared with 173 in the baseline (Montiel Olea and Pflueger, 2013). The reduced form remains significant across instrument sets, with Anderson–Rubin p-values of 0.060 or less for programs and 0.015 or less for completions (Table E8). The 2SLS estimates are stable (Table E9). A one-gigawatt increase in predicted solar investment raises solar programs by 0.078–0.085 SD and solar completions by 0.049–0.055 SD. The Hansen J test finds no evidence that the instruments imply different effects, with p-values between 0.59 and 0.96.

In California, the global solar module price does not predict county solar investment (first-stage F of 0.97), so I compare the California RPS instrument with the ITC instrument (Table E10). The two instruments are highly correlated (0.90) and give similar results. With the ITC instrument, a one-gigawatt increase in predicted solar investment raises the share of solar sections by 0.47 SD and solar sections by 0.57 SD, compared with 0.46 SD and 0.59 SD in the baseline. Non-solar sections remain unaffected under both instruments.

Table E7: First-Stage Estimates for National Solar Investment, Alternative Instruments

Instruments:	Baseline	RPS Only	ITC + Price	RPS + Price + ITC
	(1)	(2)	(3)	(4)
Effect per SD:				
Solar Potential ×State RPS$_{t-3}$ (SD)	0.737***	0.851***		0.742***
	(0.044)	(0.071)		(0.045)
Solar Potential ×Solar Module Price$_{t-3}$ (SD)	-0.269***		-0.522***	-0.267***
	(0.065)		(0.149)	(0.061)
Solar Potential ×Federal ITC$_{t-3}$ (SD)			0.607***	-0.008
			(0.064)	(0.051)
Year FE	Y	Y	Y	Y
State FE	Y	Y	Y	Y
Population Control	Y	Y	Y	Y
Joint F-Stat of Instruments	1,144	142	62	770
Effective F-Stat	173	142	21	147
Observations	24,268	24,268	24,268	24,268
Mean of Solar Investment (MW)	34	34	34	34
SD of Solar Investment (MW)	146	146	146	146

Notes. This table repeats the first stage in Table B4 with alternative instrument sets. The outcome is annual solar investment in the commuting zone in standard deviations, and the bottom rows report its mean and standard deviation in megawatts. Column 1 uses the baseline instruments, solar potential interacted with state RPS targets and with the solar module price. Column 2 uses the RPS instrument alone. Column 3 replaces the RPS instrument with solar potential interacted with the federal ITC rate. Column 4 uses all three instruments. Instruments are expressed in standard deviations. The effective F-statistic follows Montiel Olea and Pflueger (2013). The specification includes population controls and year and state fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the commuting-zone level in parentheses.

Table E8: Reduced Form Estimates for National Solar Outcomes, Alternative Instruments

Instruments:	Baseline	RPS Only	ITC + Price	RPS + Price + ITC
	(1)	(2)	(3)	(4)
Panel A: Δ Solar CIPs (SD)				
Solar Potential ×State RPS_{t-3} (SD)	0.010*	0.011***		0.007
	(0.006)	(0.003)		(0.011)
Solar Potential ×Solar Module $Price_{t-3}$ (SD)	0.000		-0.004	-0.002
	(0.009)		(0.008)	(0.010)
Solar Potential ×Federal ITC_{t-3} (SD)			0.011*	0.005
			(0.006)	(0.010)
Panel B: Δ Solar Completions (SD)				
Solar Potential ×State RPS_{t-3} (SD)	0.005**	0.006***		0.004
	(0.002)	(0.002)		(0.004)
Solar Potential ×Solar Module $Price_{t-3}$ (SD)	-0.003		-0.005	-0.003
	(0.004)		(0.004)	(0.004)
Solar Potential ×Federal ITC_{t-3} (SD)			0.004*	0.001
			(0.003)	(0.005)
Year FE	Y	Y	Y	Y
State FE	Y	Y	Y	Y
Population Control	Y	Y	Y	Y
Joint F-Stat, Δ Solar CIPs	5.63	9.22	2.84	3.96
Joint F-Stat, Δ Solar Completions	7.16	11	4.22	5.24
AR p-Value, Δ Solar CIPs	0.004	0.003	0.060	0.008
AR p-Value, Δ Solar Completions	0.001	0.001	0.015	0.001
Observations	24,268	24,268	24,268	24,268
SD of Δ Solar CIPs	0.24	0.24	0.24	0.24
SD of Δ Solar Completions	2.72	2.72	2.72	2.72

Notes. This table reports the reduced form for the instrument sets in Table E7. Outcomes are the annual changes in solar programs (Panel A) and solar completions (Panel B), in standard deviations of the annual change, reported in the bottom rows. Joint F-Stat and AR p-Value are the F-statistic and p-value from a joint test of the instruments. This is the Anderson–Rubin test, which is robust to weak instruments (Anderson and Rubin, 1949). The specification includes population controls and year and state fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the commuting-zone level in parentheses.

Table E9: Effect of Solar Investment on Community-College Solar Outcomes, Alternative Instruments

Instruments:	Baseline	RPS Only	ITC + Price	RPS + Price + ITC
	(1)	(2)	(3)	(4)
Panel A: Δ Solar CIPs (SD)				
$\widehat{SolarInvestment}_{t-2}$ (GW)	0.078***	0.085***	0.084*	0.078***
	(0.025)	(0.030)	(0.046)	(0.025)
Panel B: Δ Solar Completions (SD)				
$\widehat{SolarInvestment}_{t-2}$ (GW)	0.051***	0.049***	0.055***	0.051***
	(0.015)	(0.014)	(0.021)	(0.015)
Year FE	Y	Y	Y	Y
State FE	Y	Y	Y	Y
Population Control	Y	Y	Y	Y
Wald F-Stat	1,144	142	62	770
Effective F-Stat	173	142	21	147
Hansen J p-Value, Δ Solar CIPs	0.756		0.589	0.773
Hansen J p-Value, Δ Solar Completions	0.794		0.850	0.958
Observations	24,268	24,268	24,268	24,268
SD of Δ Solar CIPs	0.24	0.24	0.24	0.24
SD of Δ Solar Completions	2.72	2.72	2.72	2.72

Notes. This table repeats Table 11 with the instrument sets in Table E7. Effects are per gigawatt of predicted solar investment, in standard deviations of the annual change in each outcome, reported in the bottom rows. The lag is two years. The effective F-statistic follows Montiel Olea and Pflueger (2013). The Hansen J test of overidentifying restrictions is not available for Column 2, which has one instrument. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the commuting-zone level in parentheses.

Table E10: Effect of Solar Investment on Course Sections in California, Alternative Instruments

Outcome:	Δ Share Solar Sections		Δ Solar Sections		Δ Non-Solar Sections	
Instrument:	State RPS	Federal ITC	State RPS	Federal ITC	State RPS	Federal ITC
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 2SLS						
$\widehat{Solar Investment}_{t-1}$ (GW)	0.458**	0.468***	0.590*	0.565***	0.112	-0.173
	(0.210)	(0.151)	(0.304)	(0.197)	(0.348)	(0.417)
Panel B: Reduced Form						
Solar Potential ×State RPS$_{t-2}$ (SD)	0.019***		0.025***		0.005	
	(0.005)		(0.006)		(0.015)	
Solar Potential ×Federal ITC$_{t-2}$ (SD)		0.015***		0.018***		-0.005
		(0.002)		(0.002)		(0.013)
F-Stat of Instrument	13	60	16	107	0.10	0.17
AR p-Value	0.001	0.000	0.000	0.000	0.759	0.686
Panel C: First Stage (Solar Investment$_{t-1}$, SD)						
Solar Potential ×State RPS$_{t-2}$ (SD)	0.323***		0.323***		0.323***	
	(0.110)		(0.110)		(0.110)	
Solar Potential ×Federal ITC$_{t-2}$ (SD)		0.240***		0.240***		0.240***
		(0.077)		(0.077)		(0.077)
F-Stat of Instrument	8.67	9.77	8.67	9.77	8.67	9.77
SD of Solar Investment (MW)	131	131	131	131	131	131
Year FE	Y	Y	Y	Y	Y	Y
CZ FE	Y	Y	Y	Y	Y	Y
Population Control	Y	Y	Y	Y	Y	Y
Observations	2,567	2,567	2,567	2,567	2,567	2,567
Sample Mean	0.06	0.06	0.43	0.43	679	679
SD of ΔOutcome	0.12	0.12	0.67	0.67	98	98

Notes. This table repeats Tables B5 and 12 with the federal ITC instrument in place of the California RPS instrument. Panel A reports 2SLS effects per gigawatt of predicted solar investment, Panel B the reduced form, and Panel C the first stage. Outcomes are in standard deviations of the annual change, and the first stage is in standard deviations of solar investment. Instruments are expressed in standard deviations. AR p-Value is the Anderson–Rubin p-value, which is robust to weak instruments. With one instrument, the effective F-statistic (Montiel Olea and Pflueger, 2013) equals the F-statistic of the instrument. The lag is one year. The module price instrument is not used because it does not predict solar investment in California (first-stage F of 0.97). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the county level in parentheses.

E.6 Solar Covariate Balance and Baseline Controls

Table E11 compares baseline characteristics of colleges by the solar potential of their commuting zone. Each row regresses a characteristic on solar potential, both in standard deviations. Solar potential does not predict baseline solar programs or completions. Within states, however, colleges in higher-potential commuting zones had faster growth in CTE programs, higher FTE enrollment, and a larger Hispanic share before 2002, and their commuting zones

were larger and more often metropolitan. In California, colleges in higher-potential counties had fewer non-solar sections and lower enrollment (Table E12).

To test the sensitivity of the results to these differences, I add each baseline characteristic interacted with year to Equation 13 (Table E13). The completion effect is robust to every control and ranges from 0.037 to 0.112 SD. The program effect is similar with controls for population, pre-2002 CTE growth, metro status, and FTE enrollment, but falls to 0.030 SD and is no longer significant with the Hispanic share or with all controls together. Column 8 replaces the year and state fixed effects with state-by-year fixed effects, so only differences in solar potential within a state and year identify the effect. The program effect rises to 0.108 SD (0.047) and the completion effect is 0.047 SD (0.019), similar to the baseline. In California, the estimates are stable across controls (Table E14). All California counties share one RPS path, so the year fixed effects already absorb statewide changes. Column 5 adds commuting-zone-by-year fixed effects, so only counties in the same commuting zone and year are compared. The estimates for solar sections are similar but less precise, and the first stage is weaker, with an F-statistic of 5.7.

Table E11: Baseline Balance by Solar Potential, National Sample

			Solar Potential (SD)	
	Mean	SD	No FE	State FE
	(1)	(2)	(3)	(4)
Panel A. Outcomes, 1997–2001				
Solar Programs	0.01	0.08	-0.01 (0.02)	-0.00 (0.03)
Solar Completions	0.02	0.22	0.01 (0.03)	0.01 (0.04)
CTE Programs	34	22	0.13*** (0.02)	-0.00 (0.03)
CTE Completions	403	504	0.03* (0.02)	0.01 (0.03)
Annual Change in CTE Programs	1.39	3.08	0.15*** (0.01)	0.11*** (0.04)
Panel B. Institutional Characteristics, 1997–2001				
FTE Enrollment	2,973	3,715	0.23*** (0.02)	0.11*** (0.03)
% Part-time	58	18	0.27*** (0.04)	0.03 (0.02)
% Black	13	16	-0.09* (0.05)	0.07** (0.03)
% Hispanic	8.01	14	0.41*** (0.04)	0.20*** (0.03)
% Female	58	9.28	-0.02 (0.02)	-0.01 (0.02)
Panel C. Commuting Zone, 2000				
Population (000s)	1,970	3,553	0.74*** (0.17)	0.85*** (0.09)
Unemployment Rate (%)	4.29	1.31	0.15*** (0.04)	-0.04 (0.13)
% Construction	7.16	1.38	-0.05 (0.11)	-0.07* (0.04)
% Manufacturing	15	6.36	-0.19 (0.14)	0.03 (0.08)
% Mining and Oil & Gas	0.78	1.67	-0.01 (0.07)	-0.04 (0.04)
Metro Indicator	0.60	0.49	0.19*** (0.03)	0.13** (0.05)
Colleges			1,106	
Commuting Zones			471	

Notes. This table reports baseline characteristics of colleges in the national solar sample. Columns 1 and 2 report means and standard deviations. Columns 3 and 4 report the coefficient from a regression of each characteristic on solar potential, both standardized, without fixed effects and with state fixed effects. Outcomes and institutional characteristics are 1997–2001 averages for each college, and commuting-zone characteristics are from 2000. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the commuting-zone level in parentheses.

Table E12: Baseline Balance by Solar Potential, California

			Solar Potential (SD)	
	Mean	SD	No FE	CZ FE
	(1)	(2)	(3)	(4)
Panel A. Outcomes, 1997–2001				
Solar Sections	0.05	0.29	-0.05 (0.04)	-0.10 (0.08)
% Solar Sections	0.76	4.18	-0.05* (0.03)	-0.07 (0.05)
Non-Solar Sections	653	367	-0.12** (0.04)	-0.16*** (0.04)
Annual Change in Sections	32	47	0.15 (0.10)	0.07 (0.09)
Panel B. Institutional Characteristics, 1997–2001				
Enrollment (CCCCO)	19,013	12,068	-0.15** (0.06)	-0.23*** (0.05)
FTES (CCCCO)	2,139	1,278	-0.08* (0.05)	-0.13 (0.08)
Panel C. County, 2000				
Population (000s)	2,828	3,338	0.06 (0.12)	-0.24 (0.22)
Unemployment Rate (%)	5.05	2.14	0.21 (0.21)	0.07 (0.11)
Metro Indicator	0.94	0.24	0.01 (0.06)	-0.01 (0.01)
Colleges			113	
Counties			38	

Notes. This table reports baseline characteristics of California community colleges in the course-section sample. Columns 1 and 2 report means and standard deviations. Columns 3 and 4 report the coefficient from a regression of each characteristic on county solar potential, both standardized, without fixed effects and with commuting-zone fixed effects. Outcomes and institutional characteristics are 1997–2001 averages for each college, and county characteristics are from 2000. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the county level in parentheses.

Table E13: Effect of Solar Investment on Community-College Solar Outcomes, Baseline Controls

Control × Year:	Baseline	Population	Pre-2002 CTE Growth	Metro	% Hispanic	FTE Enrollment	All	State × Year FE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Δ Solar CIPs (SD)								
$\widehat{Solar Investment}_{t-2}$ (GW)	0.078***	0.085*	0.052*	0.088***	0.030	0.102***	0.030	0.108**
	(0.025)	(0.049)	(0.026)	(0.033)	(0.033)	(0.034)	(0.044)	(0.047)
Panel B: Δ Solar Completions (SD)								
$\widehat{Solar Investment}_{t-2}$ (GW)	0.051***	0.112***	0.048***	0.073***	0.037*	0.075***	0.086***	0.047**
	(0.015)	(0.032)	(0.017)	(0.016)	(0.020)	(0.019)	(0.033)	(0.019)
Year FE	Y	Y	Y	Y	Y	Y	Y	
State FE	Y	Y	Y	Y	Y	Y	Y	
State × Year FE								Y
Population Control	Y	Y	Y	Y	Y	Y	Y	Y
Wald F-Stat	1,144	512	1,116	999	748	882	433	316
AR p-Value, Δ Solar CIPs	0.004	0.048	0.083	0.007	0.346	0.001	0.148	0.047
AR p-Value, Δ Solar Completions	0.001	0.001	0.005	0.000	0.146	0.000	0.020	0.056
Observations	24,268	24,268	24,268	24,268	24,266	24,266	24,266	24,250
SD of Δ Solar CIPs	0.24	0.24	0.24	0.24	0.24	0.24	0.24	0.24
SD of Δ Solar Completions	2.72	2.72	2.72	2.72	2.72	2.72	2.72	2.72

Notes. This table repeats Table 11 with baseline characteristics interacted with year. Column 1 is the baseline. Columns 2 through 6 add, one at a time, the log of commuting-zone population in 2000, the average annual change in CTE programs in 1997–2001, a metro indicator, the Hispanic share of enrollment, and FTE enrollment. Column 7 adds all five. Column 8 replaces the year and state fixed effects with state-by-year fixed effects. Colleges missing a characteristic are assigned zero with a missing indicator interacted with year. Effects are per gigawatt of predicted solar investment, in standard deviations of the annual change in each outcome, reported in the bottom rows. AR p-Value is the Anderson–Rubin p-value, which is robust to weak instruments. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the commuting-zone level in parentheses.

Table E14: Effect of Solar Investment on Course Sections in California, Baseline Controls

Control × Year:	Baseline	Enrollment	Pre-2002 Section Growth	Both	CZ × Year FE
	(1)	(2)	(3)	(4)	(5)
Panel A: Δ Share Solar Sections (SD)					
$\widehat{SolarInvestment}_{t-1}$ (GW)	0.458**	0.499**	0.451**	0.493**	0.489*
	(0.210)	(0.227)	(0.208)	(0.226)	(0.277)
Panel B: Δ Solar Sections (SD)					
$\widehat{SolarInvestment}_{t-1}$ (GW)	0.590*	0.606*	0.631**	0.649**	0.706
	(0.304)	(0.305)	(0.306)	(0.307)	(0.534)
Panel C: Δ Non-Solar Sections (SD)					
$\widehat{SolarInvestment}_{t-1}$ (GW)	0.112	0.268	0.154	0.322	-0.372
	(0.348)	(0.262)	(0.327)	(0.270)	(0.321)
Year FE	Y	Y	Y	Y	
CZ FE	Y	Y	Y	Y	
CZ × Year FE					Y
Population Control	Y	Y	Y	Y	Y
Wald F-Stat	8.67	8.77	8.79	8.77	5.69
AR p-Value, Δ Share Solar Sections	0.001	0.001	0.001	0.000	0.000
AR p-Value, Δ Solar Sections	0.000	0.000	0.000	0.000	0.008
AR p-Value, Δ Non-Solar Sections	0.759	0.392	0.664	0.329	0.191
Observations	2,567	2,567	2,567	2,567	2,475
SD of Δ Share Solar Sections	0.12	0.12	0.12	0.12	0.12
SD of Δ Solar Sections	0.67	0.67	0.67	0.67	0.67
SD of Δ Non-Solar Sections	98	98	98	98	98

Notes. This table repeats Table 12 with baseline characteristics interacted with year. Column 1 is the baseline. Column 2 adds CCCCO enrollment in 1997–2001, Column 3 adds the average annual change in sections in 1997–2001, and Column 4 adds both. Column 5 replaces the year and commuting-zone fixed effects with commuting-zone-by-year fixed effects. Colleges missing a characteristic are assigned zero with a missing indicator interacted with year. Effects are per gigawatt of predicted solar investment, in standard deviations of the annual change in each outcome, reported in the bottom rows. AR p-Value is the Anderson–Rubin p-value, which is robust to weak instruments. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors clustered at the county level in parentheses.

E.7 Texas Earnings Attrition

Table E15 reports estimates of Equation 15 with the number of quarters with earnings in the Texas unemployment insurance records over each horizon as the outcome. Energy completers appear in about as many quarters as completers of other engineering and technician credentials in every cohort and at every horizon, and none of the energy differences is statistically significant. The largest is 0.95 quarters on a baseline of 33 quarters, for the 1998–2002 cohort over ten years. Differences are larger for some other fields. Completers in nursing, health technician, and public service programs appear in more quarters, while completers in personal care, arts, and other programs appear in fewer, with gaps of up to about 15% of the baseline mean.

Table E15: Quarters with Earnings by Cohort

Quarters with Earnings:	1-Year				5-Year			10-Year		15-Year
	1994-1997	1998-2002	2003-2007	2008-2011	1994-1997	1998-2002	2003-2007	1994-1997	1998-2002	1994-1997
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Program:										
Energy	0.02	0.08	0.01	-0.05	-0.23	0.39	0.11	-0.28	0.95	-0.40
	(0.12)	(0.10)	(0.07)	(0.06)	(0.56)	(0.38)	(0.30)	(1.03)	(0.72)	(1.58)
HC Nurse	0.41***	0.40***	0.31***	0.17***	2.19***	2.01***	1.61***	3.94***	3.94***	5.43***
	(0.04)	(0.03)	(0.02)	(0.02)	(0.17)	(0.13)	(0.11)	(0.33)	(0.25)	(0.49)
HC Technician	0.19***	0.19***	0.16***	0.13***	0.91***	0.97***	0.76***	1.65***	1.96***	2.42***
	(0.04)	(0.03)	(0.02)	(0.02)	(0.18)	(0.13)	(0.11)	(0.33)	(0.26)	(0.49)
Business	-0.08*	-0.05	-0.01	0.00	-0.40*	-0.21	-0.07	-0.83**	-0.14	-0.96*
	(0.04)	(0.03)	(0.03)	(0.03)	(0.21)	(0.15)	(0.14)	(0.39)	(0.29)	(0.55)
Computer Science / IT	-0.00	-0.03	-0.06	0.01	-0.19	-0.04	0.05	-0.27	0.18	-0.12
	(0.04)	(0.04)	(0.04)	(0.03)	(0.19)	(0.18)	(0.18)	(0.36)	(0.37)	(0.55)
Public Service	0.21***	0.16***	0.17***	0.15***	1.03***	1.05***	0.83***	2.03***	2.36***	3.05***
	(0.03)	(0.03)	(0.02)	(0.02)	(0.17)	(0.13)	(0.10)	(0.34)	(0.25)	(0.50)
Personal Care	-0.41***	-0.47***	-0.37***	-0.25***	-2.33***	-2.56***	-2.10***	-4.98***	-5.05***	-7.25***
	(0.06)	(0.05)	(0.04)	(0.04)	(0.27)	(0.21)	(0.19)	(0.52)	(0.39)	(0.73)
Arts	-0.21***	-0.29***	-0.24***	-0.14**	-1.22***	-1.13***	-1.16***	-2.77***	-2.22***	-4.53***
	(0.08)	(0.06)	(0.05)	(0.06)	(0.30)	(0.27)	(0.23)	(0.53)	(0.50)	(0.76)
Other	-0.10**	-0.08*	-0.10***	-0.06*	-0.51**	-0.49***	-0.64***	-0.90**	-0.94***	-0.93
	(0.05)	(0.04)	(0.04)	(0.03)	(0.21)	(0.19)	(0.16)	(0.40)	(0.35)	(0.58)
Cohort×Middle School FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Cohort×CC FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Test Score Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Demographic Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Grade 8 Cohort	1994-1997	1998-2002	2003-2007	2008-2011	1994-1997	1998-2002	2003-2007	1994-1997	1998-2002	1994-1997
Observations	31,248	44,505	53,785	49,765	31,248	44,505	53,785	31,248	44,505	31,248
R-Squared	0.190	0.185	0.160	0.160	0.204	0.199	0.176	0.205	0.200	0.203
Baseline Mean	3.35	3.39	3.44	3.56	16.58	16.81	16.82	32.96	32.82	48.71

Notes. This table reports estimates of Equation 15 separately for four sets of eighth-grade cohorts, with the number of quarters with earnings in the Texas unemployment insurance records over each horizon as the outcome. Column 10 reports 15-year horizons, which are observed only for the 1994–1997 cohorts. The omitted field is other engineering and technician credentials. All specifications include the individual covariates and middle-school-by-cohort and college-by-cohort fixed effects described in the text. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Standard errors are clustered at the community college by cohort level.