The Journal of INTERNATIONAL POLICY SOLUTIONS

Volume XXIII

Spring 2026

School of Global Policy and Strategy
University of California San Diego

THE JOURNAL OF INTERNATIONAL POLICY SOLUTIONS
School of Global Policy and Strategy
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0519

To contact the journal, please send inquiries to:
jips@ucsd.edu

ISSN 1937-1284

The Journal of International Policy Solutions is an annual academic journal published in San Diego, California. The views and opinions expressed in the journal do not necessarily reflect the views of the School of Global Policy and Strategy or the University of California, San Diego.

June 2026

The Journal of International Policy Solutions

Table of Contents

Letter from the Editor

Dear Readers,

Welcome to the twenty-third of the Journal of International Policy Solutions.

For us, reviving JIPS after a period of inactivity has been more than bringing back a publication. It has been an effort to preserve and renew a tradition that reflects what makes GPS special: a community of individuals who arrive here from remarkably different paths and experiences, yet are united by a common purpose—to better understand the world around us and to contribute to solving its most pressing challenges.

At GPS, we have learned alongside classmates whose journeys began in fields as varied as aerospace engineering, communications, economics, public service, and the arts. Despite these differences, we share a curiosity about policy, a commitment to rigorous analysis, and a belief that informed dialogue matters. So when we came across a dormant student journal that no one was tending, we couldn't help but imagine what it could be — a place for students to share their perspectives openly and honestly, in a way that felt both exciting and professional.

We are writing this letter at a moment that, in many ways, feels emblematic of our time. Each day brings developments that reshape our understanding of politics, economics, technology, and society. Policymakers and analysts are increasingly asked to make sense of a world characterized by rapid change, uncertainty, and competing narratives. The rise of artificial intelligence is only one example —offering extraordinary opportunities to learn, create, and work more effectively, while also raising difficult questions about ethics, governance, and the future of human decision-making.

Yet uncertainty is not unique to our generation. History reminds us that periods of profound transformation have always demanded thoughtful debate, intellectual humility, and a willingness to engage with perspectives different from our own. We believe journals such as JIPS play an important role in sustaining these conversations. They provide a space where emerging policy professionals can question assumptions, test ideas, and contribute meaningfully to discussions that extend beyond the classroom.

We are deeply grateful to everyone who submitted their work to this issue. Sharing one's ideas publicly requires vulnerability and courage, particularly in a world where the search for truth is often challenged by polarization, misinformation, and noise. For many around the world, the ability to speak, write, and publish freely remains far from guaranteed. We do not take that privilege for granted.

This issue is a testament to the intellectual curiosity, analytical rigor, and diverse experiences of our fellow GPS students. We hope it serves not only as a collection of thoughtful scholarship, but also as a reminder that meaningful policy solutions begin with asking difficult questions, listening carefully, and remaining committed to learning.

Thank you for reading, and for being part of the continued legacy of JIPS.

Warmly,

Nahal Lotfi
Editor-in-Chief

Max Avramis
Managing Editor

Pharmaceutical Industry Reform

Trinity Dubrow

EXECUTIVE SUMMARY

Brand name prescription drugs in the United States are costly. To reduce prescription drug costs while increasing innovation, lower manufacturer irresponsibility, and increase affordability, the United States must create policy that redesigns pharmaceutical R&D; with consideration to patient wellbeing and private sector interest. This memo recommends a comprehensive policy alternative with community-driven R&D;, expanded access to clinical trials for serious conditions, and value-based pricing. This alternative ties price to economic and clinical value, keeps incentives for high-quality R&D;, and ensures patient equity where there wasn't previously. Compared to other alternatives such as public manufacturing and specialty drug vouchers, the community driven R&D alternative suggested allows for the FDA and HHS to build on the United States' current R&D priorities while recentering affordability and quality patient care.

INTRODUCTION

It is no secret that the price of prescription drugs in the United States has gotten unfathomably expensive. Across all drugs, the US prices are nearly three times as high as they are in all other OECD countries. For brand-name drugs under strict patent laws, the cost is more than four times as high (Lovejoy et al, 2024). Although 80% of prescription drugs in the US are generic (Lal, 2015), patients needing specialty care are receiving inequitable access to patented drugs.

The United States' current legislative design to regulate and provide brand-name prescription drugs highly prioritizes R&D; initiatives. Clearly the innovation incentive is working, as the United States has completed 81.9% of all clinical trials since 1999. (World Health Organization, 2025). The FDA provides Intellectual Property (IP) rights to new brand-name drugs to recoup the costs of R&D and gain profit from the invention before the generic drug is available (Hickey & Ward, 2024), making innovation attractive. Other countries have adapted their pharmaceutical regulations to lower R&D costs and make new innovations more accessible (Lintern, 2025), begging the question as to why the United States has been unable to do so as well.

The most pressing issue to be addressed in relation to high prescription drug costs how the government permits monopoly pricing on brand-name drugs under IP. For many lifesaving drugs, there is only one seller and no other medical alternatives (Rajkumar, 2020). Pharmaceutical monopolies are directly correlated with FDA exclusivity and patent regulation. Unregulated monopolies, especially ones that produce lifesaving products, may adversely slow down the pace of innovation, which is a great cause for concern (Rajkumar, 2020). The design of the pharmaceutical market is harmful to patients and significantly raises taxes and premiums (Feldman, 2025).

Many pharmaceutical companies are able to extend patents and high prices through loopholes in policies and regulations, creating dysfunction in the free market. In 2023, a policy statement from the FTC stated that dozens of inhalers were improperly and illegally patented due to a regulatory loophole, leading to unaffordable drug products (Robbins, 2023). Improper patents on drug devices are not just common amongst inhalers, but other drug devices as well. Additionally, manufacturers sometimes make minor modifications to a drug that is about to lose its patent to allow for a new patent to be approved, thus delaying generic competition and accessibility (Feldman, 2025). This flags an issue about the efficacy of the patent regulation, exclusivity, drug price negotiation, and pharmacoeconomic strategies in the US.

CURRENT POLICY

Drug Patent Regulation

Pharmaceutical companies have strong legal protections that keep generic competitors from the marketplace (Jacobson, 2025). These companies can have a drug patent for up to 20 years prior to its release to the generic marketplace. Lengths of patents may be shorter or longer at the discretion of the FDA. This is intended to incentivize the innovation of new drugs and clinical research (Lal, 2015). The patent system intends to encourage competition and the invention of new prescription drugs, however there is evidence that most new patents aren't actually for new drugs-but are rather existing ones. Each of America's most profitable drugs has an average of 74 patents each (Barrueta, 2022).

Market Exclusivity

The FDA provides market exclusivity on certain prescription drugs with the intent of sustaining the research of new prescription medication options (Lal, 2015). With exclusivity, the FDA will not grant approval to drugs for similar conditions within a designated time period, leaving patients with less options. The FDA's exclusivity program is intended to incentivize innovation of new prescription drugs and provide a guaranteed money back on R&D; costs, which averages about 3 billion dollars per drug (Rajkumar, 2020). Although market exclusivity is generous and supportive of the prescription drug market, it leaves patients' needs to the wayside. Below is the list of exclusivities that are driving up costs:

- **Orphan Drug Exclusivity (ODE):** Granted to drugs that are approved to treat conditions affecting less than 200,000 in the US, or more than 200,000 with low hopes of recovering costs. It bars applications for a similar drug for the same "orphan" disease or condition for 7 years.
- **New Chemical Exclusivity (NCE):** Bars the FDA from approving drugs that have the same new chemical entity (not previously approved by the FDA) for a total of 5 years.
- **"Other" Exclusivity:** Keeps the FDA from approving, for a three-year period, any application that relies on the information supporting the approval of the drug or a change to the drug for which the information was submitted and the exclusivity granted after the reporting of completing new clinical investigations.
- **Pediatric Exclusivity (PED):** Grants the company six additional months of market protection at the end of a patent after having submitted pediatric studies on the chemical composure. (Lal, 2015).

"Delivering Most-Favored-Nation Prescription Drug Pricing to American Patients"

This May 2025 Executive Order written by Donald Trump promised to lower the price of prescription drugs with support of HHSA (Rogers, 2025). Although he promised for drug prices to lower "almost immediately", no known progress has happened to advance this executive order into action (Jacobson, 2025).

Inflation Reduction Act (IRA)

The IRA granted the federal government the authority to negotiate Medicare drug prices with pharmaceutical companies to ensure that Medicare recipients pay less. Additionally, the IRA provides Medicare Part D beneficiaries with an out-of-pocket cost cap at $2,000 annually (saving about $400 per person per year) (Hughes & Rapfogel, 2023).

HR1, "The Big Beautiful Bill"

Unfortunately, some of the Medicare price negotiation on "Orphan Drugs" that were anticipated when the IRA passed in 2022 will likely be unavailable with the passing of HR1. The orphan drug definition has been broadened to include more drugs and longer exclusivity periods, added to the bill after much lobbying from the pharmaceutical industry (Cubanski & Neuman, 2025).

POLICY CRITERIA

Expand Equity and Affordability: More than any other criteria, affordability must centerany policy created to address prescription drug spending. The current annual spending rates are anticipated to rise 3-6% worldwide, which is unsustainable and inequitable (Rajkumar, 2020).

Policies must demonstrate measurable deductions in out-of-pocket costs for patients. Equity must be measured by the policy's ability to provide access to different people across socioeconomic, geographic, gender, and race lines. Outcomes to measure the policy's effectiveness are the increased utilization of brand-name drugs, noticeable reduction in the median out-of-pocket spending after 3 years, and equal access for public and private insures.

Increased Innovation: Any accepted policy alternative must include increased innovation of new drugs without reinforcing monopoly pricing. The current patent and exclusivity laws are created with the intent of incentivizing innovation; however, some scholars suggest that they incentivize neither R&D; nor economic competition (Dosi et. al, 2023). Policy alternatives should reward development that demonstrates measurable improvements in therapeutic efficacy, access, and affordability. Possible metrics to measure this are the proportion of federally funded or collaboratively funded R&D; projects developed under non-monopolistic models.

Reduce Fiscal Irresponsibility Across the Pharmaceutical Industry: The policy alternative best fitting to address high prescription drug prices will eliminate opportunities for the private sector to prioritize profit over patient benefit. Fiscal responsibility includes ensuring that private expenditures and public spending yield proportional outcomes. Alternatives that meet these criteria will include evaluation metrics to increase price transparency, decrease patent misuse, and ensure cost-effectiveness. To measure additional responsibility in any policy alternatives, the FDA should evaluate the decrease of secondary or extended patents for non-novel modifications and the GAO could complete reports on the return on investment for federally funded research.

POLICY ALTERNATIVES

In the US, citizens overwhelmingly support federal policy efforts to reduce prescription drug costs (Sullivan, 2022). Even though the price of brand name drugs has risen, the success rate of drug development has declined, exemplifying a lowering return on investment (Austin & Hayford, 2022). Due to bipartisan support, and increasing costs, there is significant evidence that the current pharmaceutical regulation design should either be reformed or removed altogether.

Patent Reform

A patent reform policy alternative would include numerous developments to the standing patent legislation to better protect them from misuse. It would be the most efficient with the least amount of addition of administrative oversight. This patent reform program would enshrine the requirement that prices be tied to the clinical and economic value that they provide to society.

- ***Value-Based Price Ceilings:*** A value-based price ceiling would tie drug prices to the drug's measured outcomes. This reform would encourage manufacturers to align prices with clinical value. This intends to cap prices according to the drug's clinical effectiveness and its level of necessity. To implement this, CMS could create outcome-based reimbursement contracts for a variety of prescription drugs. CMS would begin integration with cancer treatment and gene therapies then continue as negotiation opportunities expand. This would force the industry to sell their drugs closer to the manufacturer's price, while also creating room for future R&D; funds.
- ***Capped Price Increases:*** The FDA could develop a methodology to ensure that brand-name drug price increases do not go over a certain percentage annually. The price increase could vary by measured outcomes, clinical effectiveness, and length on the market. Doing this would ensure that prescription drug costs do not rise above what is necessary to fund future R&D.
- ***Patent Buyout Funds:*** A patent buyout fund is a mechanism through which the federal government could purchase the exclusive IP rights to a drug or medical device with the intent ofmaking its intellectual property publicly available. The FDA could review a variety of drugs with high social value and low accessibility and choose some to buyout. This allows these drugs to be manufactured generically, therefore at lower costs. This introduces the opportunity for the pharmaceutical company to receive returns on their R&D; while further incentivizing developing high value products. This model would raise cost effectiveness while creating equitable access. Long term, this design will lead to a reduction in overall public program spending and reduce volatile pricing within the monopoly system.

International Importation of Prescription Drugs

The international importation of prescription drugs would allow for the United States to form bonds with other countries, break up monopolies in the US pharmaceutical industry, provide patients with foreign specialty treatments, and lower prices effectively and equitably. The FDA would create a program with partner countries, such as the EU, Australia, and Japan, to authorize federal importation from certified facilities abroad. All countries could agree on a safety protocol, including specific labeling and expanded international travel surveillance to ensure no tampering with the prescription drugs. The program would begin with the importation of 10 pilot drugs, specifically for conditions that have few inexpensive therapeutic options. Provisions of the pilot drugs will be given to members of communities that face medication non-adherence due to high prices of their prescriptions. This may also have an impact on improved foreign affairs with partner countries, especially if there is an agreement to lower tariffs on all prescription drug trade.

Community-Driven R&D with Parallel Clinical Trial Access

Designed by example from the activism and R&D; protocol from the AIDS crisis of the 1980s and 1990s (France, 2016), this policy alternative prioritizes lowering R&D; prices and incorporate value-based pricing as a means to continue innovation while encouraging accessibility to highly effective drugs. Patient Priority Councils (PPCs) would be chosen by the HHSA, FDA, and NIH to co-define research priorities, clinical trial criteria, trial endpoints, and risk tolerance requirements. PPCs would be composed of community members with a personal connection to serious condition, such as being a part of an awareness group, having a particular disease, or academics studying serious conditions. The PPCs will verify that clinical trials are accessible and inclusive, intentionally including underserved populations such as across ethnicities, genders, geographical location, and income, especially in cases where there is a dire need for speedy clinical trials. Parallel clinical trial access will drastically change the R&D; process and its costliness, by providing pharmaceutical manufacturers additional funds that go to R&D; for serious conditions. Parallel clinical trial access would allow eligible patients to undergo investigational treatment under structured protocols. Parallel access will additionally allow for time-limited rapid qualification of effective drugs, providing conditional approval. This would also allow the PPC to require post-market requirements, such as capped, value-based drug prices.

Public-Private Manufacturing Authority

This policy alternative would be by far the most drastic change to the pharmaceutical industry. It would establish a government-backed production design for brand-name drugs with little market competition. In this model, the Federal government, through DHSS, would contract with private pharmaceutical manufacturers to provide drugs at a transparent price. The government's participation in manufacturing would cause stabilized and reduced prices. This program would include shared risk-taking and performance-based subsidies. Private manufacturers could receive subsidies tied to providing low prices and meeting supply demands. It will reduce necessity for foreign imports and demonstrates federal leadership on drug affordability.

Specialty Drug Vouchers

A specialty drug voucher program would create almost immediate access and to the specialty drugs prescribed at an affordable cost. Specialty drug vouchers would be federally funded and redeemable at insurers or pharmacies. For people who are 500% under the FPL or less, they can receive vouchers for a specialty drug if prescribed them by their provider. After receiving legislative authorization, CMS can begin rulemaking about the voucher distribution conditions. The return on investment would happen within a few years, with more patients accessing the brand-name drugs essential to their care, there would be reduced hospitalizations and emergency care use, lessening the financial burden on the healthcare system. Additionally, this would be a politically popular, by ensuring that the people that are in dire need get the care, but also not interfering directly with the market.

LIMITATIONS

Patent Reform: There is competing evidence about the efficacy of the patent and exclusivity structure and its effectiveness, and some experts state that this design should be removed altogether. A study showed that small pharmaceutical companies may opt out of R&D; due to the limitations caused by the concentrated market system due to legal market exclusivities (Dosi et. al, 2023). This policy, although would have long term benefits, would have high up-front costs. Additionally, there is a possibility for overpayment, because evaluating the "fair" value of a patent is incredibly difficult. Short term this could also introduce some supply chain issues, as the rapid number of generic drugs being produced could interrupt current manufacturing procedures.

International Importation of Prescription Drugs: The most concerning aspect of this alternative is its potential to decrease innovation in the domestic pharmaceutical market, due to the policy's inherent design that creates more market competition. It could, however, allow patients accessibility the of foreign R&D;, therefore providing more therapeutic options. It may increase FDA costs due to any resource strain and additional oversight necessary for this program. This policy design would also be unpopular with US pharmaceutical companies, that will be averse to additional competition, and therefore bump up additional lobbying from the industry.

Community-Driven R&D with Parallel Clinical Trial Access: This alternative does open up the possibility to unintentional inequity of accessibility should PPC members have bias, however, there is opportunity to address this by equity quotas and transparency to the public about the makeup of the PPCs. Some concern may also lie in the additional administrative load- a program of this size would include. It would increase the necessity for more FDA employees and the implementation of additional electronic health records.

Public-Private Manufacturing Authority: Of all the policy alternatives, this one would receive the most political debate. Due to its reforming of the entire market, there would be lots of lobbying and legal pushback, especially from pharmaceutical companies and conservative perspectives that request limited government intervention. This would be a costly policy, returns on the investment would accrue gradually. This policy may also, in turn, disincentivize private investment into similar prescription drug products.

Specialty Drug Vouchers: Specialty drug vouchers may incentivize manufacturers to keep overall prices high, knowing that vouchers will make specialty drugs affordable for those that need them. Therefore, this policy as a standalone could become costly if there are not additional price controls and ceilings in place to monitor costs. Most concerningly, specialty drug vouchers have the potential for additional inequities, due to the inevitable patients that fall outside of the eligibility thresholds. This policy alternative may receive political pushback as many constituents find any form of socialized medicine to be unfavorable.

RECOMMENDATION

After careful evaluation of the policy alternatives and their ability to successfully meet the criteria for success, the most effective policy is the Community-Driven R&D; with Parallel Clinical Trial Access.

This option increases innovation incentives by rewarding manufacturers for R&D; of highly effective drugs, minimizes manufacturer irresponsibility, and limits excessive base prices. Community driven R&D; coinciding with parallel clinical trial access creates a structured opportunity for government action when necessary, fulfilling opportunities for equity and value-based capping on the most valuable brand-name prescriptions.

It balances private sector interests with necessary accountability to ensure affordability and accessibility for patients. Implementing this reform would position the U.S. as both a leader in drug innovation and in an equitable pharmaceutical industry.

REFERENCES

Austin, D., & Hayford, T. (2021, April). *Research and Development in the Pharmaceutical Industry.* Congressional Budget Office. https://www.cbo.gov/publication/57126

Barrueta, A. (2022, December 5). *Want to Lower Drug Prices? Reform the U.S. Patent System.* About.kaiserpermanente.org. https://about.kaiserpermanente.org/news/want-to-lower-drug-prices-reform-the-us-patent-system

Conti, R. M., Frank, R. G., & Gruber, J. (2021). *Regulating Drug Prices while Increasing Innovation. New England Journal of Medicine, 385(21), 1921–1923. https://doi.org/10.1056/nejmp2113764*

Dosi, G., Palagi, E., Roventini, A., & Russo, E. (2023). Do Patents Really Foster Innovation in the Pharmaceutical sector? Results from an evolutionary, agent-based Model. Journal of Economic Behavior & Organization, 212, 564–589. https://doi.org/10.1016/j.jebo.2023.05.039

Feldman, W. B. (2025). Patent Thickets and Product Hops: Challenges and Opportunities for Legislative Reform. The Journal of Law Medicine & Ethics, 1–6. https://doi.org/10.1017/jme.2025.54

France, D. (2017). How to survive a plague : the story of how activists and scientists tamed AIDS. Vintage Books.

Hickey, K., & Ward, E. H. (2024, January 30). The Role of Patents and Regulatory Exclusivities in Drug Pricing. Congress.gov. https://www.congress.gov/crs-product/R46679

Hughes, S., & Rapfogel, N. (2023, October 12). Following the Money: Untangling U.S. Prescription Drug Financing. Center for American Progress. https://www.americanprogress.org/article/following-the-money-untangling-u-s-prescription-drug-financing/

Jacobson, L. (2025, May 20). Trump Exaggerates Speed and Certainty of Prescription Drug Price Reductions - KFF Health News. KFF Health News. https://kffhealthnews.org/news/article/drug-price-pricing-executive-order-may-12-most-favored-nation-fact-check-timing/

Lintern, S. (2025, November 2). New drugs to be approved faster so patients don't wait years. Thetimes.com; The Sunday Times. https://www.thetimes.com/uk/healthcare/article/new-drugs-to-be-approved-faster-so-patients-dont-wait-years-k6qtfp839?utm_campaign=GBR%20SC%20FTI%20Consulting%20LS%20Round-up%203%20November%202025&utm_medium=email&utm_source=Eloqua&trk=article-ssr-frontend-pulse_little-text-block

Mulcahy, A. W., Schwam, D., & Lovejoy, S. L. (2024). International Prescription Drug Price Comparisons: Estimates Using 2022 Data Assistant Secretary for Planning and Evaluation (ASPE). In Assistant Secretary for Planning and Evaluation (ASPE) U.S. Department of Health & Human Services. RAND Corporation. https://aspe.hhs.gov/sites/default/files/documents/8e057b0a094e6f9b9d01171fce6698f4/international-price-comparisons.pdf

Neuman, T., & Cubanski, J. (2025, September 24). People with Medicare Will Face Higher Costs for Some Orphan Drugs Due to Changes in the New Tax and Budget Law | KFF. KFF. https://www.kff.org/medicare/people-with-medicare-will-face-higher-costs-for-some-orphan-drugs-due-to-changes-in-the-new-tax-and-budget-law/

Rajkumar, S. V. (2020). The high cost of prescription drugs: Causes and solutions. Blood Cancer Journal, 10(6), 71. https://doi.org/10.1038/s41408-020-0338-x

Renu Lal, Pharm.D. (2015). Patents and Exclusivity. Food and Drug Administration. https://www.fda.gov/media/92548/download

Rogers, H.-A. (2025). Most-Favored-Nation Prescription Drug Pricing Executive Order: Legal Issues. Congress.gov. https://www.congress.gov/crs-product/LSB11319?

Sullivan, S. D. (2022). Medicare Drug Price Negotiation in the US: Implications and Unanswered Questions. Value in Health, 26(3), 394–399. https://doi.org/10.1016/j.jval.2022.11.015

World Health Organization. (2025, November). Number of clinical trials by year, country, region and income group. Www.who.int; World Health Organization. https://www.who.int/observatories/global-observatory-on-health-research-and-development/monitoring/number-of-clinical-trials-by-year-country-who-region-and-income-group

Closing the Efficiency Deficit: Vietnam's Path to High-Income Status

Charles Brightman Skinner III

ROADMAP

This paper examines the significant economic success Vietnam witnessed over a brief fifty-year period. The country requires a brief introduction where its history, population, and economic standing will be stated. The paper moves on to the empirical data section, where output per worker, physical capital, and human capital are calculated and explained. TFP will be backed out as a residual. The paper will then examine several key policies that contributed to Vietnam's successful transition to a middle-income nation and how those policies will shape the nations continued rise.

HISTORICAL INTRODUCTION

Situated in Southeast Asia on the Indochina peninsula, Vietnam has a nearly 3,000-mile land border. China, Loas, and Cambodia border the country to the north and west, while the South China Sea lies to the east (Consulate Viet Nam). Positioned along key land and maritime routes, Vietnam's geography fostered commerce but also left it vulnerable to unwelcome meddling. Throughout its long history, Vietnam was plagued by foreign aggression and colonial strife. For nearly twelve centuries, the country fought wars of resistance against Chinese invaders. Early records indicate that resistance against the Chinese Qin Dynasty in the third century B.C. was the beginning of these struggles (Consulate Viet Nam). Vietnam would continue to resist Chinese occupation by various dynasties until the 10th century. Following the cessation of major warfare, Vietnam experienced a golden age of existence that led to major agricultural and construction innovations. Rice cultivation would dominate the economic priorities of the country, transforming Vietnam into a prosperous nation (Thành Cổ Loa). Vietnam's era of prosperity ended when the wave of European expansion swept over Vietnam, bringing it under colonial rule. France and Britian competed for dominance by sending missionaries to the region that were brutally expelled by the local population. This led to the French invasion and subsequent occupation of the country for nearly 100 years beginning in 1857. During this time, Vietnam's population and resources would be violently exploited, leading to resistance movements against their occupiers (UC Davis). These movements would not be successful until after the Second World War ended, when colonial powers, weakened by the conflict, began to grudgingly retract their influence abroad. Vietnam began its fight for independence against in 1946, after French forces seized the south (History.com). Following the

Geneva Accords in 1954, the country was divided in two at the 17th parallel (Consulate Viet Nam). A French and American supported regime occupied the southern half of the country, while the pro-communist government, lead Ho Chi Minh, occupied the northern half of the country. Unification of the country was intended to happen gradually but anti-communist ideology in America led to another brutal war. This unsuccessful struggle to oust the communists would persist for nearly twenty years. Exhausted from two decades of costly war, and with domestic issues taking priority, the United States pulled its remaining forces out of Saigon in 1975. In the years that followed, the country achieved unification, and the foundations of economic prosperity began to take shape.

ECONOMIC INTRODUCTION

The earliest economic records paint Vietnam as a country with a rich agricultural history, where its people were among the first to farm wet rice. Due to the tropical climate of the country, plentiful rainfall and sunlight created ideal conditions for the rice crop to flourish. Until the 19th century, Vietnam's economy was largely feudal and centered on agriculture, with land and rice fields controlled by the royal court, government officials, and local landlords (Consulate Viet Nam). As the system was not conducive to an export economy, this resulted in the local consumption of the crop. This all changed when French colonists invaded in the 1860s. With over 90% of the country's population working in agriculture, the colonists created a vast rice export market. By 1934, the peak of French economic extraction in Indochina, rice exports from the port of Saigon exceeded 1.5 million tons (Geoffrey Gunn). As most of the rice was exported and a drought soon followed, the local population suffered famine and war until 1975. With the settlement of the war with America, the country faced a bleak decade. Amid recovery from three decades of war and the issues of a centrally planned economy, the county was unable to provide adequately for its population. “Things began to change for the better in 1986, with the start of an ambitious program of economic reform known as ***Doi Moi***, or renovation” (International Monetary Fund). The IMF and World Bank drastically helped the country recover from the dire scenario of the 1980s. This paper will focus on the period of 1970 to 2023 to give a clear picture of how much Vietnam has progressed since then.

In 1985, Vietnam's nominal GDP was $14.1bn with nominal GDP per capita at $239 (EconScope). This is compared to GDP in the US at $4.3trn and nominal GDP per capita at $18,237. Since then, Vietnam has significantly closed the gap with the US, with most recent figures showing Vietnam's 2023 nominal GDP at $479.4bn and 2023 nominal GDP per capita of $4,717 compared to the US with current Nominal GDP of $29.2trn and nominal per capita GDP of $85,810. Vietnam had a population of 59 million people, with 27 million employed, in 1985 compared to the US with a population of 242 million, with 112 million employed. Current 2023 figures have the Vietnamese population at 100 million, with 55 million employed and the US population at 343 million and 168 million employed. I have compared Vietnam to the US in this paper to better show the vast developmental gap the country has closed in the last half century.

EMPIRICAL DATA

Using the most recent version of the Penn World Tables and the Barro-Lee data, this paper shows the gap between Vietnam and America in terms of output per worker, physical capital per worker, human capital per worker and TFP. The data collected looks back to 1970 and through to 2023 (2015 for human capital as that is the most recent datapoint). This paper uses the Cobb-Dougles function to find the factors for the four data points in the two separate years to show Vietnam's rapid growth. This paper will walk through how to find the four data points for those two years, using the function below:

$$Y1 = Ai\ (Ki)\alpha\ (Hi)1\text{-}\alpha$$

Output per worker in 1970 for Y

Using Vietnam's employed population of 15M and the expenditure-side real GDP at chained PPPs (in mil. 2021US$) of $41,173M, output per worker was $2,709M. The United States had an employed population of 84M and expenditure-side real GDP at chained PPPs (in mil. 2021US$) of $5,857,693M, the output per worker was $69,369M. In 1970, U.S. output per worker was 25.6 times higher than that of Vietnam, reflecting the vast income differences between the two countries at the time. Formula below:

$$YUS / YVNM = \$69{,}369M / \$2{,}709M$$

$$Y1 = 25.6$$

Output per worker in 2023 for Y

Using Vietnam's employed population of 55M and the expenditure-side real GDP at chained PPPs (in mil. 2021US$) of $1,339,399M, output per worker was $24,373M. The United States had an employed population of 168M and expenditure-side real GDP at chained PPPs (in mil. 2021US$) of $25,641,846M, the output per worker was $153,049M. The ratio of U.S. to Vietnamese output per worker was approximately 6.3 in 2023, illustrating the substantial convergence in productivity, living standards, and income levels that has occurred over the past half-century. Formula below:

$$YUS / YVNM = \$153{,}049M / \$24{,}373M$$

$$Y1 = 6.3$$

Figure below shows the output per worker from 1970 to 2023, showing drastic narrowing in recent years and a spike in the early 1970s due to the Vietnam War:

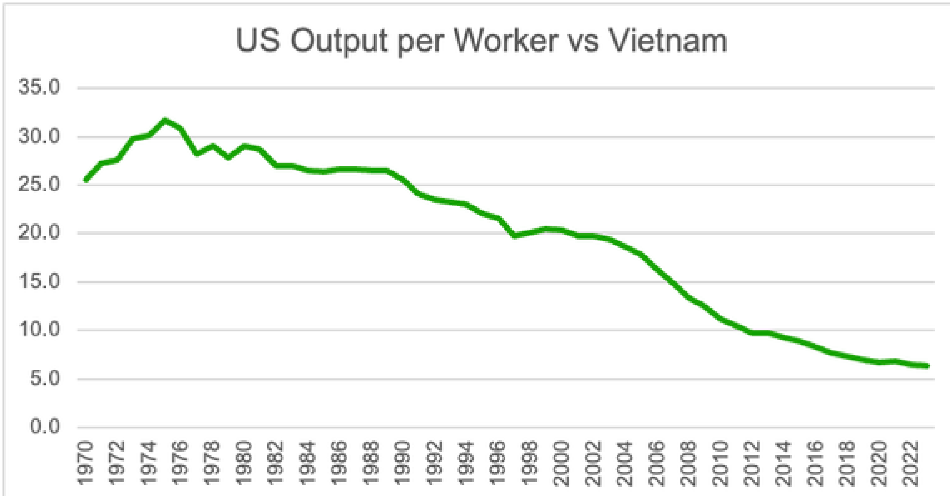

Physical capital per worker in 1970 for K

Given Vietnam's employed population of 15M and physical capital at current PPPs (in mil. 2021US$) of $37,267M, physical capital per worker was $2,452M. The United States had an employed population of 84M and physical capital at current PPPs (in mil. 2021US$) of $21,169,220M, the physical capital per worker was $250,694M. Using an output elasticity of capital of one-third, the U.S. physical capital per worker produced a factor of 4.7 compared to Vietnam in 1970. In 1970, the U.S. physical capital per worker was about five times that of Vietnam, indicating that the average American worker had roughly five times the productive capacity. Vietnam's heavy reliance on manual labor as well as fighting a war greatly reduced the country's ability to produce goods efficiently. Formula below:

KUS / KVNM = $250,690M / $2,452M = 102.30.333

Ki = 4.7

Physical capital per worker in 2023 for K

Given Vietnam's employed population of 55M and physical capital at current PPPs (in mil. 2021US$) of $4,864,515M, physical capital per worker was $88,517M. The United States had an employed population of 168M and physical capital at current PPPs (in mil. 2021US$) of $80,823,096M, the physical capital per worker was $482,412M. Using an output elasticity of capital of one-third, the U.S. physical capital per worker produced a factor of 1.8 compared to Vietnam in 2023. By 2023, Vietnam was able to greatly narrow the productive capacity gap, showing the country's significant investment in physical infrastructure over the years. Formula below:

KUS / KVNM = $482,412M / $88,517M = 5.40.333

Ki = 1.8

Figure below shows the physical capital from 1970 to 2023, showing drastic narrowing in recent years:

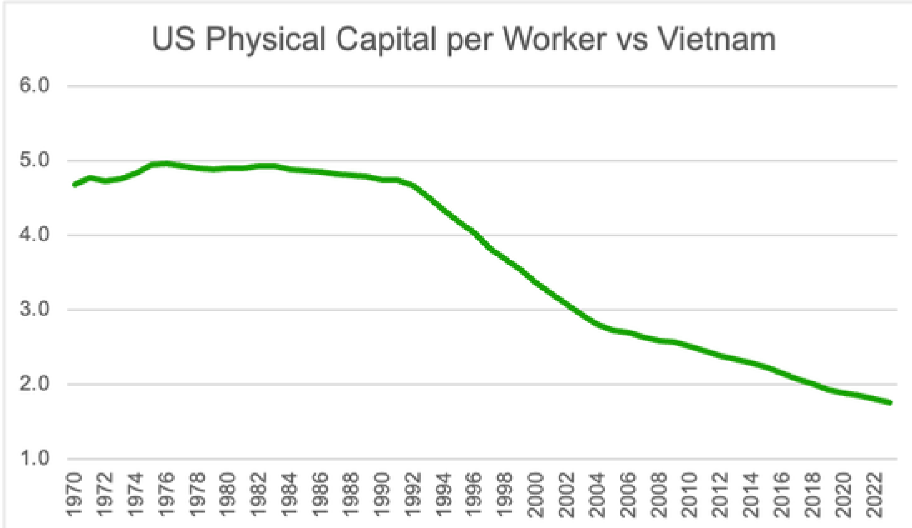

Human capital per worker in 1970 for H

Using the Barro-Lee data, this paper computes the human capital contribution to each country's economy by looking at return per year of schooling. In 1970, Vietnam average years of schooling across its entire population was 3.6 years, compared to 10.8 years in America. Using the Mincer earnings function where the coefficient 0.10 signifies that each additional year of education is associated with an approximate 10% increase in earnings. This indicates that the United States, with an average of 10.8 years of schooling, has a human capital factor of 2.9, whereas Vietnam's 3.6 years of schooling corresponds to a factor of 1.4. This calculation highlights the difference in human capital between the two countries. Using an output elasticity of human capital of two-thirds, a U.S. worker represents 1.6 times the human capital of a Vietnamese worker, reflecting differences in years of schooling and their impact on productivity potential in the workplace. Calculation below:

*HUS = exp(0.10*10.8) = 2.9*
*HVNM = exp(0.10*3.6) = 1.4*
HUS / HVNM = 2.9 / 1.4 = 2.10.667
Hi = 1.6

Human capital per worker in 2015 for H

Again, using the Barro-Lee data and the Mincer function, this paper computes the human capital contribution to each country's economy by looking at return per year of schooling. In 2015, the exponential function indicates that the United States, with an average of 13.2 years of schooling, has a human capital factor of 3.7, whereas Vietnam's 7.6 years of schooling corresponds to a factor of 2.1. This calculation shows how the gap has somewhat narrowed. Using an output elasticity of human capital of two-thirds, in the new average, a U.S. worker represents 1.5 times the human capital of a Vietnamese worker. Calculation below:

*HUS = exp(0.10*13.2) = 3.7*
*HVNM = exp(0.10*7.6) = 2.1*
HUS / HVNM = 3.7 / 2.1 = 1.80.667
Hi = 1.5

Given theminimalchangein human capital, I did not include a graph.

Finding Total Factor Productivity (TFP) for A

Now that we have the data points for the three key inputs for 1970 and 2023, we can back out TFP. This factor will show the productivity variation, not represented by differences in labor and capital, between the United States and Vietnam in 1970 and 2023. Using the formula below we can back out TFP as a residual:

Y1 = Ai (Ki)α (Hi)1-α
1970: 25.6 = Ai (4.7 x 1.6) Ai = 25.6 / (7.5) =
Ai = 3.4
2023: 6.3 = Ai (1.8 x 1.5) Ai = 6.3 / (2.7) =
Ai = 2.3

Figure below shows the TFP from 1970 to 2023, showing narrowing in production variation to the US in recent years:

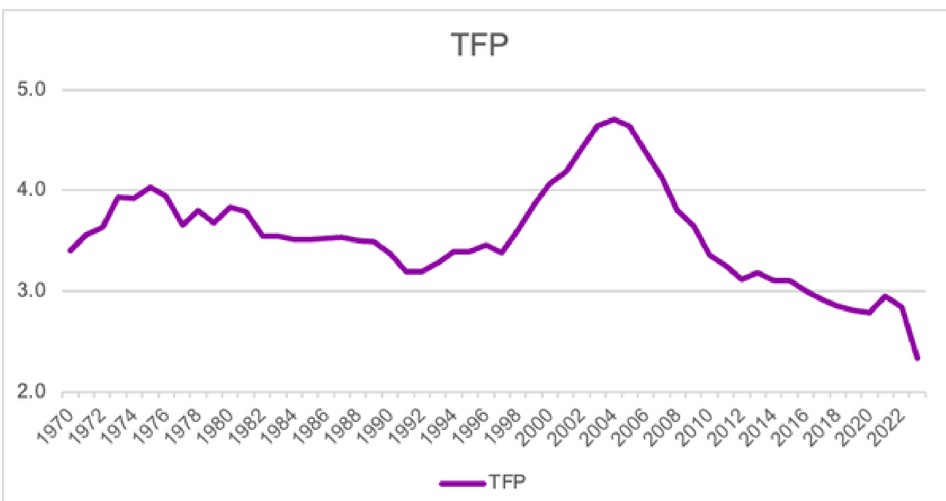

RESULTS

These calculations utilize the development accounting framework to show where the main gaps are, breaking down the US-Vietnam development gap into four specific factors. The US-Vietnam output per worker gap in 1970 was 26.6 which, by 2023, had narrowed to a factor of 6.3. This shows the vast improvement in incomes and quality of life the Vietnamese people experienced since 1970.

Significant investment and improvement in the infrastructure of the country resulted in the physical capital gap seeing strong progress, decreasing from 4.7 in 1970 to 1.8 in 2023. Despite the improvements, the existing output and physical capital gap can be attributed to a further need to expand roads and grid access out to rural areas. Areas where agriculture is still a main staple of the economy. The country's heavy exposure in agriculture, where 33% of the working population is employed, is a clear driver of its gap with the US. The country has set goals to accelerate the transition into higher value industries, but this transition will be very slow.

Human capital remained relatively stable, with the gap narrowing slightly from a factor of 1.6 in 1970 to 1.5 in 2015, based on the most recent Barro-Lee study. This is an interesting finding, given how other developing nations compare to the United States. Further research suggests that Vietnam's long-standing emphasis on education may account for this result. Despite having impressive access to primary and lower secondary education, the country struggles to provide adequate upper secondary and tertiary education opportunities. The following policy section will examine this human capital advantage in greater depth.

Over the 50-year period, TFP proved to be another variable that changed notably, with the relative TFP level improving from a factor of 3.4 in 1970 to 2.3 in 2023. This TFP factor still represents a considerable gap in the technological and innovation landscape of the country. This gap is highly attributed to lack of access to reliable internet as well as electrical grid insecurities that can cause rolling blackouts, disrupting productivity. As the country transitions to a more industrially focused economy the TFP gap will likely narrow even more. Since 1970 Vietnam has made significant progress in closing the output and physical capital gaps with the US, while human capital exhibited modest progress due to its already strong position. To accelerate this success and narrow the gap with advanced economies like the United States, this paper will now examine key policy levers necessary for maximizing Vietnam's future growth potential.

Resulting gaps summarized by category from 1970 to 2023 in the table below:

Year	Output Gap	P Capital Gap	H Capital Gap	TFP Gap
1970	25.6	4.7	1.6	3.4
2023	6.3	1.8	1.5	2.3

POLICY

Vietnamese Electrification and Power Generation

Following the US-Vietnam war, the electrical infrastructure of the country was completely destroyed. The reforms of the 1980s identified several key areas where significant investment was needed, and electrification of both urban and rural areas was deemed vital. This led to the founding of Vietnam Electricity in the early 1990s that focused on renewable hydroelectric power for rural areas to keep costs low (Borgen Project). The program proved to be a resounding success, as of 2017, 100% of Vietnam's population has access to electricity. To enhance grid reliability, the government approved a comprehensive new energy plan in 2023. The Vietnam's Power Development Plan VIII (PDP8) prioritizes a massive scale-up of wind and solar power alongside the development of pumped-storage hydropower and Battery Energy Storage Systems (BESS) to create a storable, dependable, and sustainable foundation for its future energy needs (Norton Rose Fulbright). This is a key initiative for further reducing the technology and physical capital gap between America and Vietnam.

Several BESS applications have already been built and are serving as flagship examples of the success of this initiative in both urban and rural areas. A 750 kW BESS project at the PECC2 Innovation Hub in Ho Chi Minh City serves as a case study for application in a commercial setting (Hailongjsc). By utilizing the building's existing energy management system, in this case the roof located solar panels, the system can charge during off-peak times and discharge during peak hours. Lowering energy costs and establishing a reliable power backup system. This system is also being applied successfully in very remote areas as well. The 650 kW system on Bach Long Vi Island in the Gulf of Tonkin is an example of bringing reliable power to places with a long history of grid instability. These are two very small but successful projects that are just the beginning of a massive initiative that targets 10,000 MW and 16,300 MW of battery backup capacity by 2030. This is coupled with other renewable energy projects of equal magnitude.

The Power Development Plan VIII (PDP8) is calling for massive investment and buildout of additional renewable energy projects to support the existing grid. Current 2030 targets for onshore and nearshore wind capacity are between 26,000 MW and 38,000 MW. Offshore wind capacity is projected at a more fixed target of 17,000 MW (Norton Rose Fulbright). This massive capital investment in both renewable substitutes and new transmission infrastructure is only justified due to the need for greater grid reliability in Vietnam's industrial zones. Similar policy was applied in rural Zambia and Tanzania where industrial areas running inefficiently on diesel generators were electrified, bringing significant increases to production efficiency (Walter & Moneke). The study found that areas with pre-existing use of electricity benefited substantially more than areas where it was newly introduced, as is relatable to Vietnam. This study did not describe a discernable macroeconomic impact but focused on the microeconomic impact of firms in the region. A study conducted by Apeti and Li quantified the effects power outages have on the TFP of developing nations. The study used firm level data to estimate the causal impact of power reliability on productivity across developing nations. The paper found that firms in countries with grid unreliability saw an 11% decrease in workplace productivity. These findings can be directly attributed to Vietnam's current electrical status. With nearly 100% of the country having access to some form of electricity, expanding upon the existing infrastructure will have a significant positive effect on TFP. Taking the previously calculated TFP gap of 2.3 for 2023 and applying the 11% boost, the new TFP gap with the US is 2.1. The output gap narrows to 5.7 from 6.3 once the new TFP number is factored in (with all K and H at above 2023 levels). Calculations below:

2023: Old Ai = 2.3 = (2.3 / 1.11) =

New Ai = 2.1

Y1 = Ai (Ki)α (Hi)1-α

Y1 = 2.1 x 1.8 x 1.5

New Y1 = 5.7

Vietnamese Education

Education has been a very important policy initiative for Vietnam since the country gained independence in 1945. The government prioritized education from the onset by launching a universal 12-year general education system that, by 1958, had nearly eliminated illiteracy in the country (Vietnam+). Given the country's focus on education, Vietnam's human capital ratio has remained close to that of developed nations, even during periods of war and economic hardship. Vietnam's early investment in education established an impressive base of human capital that would later underpin the country's transition to the modern, market-oriented economy that is seen today. Looking ahead, the country is setting ambitious goals in its upper secondary and tertiary education programs. "By 2035, Vietnam aims to complete universal upper secondary education and see at least two universities ranked among the world's top 100 in specific fields" (Vietnam+). Through these ongoing reforms, Vietnam seeks to bridge its remaining human capital and technological gaps, strengthening its foundation for future economic success.

Vietnam's ambitious goals for the future of its education are noted in Resolution No. 71-NQ/TW, issued by the politburo in August of 2024 (Vietnam.VN). This initiative sets out a strategic plan for further education reform through 2045. Two specific quantitative goals in the initiative focus on achieving a score of above 0.8 on the United Nations human development index (HDI) and reducing the education inequality index below 10%. Surpassing 0.8 HDI would put Vietnam in the "Very High Human Development" tier (UNDP). To achieve this, Vietnam needs to increase its current years of schooling for its adult population to levels closer to that of the US and European countries (12 to 13 years of schooling). Vietnam will also need to focus on reducing inequality in its education system by focusing on rural education programs. These programs are already in the works to ensure universal education for all children in the country regardless of location. Achieving an education inequality index below 10% is an ambitious goal, in 2019 their Inequality-adjusted HDI (IHDI) sits at 16.2% (UNDP Briefing Note). If these goals are achieved, a notable narrowing of Vietnam's human capital gap will occur. I am assuming in the below calculation that if Vietnam can achieve 12 average years of schooling for its adult population the resulting human capital per worker increases dramatically to 3.3 from 2.1. Using an output elasticity of human capital of two-thirds, in the new average, a U.S. worker represents only 1.1 times the human capital of a Vietnamese worker. The output gap narrows dramatically to 4.6 from 6.3 once the new TFP number is factored in (with all K and A at above 2023 levels). Calculation below:

Old HVNM = exp(0.10*7.6) = 2.1

HVNM = exp(0.10*12) = 3.3

HUS / HVNM = 3.7 / 3.3 = 1.10.667

New Hi = 1.1

Y1 = Ai (Ki)α (Hi)1-α

Y1 = 2.3 x 1.8 x 1.1

New Y1 = 4.6

National Digital Transformation Program

Vietnam has benefited from an economic model centered on low-cost labor and capital accumulation. This approach successfully attracted large inflows of foreign investment to develop export-oriented industries. However, the model now faces diminishing returns as the country seeks to improve total value productivity. In 2021, the government established "the Vietnam National Committee on Digital Transformation" (International Trade Administration), a strategy focusing on accelerating the digitalization of both the public and private sector through the adoption of core technologies like AI, IoT, and 5G. This policy aims to bridge the remaining TFP gap and accelerate the country's ascent to upper-income status.

With an annual growth rate approaching 20%, Vietnam has emerged as the fastest-growing digital economy in Southeast Asia. This acceleration has necessitated a wave of new infrastructure policies designed to accommodate the sector's rapid expansion. Access to fast, low latent internet will be a key factor in supporting the digital sector's continued growth. Approved in January 2024, Decision 36/QD-TTg (Vietnam Briefing), suggests key build out targets for telecommunication and data center infrastructure. By 2030, the directive plans to offer broadband network coverage with speeds of over one gigabit per second to 100% of users and 5G network coverage to 99% of the population. To reduce internet bottlenecks that have slowed development in the past, the construction of 6 to 8 new international undersea fiber optic cables is also planned for 2030. The country plans to complete three data center clusters that target to give 70% of Vietnamese businesses access to cloud computing services. These strategic measures aim to bolster the Vietnamese digital economy, accelerating its convergence with technology-centric societies such as China and the US.

Research into the digitalization of developing countries confirms that robust policy implementation acts as a catalyst for positive economic outcomes. The paper by Hjort & Tian quantifies the extent that digitalization has on emerging markets. The effects range from vast improvements in firm performance to the expansion of the labor market coupled with an increase in wages. For the purposes of this paper, I will focus on the effects that successful digital adoption has on wages. The paper notes that a 10% increase in 3G internet coverage in a region increases wages by 2.1%. This indicates that access to faster internet allows for higher paying jobs as well as the ability to perform more efficiently in the workplace. Given the marginal product theory, I have applied the 2.1% wage increase to output per worker since real wages are roughly equal to the marginal product of labor. By increasing the existing output per worker number by 2.1% the output gap slightly narrows to 6.1 from 6.3. Calculation below:

YUS / YVNM = $153,049M / $24,373M

Old Y1= 6.3

YUS / YVNM = $153,049M / ($24,373M x 1.021)

YUS / YVNM = $153,049M / 24,885M

New Y1= 6.1

Policy Results

The cumulative impact of Vietnam's new policy architecture is quite meaningful when applied to the development accounting framework. Energy, educational, and digital reforms reveal a clear trajectory toward high-income status. By addressing grid reliability through the Power Development Plan VIII (PDP8) and Battery Energy Storage Systems (BESS), Vietnam can recover an estimated 11% productivity loss currently attributed to power instability. Applying this boost to the 2023 data improves the TFP gap from 2.3 to 2.1, resulting in the overall Output per Worker gap narrowing from 6.3 to 5.7. The most significant potential convergence comes from human capital accumulation. If reforms under Resolution No. 71-NQ/TW succeed in raising the average years of schooling to 12 years, the human capital gap between the US and Vietnam would shrink to 1.1 from 1.5. This adjustment would dramatically reduce the Output Gap to 4.6 from 6.3. This highlights education as the most powerful lever for long-term growth. As Vietnam transitions from an agrarian economy to one dominated by industry, the demand for a highly skilled workforce will only intensify. With strategic policies like Resolution No. 71 in place, the country is well-positioned to achieve the substantial gains in human capital necessary to meet this need. The implementation of the National Digital Transformation Program, aimed at expanding 5G and fiber optic coverage, is projected to increase wages and labor productivity by 2.1%. When applied to the existing data, this efficiency gain narrows the Output Gap slightly to 6.1 from 6.3. Vietnam has come a remarkable distance since the complete infrastructure collapse of the 1980s post-war era but sustaining this trajectory now requires moving beyond simple access to electricity to prioritizing high-speed connectivity and grid reliability. These calculations illustrate that while digital and energy improvements offer meaningful efficiency gains, substantial investment in human capital remains the critical factor for Vietnam to achieve high-income, tiger economy status.

Summary table below describes the output result of each scenario:

Scenario	Policy result	Equation	US to VNM Output Gap
A (Energy Fix)	11% boost in TFP	Old A = 2.3 = (2.3 / 1.11) = New A = 2.1	**5.7**
B (Education Fix)	12 avg. yrs of schooling	VNM = exp(0.10*12) = 3.3 US / VNM = 3.7 / 3.3 = 1.1^2/3 New H = 1.1	**4.6**
C (Digital Fix)	Internet connectivity improves wages by 2.1%	US / VNM = $153,049M / ($24,373M x 1.021) US / VNM = $153,049M / 24,885M = New Y = 6.1	**6.1**

Resulting gaps summarized by category from 1970 to 2023 in the table below (included again here for reference):

Year	Output Gap	P Capital Gap	H Capital Gap	TFP Gap
1970	25.6	4.7	1.6	3.4
2023	6.3	1.8	1.5	2.3

Data Sources

1. Consulate of the Socialist Republic of Viet Nam in New York. “About Viet Nam.” Consulate of the Socialist Republic of Viet Nam in New York, https://www.vietnamconsulate-ny.org/about/1/about-vietnam.html

2. MANAGEMENT DEPARTMENT OF CỔ LOA VESTIGE SITE. “Resistance Against the Qin Army.” Thành Cổ Loa Website, 23 Dec. 2024, https://thanhcoloa.vn/en/resistance-against-the-qin-army

3. Consulate of the Socialist Republic of Viet Nam in New York. “Brief History of Viet Nam.” Consulate of the Socialist Republic of Viet Nam in New York, https://www.vietnamconsulate-ny.org/about/2/brief-history-of-vietnam.html

4. Hoàng, Bryan. ““The Civilizing Mission:” French Colonialism in Vietnam (1858-1954).” Vietnamese American Experiences Model Curriculum, We Are California, University of California, Davis, https://camodelcurricula.ucdavis.edu/vietnamese-american-experiences/civilizing-mission-french-colonialism-vietnam-1858-1954#:~:text=Modern%20historians%20understand%20this%20as,their%20international%20trade%20and%20capitalism.

5. HISTORY.com Editors. “Vietnam Declares Its Independence from France / September 2, 1945.” HISTORY, A&E Television Networks, 31 Jan. 2025, https://www.history.com/this-day-in-history/september-2/vietnam-independence-proclaimed.

6. Consulate of the Socialist Republic of Viet Nam in New York. “Brief History of Viet Nam.” Consulate of the Socialist Republic of Viet Nam in New York, https://www.vietnamconsulate-ny.org/about/2/brief-history-of-vietnam.html

7. Government of Viet Nam. “The Period Before Western Capitalists’ Penetration.” Viet Nam Government Portal, https://vietnam.gov.vn/economy-68968.

8. Gunn, Geoffrey “The Great Vietnamese Famine of 1944-45 Revisited.” Asia-Pacific Journal: Japan Focus, vol. 19, no. 11, 2023, https://apjjf.org/wp-content/uploads/2023/11/article-2768.pdf.

9. International Monetary Fund. “Vietnam: Raising Millions Out of Poverty.” August 2018, www.imf.org/en/Countries/VNM/vietnam-raising-millions-out-of-poverty.

10. EconScope. “Vietnam: GDP (Nominal).” EconScope, https://econscope.com/economies/viet-nam/gdp-nominal.

11. Feenstra, Robert, Heston, Alan. Penn World Table, Version 11.0. Groningen Growth and Development Centre, University of Groningen, https://www.rug.nl/ggdc/productivity/pwt/pwt-database.

12. Feenstra, Robert, Heston, Alan. Penn World Table, Version 11.0. Groningen Growth and Development Centre, University of Groningen, https://www.rug.nl/ggdc/productivity/pwt/pwt-database.

13. Barro, Robert and Jong-Wha Lee, 2021, "A New Data Set of Educational Attainment in the World, 1950-2015." Journal of Development Economics, vol 104, pp.184-198.

14. Barro, Robert and Jong-Wha Lee, 2021, "A New Data Set of Educational Attainment in the World, 1950-2015." Journal of Development Economics, vol 104, pp.184-198.

15. Parikh, Treya. “The Electrification of Vietnam.” The Borgen Project, 16 Mar. 2020, https://borgenproject.org/the-electrification-of-vietnam/.

16. Trung, Vu Le and Anh, Vu Ha. “Development of Battery Energy Storage Systems in Vietnam.” Norton Rose Fulbright, 31 Oct. 2025, https://www.nortonrosefulbright.com/en-us/knowledge/publications/7eb0008e/development-of-battery-energy-storage-systems-in-vietnam.

17. Case Study BESS at PECC2 Smart Building: Energy Future for Vietnamese Enterprises." Hailong JSC, 19 Sept. 2025, https://en.hailongjsc.vn/case-study-bess-pecc2-innovation-hub/.

18. Trung, Vu Le and Anh, Vu Ha. “Development of Battery Energy Storage Systems in Vietnam.” Norton Rose Fulbright, 31 Oct. 2025, https://www.nortonrosefulbright.com/en-us/knowledge/publications/7eb0008e/development-of-battery-energy-storage-systems-in-vietnam.

19. Walter, Torsten Figueiredo, and Niclas Moneke. "When Does Electrification Work? Evidence from Sub-Saharan Africa." 9 Aug. 2024.

20. Apeti, Ablam Estel, and Alpha Ly. "Power constraints and firm-level total factor productivity in developing countries." Energy, vol. 309, 2024, 132871.

21. VietnamPlus. “Vietnam’s Education Reforms Shape Path for Comprehensive Development.” Vietnam+, 27 Aug. 2025, https://en.vietnamplus.vn/vietnams-education-reforms-shape-path-for-comprehensive-development-post325474.vnp.

22. VietnamPlus. “Vietnam’s Education Reforms Shape Path for Comprehensive Development.” Vietnam+, 27 Aug. 2025, https://en.vietnamplus.vn/vietnams-education-reforms-shape-path-for-comprehensive-development-post325474.vnp.

23. Communist Party of Viet Nam. “Toàn văn Nghị quyết Số 71/NQ-TW của Bộ Chính trị về Đột phá phát triển giáo dục và đào tạo.” Vietnam.vn, https://www.vietnam.vn/en/toan-van-nghi-quyet-so-71-nq-tw-cua-bo-chinh-tri-ve-dot-pha-phat-trien-giao-duc-va-dao-tao.

24. "Vietnam." Human Development Reports, United Nations Development Programme, hdr.undp.org/data/country-profiles/deu.

25. United Nations, Vietnam. "Brief Viet Nam - HDI Update." United Nations, Vietnam, Dec. 2019, vietnam.un.org/sites/default/files/2019-12/2.%20Brief%20Viet%20Nam%20-%20HDI%20update_EN_06%20Dec.pdf.

26. U.S. International Trade Administration. "Vietnam—Digital Economy." Country Commercial Guide: Vietnam. U.S. Department of Commerce, 2024, www.trade.gov/country-commercial-guides/vietnam-digital-economy.

27. Vietnam Briefing. "Information and Communication Infrastructure Master Plan Approved." Vietnam Briefing News, www.vietnam-briefing.com/news/information-communication-infrastructure-master-plan-approved.html/

28. Hjort, Jonas, and Lin Tian. "The Economic Impact of Internet Connectivity in Developing Countries." Annual Review of Economics, vol. 17, 2025, pp. 111

Political Trust and Elite Cue Receptivity in Southeast Asian Attitudes toward the United States and China

Lucas Zeng

ABSTRACT
How does political trust shape public opinion toward major powers? Existing research on public diplomacy and foreign policy opinion emphasizes the effects of elite cues, international events, and external influence efforts, but pays less attention to how domestic political trust conditions the reception of those signals. This paper argues that political trust increases the extent to which public attitudes align with the foreign policy orientation of national elites. To test this argument, I combine individual-level data from Wave 5 of the Asian Barometer Survey with country-level elite opinion data from the ISEAS State of Southeast Asia Survey across seven ASEAN countries. Using a stacked OLS model, I examine whether trust conditions the relationship between elite foreign policy orientation and public opinion toward the United States and China. The results show that individuals with higher levels of trust in government are more likely to hold attitudes toward the United States that align with elite preferences, while the corresponding effect for China is much weaker.

Keywords: Public opinion, Southeast Asia, Public diplomacy, United States, China, Trust, Elite cues

INTRODUCTION
Over the last couple of decades, research on public diplomacy has developed a substantial body of literature as academics and government officials alike recognize public opinion as a consequential factor in foreign policy, affecting both the choices leaders make and the credibility of the policies they pursue (Aldrich et al., 2006; Milner and Tingley, 2011; Goldsmith and Horiuchi, 2012; Hartig, 2016; Peterson, 2002). Thus, public opinion is not merely a reaction to global events but plays an active role in shaping international outcomes. Leaders anticipate public responses when formulating foreign policy, and foreign policy preferences are better understood as the product of domestic political processes rather than just strategic responses to external threats or material interests alone (Kertzer and Zeitzoff, 2017; Saunders, 2022). In order to utilize this growing understanding of public diplomacy, research finds that states themselves seek to shape foreign public opinion as a kind of strategic resource (Nye, 2004; Goldsmith, Horiuchi and Matush, 2021; Chapple-Sokol, 2013).

Governments often invest in high-level diplomatic visits, cultural exchange programs, international broadcasting, and symbolic gestures hoping that favorable foreign public opinion can facilitate cooperation and increase their nation's international reputation (Melissen, 2005; Cull, 2008; Nye, 2008; Nye and Kim, 2019; Shafie, 2025). Despite these advances, much of the literature on the relationship between public opinion and foreign policy preferences focuses on the impact of external stimuli, such as diplomatic initiatives, economic ties, or military incidents, while paying comparatively little attention to which domestic factors might influence the public's foreign policy preferences. More specifically, there is a gap of existing work theorizing how trust in domestic political institutions influences public opinion on foreign powers.

This omission is notable given the extensive evidence for two factors. Firstly, it is well-established that trust in domestic political institutions affects political participation, regime legitimacy, and support for domestic policy outcomes (Hetherington, 2005; Levi and Stoker, 2000; Easton, 1975). If trust in domestic institutions structures how the public perceives domestic policy, it plausibly also shapes how they evaluate foreign policy. Secondly, research on public diplomacy has noted that the public's opinion on international events is often interpreted through the framing and cues of domestic political elites (Zaller and Chiu, 1996; Guisinger and Saunders, 2017; Berinsky, 2009; Kertzer, 2017).

This gap is especially consequential in the context of contemporary great-power competition. The rivalry between the United States and China has intensified efforts by both states to cultivate favorable foreign public opinion, particularly in regions where alignment is contested. Public attitudes in these third countries matter not only for the fostering of general soft power, but also in judging the feasibility of security cooperation or economic agreements. Southeast Asia represents a particularly important context in which to examine these dynamics. The region occupies a central position in the US-China competition, yet its governments and publics have historically emphasized non-alignment, sovereignty, and strategic hedging rather than exclusive partnership with any single great power (Goh, 2016). At the same time, Southeast Asian countries exhibit substantial variation in regime type, media environments, and levels of public trust in political institutions, which makes the region well suited as a case study to examine how domestic political trust shapes foreign attitudes.

In this article, I investigate whether the public in Southeast Asian countries' trust in domestic institutions impacts their alignment with the political elites' foreign policy views on China and the United States. To answer this question, I draw from the individual-level survey data of Wave 5 of the Asian Barometer Survey (ABS) and the 2021 ISEAS opinion survey on Southeast Asian elites

(Asian Barometer Project, 2021; Seah et al., 2021). The ABS is a large, cross-national survey that captures political attitudes, institutional trust, and foreign policy perceptions across a diverse set of Asian countries. Wave 5 includes samples from multiple Southeast Asian states and contains detailed questions on trust in political institutions as well as evaluations of the influence of both the United States and China, and provides 4,140 responses from the public in multiple Southeast Asian countries from 2018 to 2021 answering the relevant survey questions, providing an opportunity to investigate if there is a statistically significant relationship between an individual's level of trust in their own government, and their alignment of domestic elites' opinion of foreign powers such as the United States and China.

HYPOTHESIS

This study argues that political trust shapes how individuals interpret elite foreign policy signals. Citizens often rely on domestic political elites to make sense of complex foreign policy issues, particularly when those issues are distant or strategically ambiguous. If individuals trust their government, they may be more likely to view elite foreign policy positions as credible and legitimate guidance. As a result, political trust should strengthen the relationship between elite foreign policy orientation and public attitudes toward major powers. Should this logic hold, the analysis should result in the following finding: Individuals who express higher levels of trust in their government will be more likely to hold foreign policy attitudes that align with the foreign policy orientation of national elites.

RESEARCH DESIGN

Public Trust in Government and Foreign Policy Preferences

To examine how political trust relates to foreign policy attitudes, I draw on individual-level survey data from Wave 5 of the ABS (Asian Barometer Project, 2021). The ABS is a cross-national public opinion survey designed to measure political values, institutional trust, and democratic attitudes across Asian societies. Wave 5 of the survey was conducted between 2018 and 2021 and includes nationally representative samples from several Southeast Asian countries. The survey employs stratified multistage probability sampling within each country to ensure demographic representativeness.

The primary independent variable in this study is trust in the national government. Respondents were asked how much they trust their national government, with responses recorded on a six-point scale ranging from complete distrust to complete trust. This measure captures citizens' confidence in the competence and legitimacy of their governing institutions, which prior research identifies as a key component of political support (Easton, 1975; Levi and Stoker, 2000; Hetherington, 2005).

The dependent variables measure respondents' evaluations of the international influence of the United States and China. In the ABS questionnaire, respondents were asked to rate how positive or negative they believe each country's influence is on their nation on a ten-point scale. These measures capture individual attitudes toward the relative influence and role of major powers in their domestic political and economic environment.

The analysis focuses on respondents from Southeast Asian countries included in Wave 5 for which comparable elite opinion data are available. After restricting the sample to these countries and removing responses with missing values for key variables, the final dataset contains 4,140 individual respondents. The models also include a range of control variables drawn from the ABS survey that may plausibly influence foreign policy attitudes, including demographic characteristics, political engagement, and economic perceptions. Additionally, the analysis incorporates country-, region-, and urban–rural-level fixed effects to account for structural differences in political and socioeconomic environments across survey respondents.

Elite Foreign Policy Preferences

To measure elite foreign policy orientation, this study draws on the ISEAS State of Southeast Asia Survey Report 2021, an annual survey conducted by the ISEAS–Yusof Ishak Institute. The survey collects responses from regional policy elites, including government officials, business leaders, academics, journalists, and civil society representatives. These respondents are selected for their professional engagement with regional political and economic issues and therefore provide a useful indicator of elite-level perceptions of geopolitical alignment in Southeast Asia (Seah et al., 2021).

The ISEAS survey asks respondents which major power they would prefer their country to align with if forced to choose between the United States and China. Responses are reported as country-level percentages indicating the proportion of elite respondents favoring alignment with each power. These percentages provide a measure of the relative orientation of national policy elites toward the United States and China.

In this study, these elite preferences are implemented as country-level measures of elite support for the United States and China, respectively. Because these measures vary only at the country level, they capture differences in elite foreign policy orientation across Southeast Asian states rather than variation among individuals within countries. The elite preference variables are merged with the ABS dataset by country, assigning each survey respondent the corresponding elite orientation for their country. Combining individual-level public opinion data with country-level elite preference measures allows for an examination of whether citizens' trust in their domestic political institutions conditions the relationship between elite foreign policy orientation and mass attitudes toward major powers. In other words, this design makes it possible to test whether individuals who express greater trust in their government are more likely to hold foreign policy views that align with the preferences of national elites.

Empirical Strategy

To evaluate whether political trust conditions the relationship between elite foreign policy orientation and public opinion, I estimate a series of ordinary least squares (OLS) regression models that combine the individual-level survey data from the ABS with the country-level elite preference measures from the ISEAS survey. Because the dependent variables measuring attitudes toward the United States and China are recorded on comparable ten-point scales, the analysis employs a stacked data structure that allows both outcomes to be analyzed within a single model.

In the stacked dataset, each survey respondent contributes two observations: one capturing their evaluation of the United States and one capturing their evaluation of China. An indicator variable identifies which country the observation refers to. This structure allows the model to estimate whether the conditioning effect of political trust differs between attitudes toward the United States and attitudes toward China while maintaining a common set of controls and fixed effects.

The key explanatory variables are political trust, elite foreign policy orientation, and their interaction. Elite orientation is implemented as the country-level percentage of elite respondents favoring alignment with either the United States or China in the ISEAS survey. Because elite preferences vary only at the country level, they capture differences in national elite orientation rather than individual-level variation within countries. The model includes an interaction between political trust and elite orientation to test whether individuals who express higher levels of trust in their government are more likely to hold attitudes that align with elite foreign policy preferences. Formally, the empirical model can be written as:

$$\begin{aligned} Opinion_{ict} = \beta_0 + \beta_1 Trust_i + \beta_2 EliteSupport_c + \beta_3 Target_t \\ + \beta_4 (Trust_i \times EliteSupport_c) \\ (1) + \beta_5 (Trust_i \times EliteSupport_c \times Target_t) \\ + \mathbf{X}_i' \boldsymbol{\theta} + \gamma_c + \delta_r + \lambda_u + \varepsilon_{ict} \end{aligned}$$

Here, $Opinion_{ict}$ represents respondent *i*'s evaluation of country *t* (the United States or China) in country *c*. $Trust_i$ measures the respondent's level of trust in their national government. $EliteSupport_c$ captures elite foreign policy orientation in country *c*. and $Target_t$ indicates whether the observation refers to the United States or China. The vector $\boldsymbol{X_i}$ includes individual-level control variables drawn from the ABS.

The terms γ_c, δ_r, and λ_u represent country, regional, and urban–rural fixed effects, respectively. The coefficient on the interaction between political trust and elite support captures whether trust conditions the relationship between elite orientation and mass attitudes. The three-way interaction with the target indicator tests whether this conditioning effect differs between attitudes toward the United States and attitudes toward China.

The identification of the interaction effect relies on two sources of variation. First, individual respondents within each country vary in their level of trust in their national government. Second, elite foreign policy orientation varies across countries as measured by the ISEAS survey. The interaction between these two variables therefore captures whether individuals with higher levels of political trust are more likely to express attitudes toward major powers that reflect the foreign policy orientation of elites in their country. Because elite preferences vary only at the country level, the models include country fixed effects and cluster standard errors at the country level to account for unobserved heterogeneity and within-country correlation in the error terms. This approach isolates the conditional relationship between political trust and elite orientation while controlling for structural differences across countries.

Table 1: Effect of Trust in Government on Alignment with Domestic Political Elite Foreign Preferences Using OLS Regression Estimates.

	Stacked OLS
Trust in Government	−0.049 (0.051)
Elite Support	0.030*** (0.007)
United States Indicator	0.184 (0.368)
Trust × Elite Support	−0.005 (0.004)
Trust × United States	0.115 (0.105)
Trust × Elite Support × United States	0.019*** (0.004)
N	4,140
R^2	0.102
Adjusted R^2	0.091
Country fixed effects	Yes
Region fixed effects	Yes
Urban/rural fixed effects	Yes
Country-clustered SEs	Yes

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure 1: Predicted Foreign Opinion at Low and High Levels of Trust in Government.

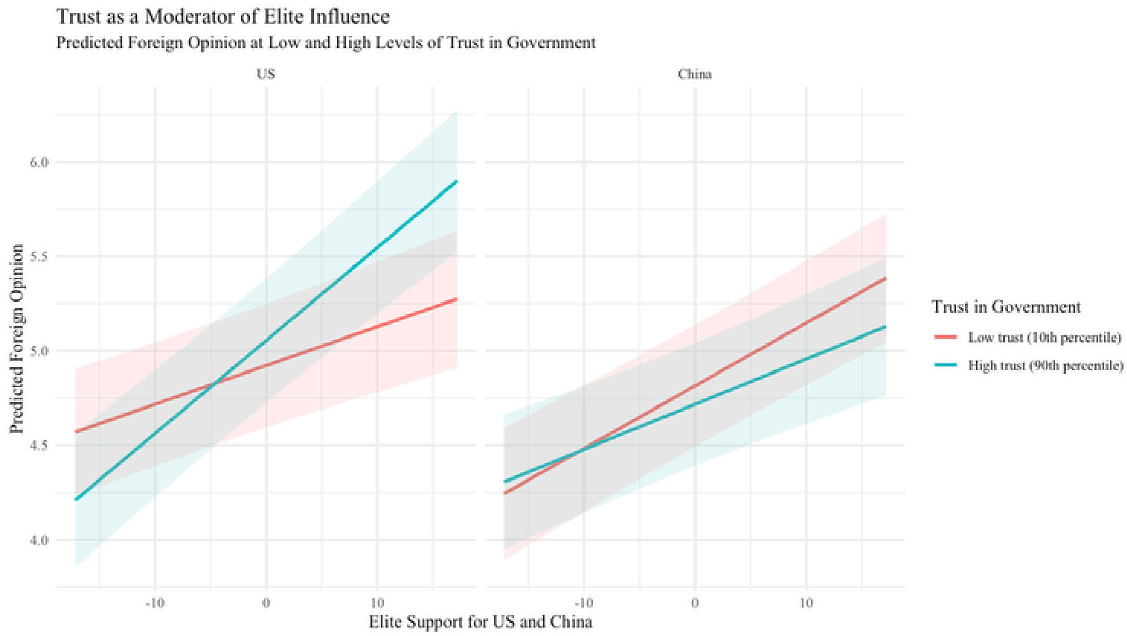

RESULTS

Political Trust and Mass–Elite Alignment

The central expectation of this study is that political trust conditions the relationship between elite foreign policy orientation and public opinion toward major powers. More specifically, individuals who express higher levels of trust in their national government should be more likely to hold attitudes that align with the foreign policy orientation of their country's elites. Table 1 provides support for this argument, but also shows that the strength of this relationship differs substantially between attitudes toward the United States and attitudes toward China.

The results indicate that political trust does not operate as a simple direct predictor of foreign attitudes. The coefficient on trust in government is statistically insignificant in the stacked model, suggesting that trust alone does not make respondents uniformly more favorable toward either the United States or China. Instead, the relationship between trust and foreign opinion depends on elite foreign policy orientation. This is consistent with the argument advanced in the hypothesis: political trust matters not because it directly shifts attitudes in a particular geopolitical direction, but because it coditions how strongly elite preferences are reflected in mass opinion. It is also worth noting that the estimated relationship is asymmetric across the two major powers. The conditioning effect of trust is strongest in evaluations of the United States while the same effect is much weaker in evaluations of China.

Evidence from the OLS Model

Table 1 displays the stacked OLS model, which jointly estimates public evaluations of the United States and China and includes standard errors clustered at the country level. In this specification, China serves as the baseline category, meaning that the coefficient on the interaction between trust and elite support captures the conditional relationship for attitudes toward China, while the three-way interaction with the United States indicator captures how that relationship differs for attitudes toward the United States.

The baseline interaction between trust and elite support is negative and statistically insignificant. This indicates that, for attitudes toward China, there is little evidence that individuals with higher levels of trust in government are systematically more likely to align with elite foreign policy orientation. In other words, the relationship between elite support for China and mass opinion toward China does not appear to be strongly conditioned by political trust.

The pattern is notably different for the United States. The three-way interaction between trust, elite support, and the United States indicator is positive and highly statistically significant. This is the main empirical result of the analysis. It indicates that the conditional effect of trust is significantly stronger for attitudes toward the United States than for attitudes toward China. Substantively, this means that in countries where elites are more supportive of the United States, individuals who trust their government more also tend to express more favorable attitudes toward the United States. Where elite support for the United States is weaker, this relationship declines accordingly.

Figure 1 provides an illustration of this relationship by plotting predicted values from the model across different levels of elite support. On the left side, Figure 1 shows that survey respondents with high levels of trust in their government (in the 90th percentile) align more strongly with domestic elite opinions of the United States. The slope for this bloc of survey respondents is clearly steeper, and have a much higher opinion of the US when elites from their country have a high opinion of the US, and vice versa. By contrast, survey respondents with low levels of trust in their government (in the 10th percentile) do not align as strongly with domestic elite opinions of the United States. The flatter slope representing this bloc of survey respondents indicates that their opinions on the US are not as influenced by elite opinion on the US as high trust respondents.

The plot on the right side of Figure 2 shows the corresponding pattern for China. Here, the slopes are much closer together, and the conditioning effect is substantially weaker. Visually, there is not much difference between respondents with low trust in government and high trust in government in the relationship between elite support for China and predicted opinion of China.

DISCUSSION

The central finding of this study is that political trust appears to condition the extent to which public attitudes toward major powers reflect the foreign policy orientation of national elites. This pattern is most clearly visible in evaluations of the United States. In countries where elites are more supportive of the United States, individuals who report higher levels of trust in their government are also more likely to express favorable views of the United States. By contrast, where elite support for the United States is weaker, the relationship between political trust and pro-US opinion is correspondingly attenuated. The results therefore suggest that trust in government does not operate simply as a direct predictor of foreign attitudes. Instead, it appears to shape how strongly elite foreign policy preferences are transmitted into mass opinion.

The evidence presented here is consistent with the view that trust functions as a mechanism of elite cue receptivity. Individuals who trust their domestic political institutions may be more likely to regard official foreign policy signals as credible, legitimate, or worthy of deference. In this sense, political trust may strengthen the link between domestic institutional legitimacy and public evaluations of external powers. Foreign policy attitudes are therefore not only responses to international events or material interests, but also mediated through the domestic political relationship between citizens and political elites.

There is the question of why this effect is so much stronger regarding the United States than China. One plausible interpretation of the asymmetric results is that elite foreign policy cues are clearer and more politically legible in the case of the United States than in the case of China. Research on elite cue-taking suggests that cue effects vary across issues depending on how salient, contested, and interpretable they are to the public. Elite signals tend to have stronger effects when they provide meaningful guidance on issues that citizens otherwise find difficult to evaluate, such as foreign policy, but these effects can weaken when citizens already have mixed or complicated views (Guisinger and Saunders, 2017). In the Southeast Asian context, this may matter because public views of the United States and China are shaped by different political and strategic concerns.

Scholars of Southeast Asian foreign policy have characterized regional responses to the US–China rivalry as forms of hedging rather than straightforward alignment, precisely because governments seek to avoid clear one-sided commitments while maintaining relationships with both powers. For the most part, states in Southeast Asia often either align clearly with the United States, such as in the case of the Philippines or Vietnam, or hedge their preferences between the two powers to **maxamize** returns and minimize risk (Goh, 2007; Kuik, 2016). This helps explain why the trust-conditioned alignment effect is strongest for attitudes toward the United States. Where elite support for theUnited States is higher, trust in government appears to make citizens more receptive to those elite signals. For China, however, trust does not generate the same degree of alignment, likely because the meaning of elite support for China is less straightforward in a region characterized by strategic hedging.

Implications

These findings have broader implications for how scholars understand public diplomacy and the domestic foundations of foreign policy opinion. Much of the public diplomacy literature emphasizes the role of external efforts to shape foreign publics through diplomatic visits, symbolic messaging, media outreach, and cultural engagement (Nye, 2004; Cull, 2008; Goldsmith, Horiuchi and Matush, 2021). That work has shown that states attempt to cultivate favorable images abroad because public opinion in foreign countries can influence the credibility and political feasibility of international cooperation. The results of this study suggest, however, that the effectiveness of such efforts may depend not only on the content of public diplomacy efforts, but also on the domestic political context in which those efforts are received.

Therefore, the analysis suggests that foreign policy opinion is filtered through domestic political trust. Individuals do not respond to international actors in a political vacuum. Rather, their views of major powers appear to be shaped in part by whether they trust the domestic institutions and elites that interpret those external relationships for them. This implies that public diplomacy may be most effective not when it directly changes attitudes on its own, but when it reinforces or interacts with elite signals that citizens already regard as credible.

This point also contributes to the broader literature on foreign policy opinion. Research in international relations and political behavior has long shown that mass attitudes toward foreign affairs are not simply the product of material interests or exogenous international events, but are shaped by elite cues, partisan framing, and information environments (Zaller, 1992; Baum and Potter, 2008; Berinsky, 2009; Kertzer and Zeitzoff, 2017). The findings here extend that insight by suggesting that political trust is an important part of the mechanism through which elite cues operate. Trust may increase the willingness of citizens to accept official foreign policy signals as meaningful guidance, thereby strengthening the connection between elite foreign policy orientation and public opinion.

Limitations

Several limitations should be kept in mind when interpreting these findings. First, the analysis is based on cross-sectional observational data, which limits the ability to make strong causal claims. Although the models include a wide range of controls as well as country, region, and urban–rural fixed effects, the results should still be understood as evidence of association rather than definitive proof that political trust causes greater mass–elite alignment in foreign policy attitudes. It remains possible that unobserved factors, such as media exposure, nationalism, or broader regime support, shape both trust in government and views of foreign powers.

Second, the elite opinion measure is available only at the country level and is based on a relatively small number of Southeast Asian cases. This means that the key interaction is identified from cross-country variation in elite foreign policy orientation rather than from within-country variation over time. As a result, the analysis is necessarily limited in the degree to which it can disentangle elite signaling from other country-level characteristics that may covary with it. The use of country-clustered standard errors helps address concerns about within-country correlation, but the small number of country clusters still calls for caution in statistical inference.

Additionally, the ISEAS survey captures elite preferences in aggregate form rather than the specific messages or frames that the public are actually exposed to. In other words, the measure provides a useful indicator of national elite orientation, but it does not directly observe the mechanisms through which elite cues reach the public. Future research could improve on this by incorporating speeches, media coverage, party statements, or public diplomatic messaging in order to trace more directly how elite foreign policy signals are transmitted and interpreted.

A further limitation of this study is temporal. The ABS and ISEAS data used here reflect attitudes and elite orientations measured around 2021 and therefore do not capture subsequent shifts in Southeast Asian perceptions of the United States and China. This matters because regional views of the United States may have changed meaningfully in response to the policy choices of the Trump administration after 2025, such as high threatened tariff rates on several Southeast Asian exporters, and the sharp reduction in US foreign assistance and diplomatic spending. These actions plausibly altered how both elites and publics in Southeast Asia evaluate the credibility, reliability, and desirability of the United States as a strategic and economic partner. The ISEAS State of Southeast Asia 2025 Survey reports that “the new US leadership” was among the region’s top geopolitical concerns, indicating that shifts in Washington were already shaping elite perceptions by early 2025. The same survey also shows that views of the major powers remain fluid, reinforcing the possibility that the relationships estimated in this paper may not fully describe the current political environment in Southeast Asia (Seah et al., 2025).

CONCLUSION

This paper has examined whether political trust conditions the relationship between elite foreign policy orientation and public attitudes toward major powers in Southeast Asia. Drawing on individual-level data from the ABS and country-level elite opinion data from the ISEAS State of Southeast Asia Survey, I argued that trust in domestic political institutions does not simply make citizens more favorable toward one foreign power or another. Instead, political trust shapes how strongly elite foreign policy preferences are reflected in mass opinion.

The empirical results provide support for this argument. Using a stacked OLS model with country-clustered standard errors, the analysis shows that the relationship between elite foreign policy orientation and public opinion is significantly stronger among individuals with higher levels of trust in government. At the same time, this effect is not uniform across the two major powers examined here. The conditioning effect of trust is substantially stronger for attitudes toward the United States than for attitudes toward China. This suggests that political trust functions as a mechanism of elite cue receptivity, but that the strength of this mechanism depends on the foreign policy target under consideration.

These findings contribute to the literature in three ways. First, they extend research on elite cue-taking by identifying political trust as an important factor shaping the relationship between elite preferences and mass foreign policy opinion. Second, they contribute to studies of public diplomacy by showing that foreign attitudes are not only shaped by external efforts at influence, but also filtered through domestic political institutions and the degree of confidence citizens place in them. Third, they add to research on Southeast Asia's role in great-power competition by showing that domestic institutional legitimacy helps structure how publics interpret the competing international positions of the United States and China.

At a broader level, the findings underscore that foreign policy opinion is not simply a product of international events or material interests. It is also mediated by the domestic political relationship between citizens and the state. If trust in government affects whether citizens adopt or resist elite foreign policy signals, then the domestic legitimacy of political institutions may be an important but underappreciated factor in shaping how publics respond to geopolitical competition. In this sense, understanding public attitudes toward major powers requires attention not only to what external actors do, but also to the domestic institutional environments through which those actions are interpreted.

The limitations of this study also point to several directions for future research. Newer survey data would help determine whether these relationships have shifted in response to more recent geopolitical developments, including changes in U.S. trade and diplomatic policy. Future work could also move beyond aggregate elite measures to examine the specific channels through which elite foreign policy cues reach mass publics, such as speeches, media coverage, and public diplomacy campaigns. More broadly, additional research across a larger set of countries and over time would help clarify when and under what conditions political trust most strongly amplifies mass–elite alignment in foreign policy attitudes.

REFERENCES

Aldrich, John H., Christopher Gelpi, Peter Feaver, Jason Reifler and Kristin Thompson Sharp. 2006. "Foreign Policy and the Electoral Connection." Annual Review of Political Science 9(1):477–502. 1

Asian Barometer Project. 2021. "Asian Barometer Project, Wave 5 Survey Data, 2018– 2021.". Co-directed by Professor Yun-han Chu and supported by the National Science and Technology Council, Academia Sinica, and National Taiwan University. URL: https://asianbarometer.org/datar?page=d10 2, 3

Baum, Matthew A. and Philip B.K. Potter. 2008. "The Relationships Between Mass Media, Public Opinion, and Foreign Policy: Toward a Theoretical Synthesis." Annual Review of Political Science 11(1):39–65. 14

Berinsky, Adam J. 2009. In Time of War. University of Chicago Press. 2, 14 Chapple-Sokol, Sam. 2013. "Culinary Diplomacy: Breaking Bread to Win Hearts and Minds." The Hague Journal of Diplomacy 8(2):161–183. 1

Cull, Nicholas J. 2008. "Public Diplomacy: Taxonomies and Histories." The Annals of the American Academy of Political and Social Science 616:31–54. 1, 13

Easton, David. 1975. "A Re-Assessment of the Concept of Political Support." British Journal of Political Science 5(4):435–457. 1, 4

Goh, Evelyn. 2007. "Great Powers and Hierarchical Order in Southeast Asia: Analyzing Regional Security Strategies." International Security 32(3):113–157. 13

Goh, Evelyn. 2016. "Southeast Asian Strategies toward the Great Powers: Still Hedging after All These Years?" The Asan Forum 4(1). 2

Goldsmith, Benjamin E. and Yusaku Horiuchi. 2012. "In Search of Soft Power: Does Foreign Public Opinion Matter for US Foreign Policy?" World Politics 64(3):555–585. 1

Goldsmith, Benjamin E., Yusaku Horiuchi and Kelly Matush. 2021. "Does Public Diplomacy Sway Foreign Public Opinion? Identifying the Effect of High-Level Visits." American Political Science Review 115(4):1342–1357. 1, 13

Guisinger, Alexandra and Elizabeth N. Saunders. 2017. "Mapping the Boundaries of Elite Cues: How Elites Shape Mass Opinion across International Issues." International Studies Quarterly 61(2):425–441. 2, 13

Hartig, Falk. 2016. "How China Understands Public Diplomacy: The Importance of National Image for National Interests." International Studies Review 18(4):655–680. 1

Hetherington, Marc J. 2005.Declining Political Trust and the Demise of American Liberalism.Princeton University Press.1, 4

Kertzer, Joshua D. 2017. "Resolve, Time, and Risk." International Organization 71(S1):S109–S136. 2

Kertzer, Joshua D. and Thomas Zeitzoff. 2017. "A bottomup theory of public opinion about foreign policy." American Journal of Political Science 61(3):543–558. 1, 14

Kuik, Cheng-Chwee. 2016. "How Do Weaker States Hedge? Unpacking ASEAN states' alignment behavior towards China." Journal of Contemporary China 25(100):500–514. 13

Levi, Margaret and Laura Stoker. 2000. "Political Trust and Trustworthiness." Annual Review of Political Science 3(1):475–507. 1, 4

Melissen, Jan. 2005. The New Public Diplomacy. London: Palgrave Macmillan UK. 1

Milner, Helen V. and Dustin H. Tingley. 2011. "Who supports global economic engagement? The sources of preferences in American foreign economic policy." International Organization 65(1):37–68. 1

Nye, Joseph S. 2004. "Soft Power and American Foreign Policy." Political Science Quarterly 119(2):255–270. 1, 13

Nye, Joseph S. 2008. "Public Diplomacy and Soft Power." The Annals of the American Academy of Political and Social Science 616:94–109. 1

Nye, Joseph S. and Youna Kim. 2019. Soft power and the Korean Wave. Routledge. 1

Peterson, Peter G. 2002. “Public Diplomacy and the War on Terrorism.” https://www.foreignaffairs.com/united-states/public-diplomacy-and-war-terrorism. 1

Saunders, Elizabeth N. 2022. “Elites in the making and breaking of foreign policy.” Annual Review of Political Science 25(1):219–240. 1

Seah, Sharon, Hoang Thi Ha, Melinda Martinus and Pham Thi Phuong Thao. 2021. The State of Southeast Asia: 2021 Survey Report. Technical report ASEAN Studies Centre, ISEAS-Yusof Ishak Institute. 2, 5

Seah, Sharon, Joanne Lin, Melinda Martinus, Kristina Fong, Pham Thi Phuong Thao and Indira Zahra Aridati. 2025. State of Southeast Asia: 2025 Survey Report. Technical report ASEAN Studies Centre, ISEAS - Yusof Ishak Institute. 16

Shafie, Tara. 2025. “K-pop Idols as Diplomats: South Korean Celebrities and Soft Power.” Athens Journal of Politics amp; International Affair 1(3):179–202. 1

Zaller, John and Dennis Chiu. 1996. “Government’s little helper: U.S. press coverage of foreign policy crises, 1945–1991.” Political Communication 13(4):385–405. 2

Zaller, John R. 1992. The Nature and Origins of Mass Opinion. Cambridge University Press. 14

Strengthening Social Protection in Guyana

Michael Khan

BRIEF DESCRIPTION OF THE ECONOMIC CONTEXT OF GUYANA

For much of recent history, Guyana has been a poor country. The country's economy has been characterized by high levels of poverty as well as reliance on the agricultural sector. It has been one of the worst-performing countries economically in the Western Hemisphere. All of this changed, however, when offshore oil was discovered in 2015. This discovery led to a rapid transformation of Guyana's economy, making it one of the fastest-growing economies in the world. Guyana is now ranked #1 in the world for oil production per capita, and currently produces ~900,000 barrels per day. Due to their newly found natural resource endowments, Guyana has had an average real GDP growth rate of 47% since 2022.[1] GDP per capita was just 6.95 thousand USD in 2020, and as of 2025 it had climbed to 31.38 thousand USD. While their growth is mainly due to oil revenues, Guyana has used these revenues to invest in infrastructure and public services, supporting the growth of the economy. Non-oil sectors are also experiencing growth, namely in agriculture, mining, and construction. Additionally, Guyana established a Natural Resource Fund in 2019, intending to use oil rents to support sustainable development for the future.

Inflation rose rapidly in Guyana in 2021, in part due to the pandemic but also as a result of its economic growth that rose faster than the provision of goods. Levels stayed high until 2023 and have cooled down to moderate levels since. While inflation has cooled, food prices have continued to increase by more than 75% on average since 2021.[2] Even amidst rapid economic growth, cost-of-living concerns remain significant for many Guyanese.

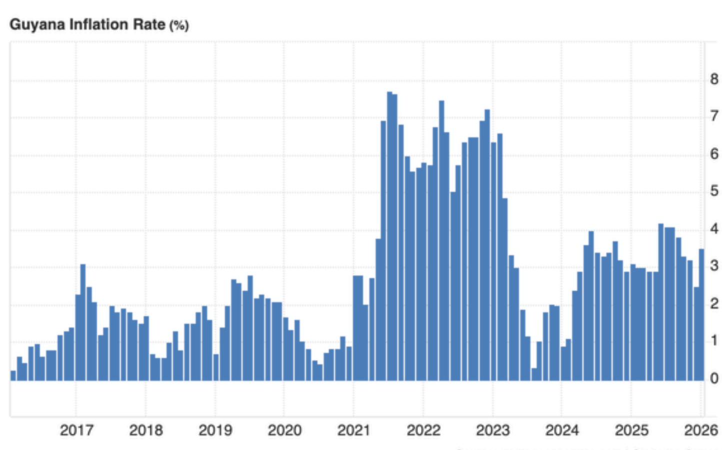

Source: Trading Economics

DESCRIPTION OF THE POVERTY AND INEQUALITY SITUATION IN GUYANA

While Guyana has experienced incredible economic growth over the past few years, inequality remains a problem in the country. Inequality in Guyana is more class-based, but the top decile shows an over-representation of Indo and Indigenous-Guyanese, while Afro and Mixed-Guyanese are over-represented in the bottom 90%.[3] Even though oil revenues have increased metrics like GDP per capita, there haven't been substantial decreases in poverty, which implies that benefits are not reaching marginalized populations.

Guyana maintains high levels of poverty, even as the country has experienced strong economic growth. An IDB report from 2024 ranked Guyana 3rd in poverty levels in South America, only behind Venezuela and Honduras.[4] More than half the population (58%) is considered to be poor, with 32% of the population considered to be in extreme poverty. Poverty in Guyana is characterized by higher levels of rural poverty than in most other countries in South America, with over 80% of the poor residing in rural areas.[5] This reflects a systemic inequity for poor households in rural areas and demonstrates the need for advancements in the country's social protection system to reach these households more effectively.

The oil boom in Guyana has increased average income, but high levels of inequality and poverty remain. Oil revenues must be used for the enrichment of all, and the government must take this chance to lift more Guyanese out of poverty through effective social protection.

ANALYSIS OF GUYANA'S SOCIAL PROTECTION SYSTEM, W/ EMPHASIS ON THE AREA OF PROPOSAL

Guyana's Social Protection programs are primarily run by the Ministry of Human Services and Social Security, and it is important to recognize that the social protection system has received much investment and improved greatly due to the country's oil revenues, as well as a 350 million USD loan from the Inter-American Development Bank. The loan from the IDB has the goal of supporting changes to the social protection system that were implemented in 2023, seeking to improve the efficiency and reach of the country's social safety net. Main objectives include the digitization of social services, expansion of public assistance for disabled people, job training, and women's empowerment.[6]

Substantial oil revenues have allowed for the vertical expansion of the current social safety net. The 2026 budget increased the monthly allowance of public assistance, old-age pensions, and various direct cash transfers, many of which are focused on school-aged children. School-aged children will now be eligible for 85,000 GYD, which is a combination of a cash grant, uniform voucher, and transportation grant.7 Additionally, universal programs like the National Cash Grant give 100,000 GYD to every Guyanese citizen 18 years or older. These programs demonstrate a recognition of the importance of supporting young people in an effort to improve education and labor outcomes in Guyana.

An important focus of recent funding and policy has been on the improvement of the pension system. In Latin America and the Caribbean, the region still lacks a wide-reaching social safety net for senior citizens. As of 2022, 34.5% of people over 65 years old in the region did not have an income, meaning they were not being covered by either labor-related income or pensions.[8] This is a serious coverage gap, which was laid bare by the pandemic in 2020 and 2021. After COVID demonstrated the fragility of Guyana's social safety net, substantial increases in the universal old-age pension followed. The monthly allowance of the pension has gone up to 46,000 GYD in 2026 from 20,500 GYD in 2020. A 20,000 GYD transportation grant was also added for old-age pensioners. While Guyana's non-contributory benefit system has seen substantial progress, the same can not be said for the contributory system, the National Insurance Scheme.

Guyana's National Insurance Scheme was first established in 1969. It is currently controlled by a Board of Directors and is overseen by the Ministry of Finance. The goal of the scheme, since its inception, has been to provide social security for all Guyanese workers. Unfortunately, its recent history has been marked by significant difficulties. While there had been problems of decline in contributors since the 1980's, the system came under noticeable strain in the 2010s. The Eighth Actuarial Review of the NIS stated plainly, "The National Insurance Scheme is nearing a crisis stage… the entire Fund will be exhausted in less than 10 years if contribution rate increases and benefit reforms are not made immediately."[9] The scheme had been put under immense strain from a combination of downward-trending contributors, non-compliance from employers and the self-employed, and underfunding.[10] Additionally, poor investments using funds from the scheme led to more financial difficulties. Questionable investments into companies like Berbice Bridge (domestic construction) as well as CLICO (an insurance firm from Trinidad & Tobago) led to criticisms about the misuse of funds that Guyanese workers were contributing to the scheme. The financial health of the NIS became so dire that Guyana sought help from the IDB, but many of the structural reforms suggested were never implemented.

To remedy some of these issues and continue the scheme, President Irfaan Ali gave a 10 billion GYD injection into the scheme to keep it afloat in 2024. The minimum pension amount was also increased to 43,000 GYD from 35,000 in 2024, which President Ali saw as necessary after visiting elderly communities. Another effort to increase the reach and effectiveness of the program was the One-Off Payment Program. This program gave some financial assistance to those over 60 who have not yet reached the 750 contribution threshold for a full pension. Those who were between the 500-749 contribution threshold were still able to access some financial assistance as a result of the program. The One-Off Payment Program has allowed more Guyanese to access social assistance and has benefited tens of thousands of Guyanese senior citizens.

DIAGNOSIS OF A PENDING NEED/CHALLENGE TO BE ADDRESSED

A major challenge to achieving effective social protection in Guyana is high levels of informal work and gig work. This is a common problem in developing countries around the world, and it is no different in Guyana. Prime Minister Mark Phillips noted that around half (48.3%) of Guyana's workforce is in the informal sector.[11] With such a large informal sector, many workers fall through the cracks of the social safety net. The current amount of the social pension is difficult to live on, and should be considered a supplement as opposed to the sole form of income for a senior citizen.

Another barrier to reaching the informal workforce is that the majority of Guyanese live in rural areas. More than 70% of the population lives in rural areas, and as mentioned above, more than 80% of Guyana's poor live in rural areas. Reaching the informal sector will require the ability to better target rural areas. A significant barrier to targeting in rural areas has been Guyana's geography. A majority of the country is characterized by dense hinterland forest, with many areas lacking sufficient infrastructure. Recent data finds that almost 100,000 people are off-grid, making it difficult to reach these communities through traditional means.[12] Many of these households have limited knowledge and/or access to social services, as opposed to urban areas. The challenge remains that it is not as easy for informal workers to be a part of the National Insurance Scheme. Informal workers in Guyana are classified as self-employed and have to pay 12.5% of their declared income towards the scheme. As informal workers' income is not as stable as that in the formal sector, that 12.5% of their declared income can become harder to contribute to an individual during periods of economic tumult. Yet, gig work continues to expand, and it becomes increasingly vital to ensure that these workers who cannot secure employment in the formal sector are incentivized to become contributors in the NIS.

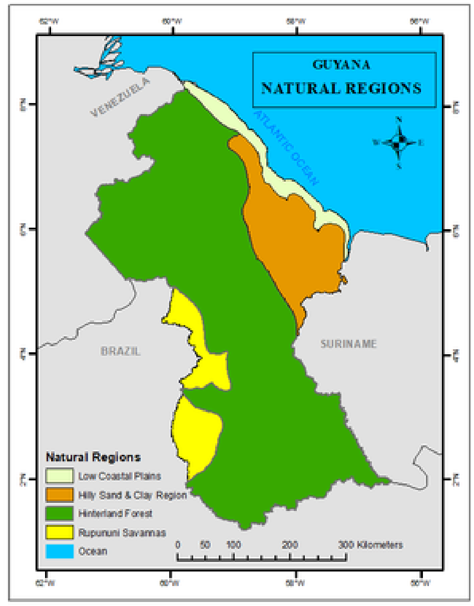

Source: Guyana Lands and Surveys Commission

PROPOSAL FOR ACTION

Looking at the whole of Guyana's social protection system, I believe that the most benefit can be realized through improving the contributory pension and social pension schemes, with a focus on the National Insurance Scheme. The social pension is important to supplement income in old age, but a better-organized contributory scheme has the opportunity to improve the welfare of Guyanese for generations to come. The following suggestions will make significant changes to the National Insurance Scheme and smaller changes to the social pension. It is necessary that we take the opportunity given by the country's oil revenues to enact reforms that will serve the Guyanese people for generations to come, and repair a structurally weak pension system.

DB/DC Hybrid Plan

Firstly, the National Insurance Scheme should become more flexible to accommodate informal work. Informal workers can not fully rely on the Old-Age Pension, and as such a large portion of Guyana's economy is characterized by informal work, contributory schemes need to reflect this reality. The NIS is currently a Defined Benefit plan, which is sensible for the formal sector, but I propose including a Defined Contribution option for the informal and gig worker sector. A Defined Contribution option will give informal workers the flexibility they need to pay what they can, and they are not locked into an inflexible pay scheme. While DB plans do offer predictability, a DC option is necessary for increasing uptake in the scheme. Increasing participation in the system is a virtuous cycle, as higher participation in the scheme will lead to lower costs and encourage further participation. Therefore, a hybrid DB/DC plan will allow for more contributions from self-employed/informal workers, which has been a structural issue in the NIS since the 1980's.

Improve Targeting and Outreach Methods for Rural Populations

The main barrier to reaching rural communities with social protection is administrative capacity. This leads to the missing middle, seen in the figure, where social protection systems can reach some of the poor and those in the formal sector, but miss out on a large portion of the informal sector, especially in rural areas. In Guyana's context, it is administratively costly to reach these communities, especially for a contributory scheme like the NIS. Much of the informal work in the rural areas consists of farming, logging, and mining, many of which have minimal regulation and access to contributory insurance schemes. I propose investing in the improvement of the administrative capacity of the NIS and the Old-Age Pension, which is often a limiting factor to reaching rural communities. Therefore, the Ministry of Finance and the Ministry of Human Services and Social Security should work with informal operations in the mining and logging sector to increase uptake of the NIS and the Old-Age Pension. With the option of flexible payments through the Defined Contribution option, it should yield an increase in contributors, even in the informal sector in rural areas.

Figure 2: **Gaps in coverage especially in the informal and rural economy**

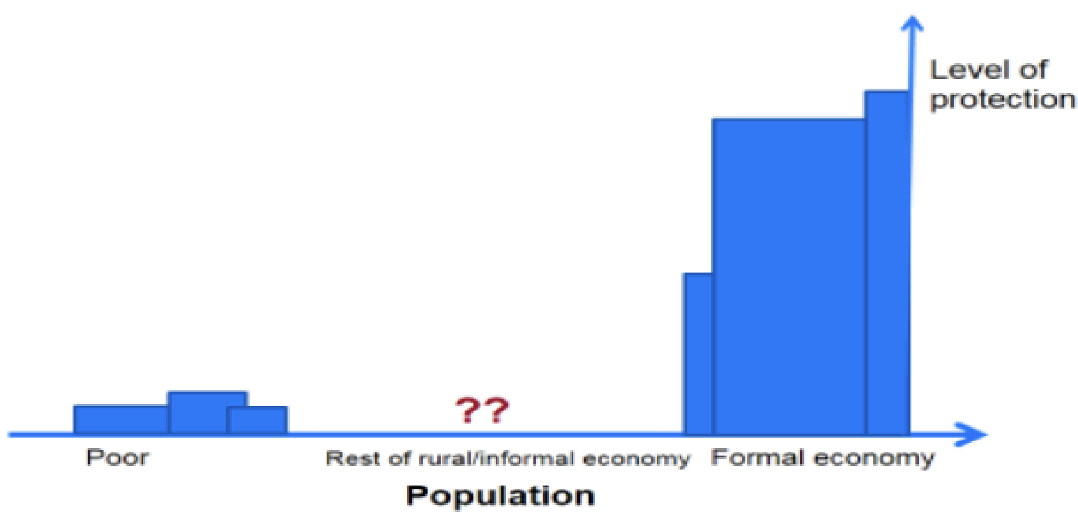

Source: United Nations Department of Economic and Social Affairs

NIS-Specific Fund from Oil Revenues

Relying on discretionary increases and funding injections, like the 10 billion GYD injection in 2024, does not ensure long-term sustainability of Guyana's pension scheme. I propose the establishment of a fund that is solely dedicated to funding the NIS, which can be initially funded by the country's Natural Resource Fund. As of 2025, Guyana's Natural Resource Fund has reached 3.6 billion USD, and oil production continues to ramp up and will continue for the foreseeable future.[13] Now is the time to ensure that these funds are being used for progressive development in Guyana. There is currently more than enough money in the Natural Resource Fund to create a fund specifically for the purpose of sustaining NIS funding and ensuring that the scheme is properly funded, as opposed to engaging in last-minute ad hoc funding. Enforcing the perception that the NIS is stable is likely to increase trust amongst Guyanese workers and eventually increase contributions.

Automatic Index to Inflation

Additionally, both the NIS and the non-contributory pension scheme must be automatically indexed to inflation. A common problem of social protection schemes, even when they are initially effective, is that inflation erodes the purchasing power of the money that individuals receive. Once inflation outstrips the increase in the allowance, even a program with excellent targeting and coverage becomes much less effective. The importance of avoiding this outcome can not be understated, especially in the context of Guyana. Rapid oil sector growth has led to the money supply increasing faster than the provision of goods and housing, leading to the average Guyanese household spending more for the same. This has a larger effect on individuals who receive a fixed income through pensions. As we have seen, the government has continued to increase the monthly allowance of both the NIS and the universal pension, but these discretionary increases do not always keep pace with the cost of living. Annually indexing pension payouts to inflation will ensure that the monthly allowance is sufficient to support the needs of those who receive pensions.

Changing of Age Requirements for Universal Pension and NIS

Access to pension payouts should reflect the realities of retirement and life expectancy in Guyana. As mentioned above, the standard retirement age for Public Service workers is 55 years old. Additionally, the average male life expectancy is 68 years old. The NIS is payable at 60, but for many, especially those who worked in physically demanding jobs, that could mean waiting for years before they can access their benefits. I propose a more flexible retirement window, with those who need to access benefits earlier receiving less monthly than those who can work until 65. Creating a system that enhances benefits the closer you work to 65, incentivizes those who can work longer to do so.

Creation of Investment Management Board

It is vital that the NIS learns from the past of poor investment and establishes an investment management board. Countries like Norway and Canada are models for sustainable investment, and it would be prudent for Guyana to learn from their success. Norway's IM Board, the Norges Bank Investment Management, invests its oil revenues on a global scale, distancing the success of the fund from purely domestic economic performance. A future NIS IM Board should seek to follow suit and not invest in domestic companies like Berbice Bridge, but seek to diversify investments globally to ensure long-term growth.

CONCLUSION

These changes demonstrate a recognition of the weaknesses of the NIS and the universal pension system. While funding for both of these programs has increased, it is not enough to simply throw money at the problem. At a time when the Guyanese government is flush with cash, we should take this time to implement structural reforms and create a social protection system that allows all Guyanese to contribute to a system that will benefit them in their old age and allow them to live with dignity and respect. Your respective ministries have the chance to transform a low-coverage, inefficient pension system into a sustainable, inclusive system that will remain a bedrock of the Guyanese social contract for decades to come.

SOURCES

A Forbes Burnham Legacy Under Pressure: The 56-Year History of Guyana's National Insurance Scheme, Guyana Business Journal, October 27, 2025. https://guyanabusinessjournal.com/2025/10/a-forbes-burnham-legacy-under-pressure-the-56-year-history-of-guyanas-national-insurance-scheme/ . 9.

Chabrol, Denis, Basic food items up 75% in past four years, Demerara Waves, July 4, 2025. https://demerarawaves.com/2025/07/04/basic-food-items-up-75-in-past-four-years/. 2.

Chang, Jillie et al., Ten Findings about Poverty in Latin America and the Caribbean, IDB, November, 2024. 4.

Constantine, Collin, Income inequality in Guyana: Class or ethnicity? New evidence from survey data, World Development, October 10, 2023. https://www.sciencedirect.com/science/article/pii/S0305750X23002474#:~:text=The%20main%20findings%20are%20as,Indigenous%2DGuyanese%20(1.3). 3.

IDB Approves US$350 Million Loan to Guyana to Strengthen Social Safety Net and Empower Vulnerable Populations, IDB, June 3, 2025. https://www.iadb.org/en/news/idb-approves-us350-million-loan-guyana-strengthen-social-safety-net-and-empower-vulnerable#:~:text=Enhanced%20Senior%20Citizens'%20Pension%20Program,and%20underserved%20populations%20in%20Guyana. 6. 10.

IMF Executive Board Concludes 2025 Article IV Consultation with Guyana, International Monetary Fund, May 7, 2025. https://www.imf.org/en/news/articles/2025/05/07/pr-25132-guyana-imf-executive-board-concludes-2025-article-iv-consultation. 1.

Natural Resource Fund, International Forum of Sovereign Wealth Funds. https://www.ifswf.org/members/natural-resource-fund. 13.

Nearly half of Guyana's workforce engaged in informal work, Stabroek News, April 11, 2025. https://www.stabroeknews.com/2025/04/11/business/nearly-half-of-guyanas-workforce-engaged-in-informal-work-pm/. 11.

Unconnected Communities, Guyana Energy Agency. https://gea.gov.gy/renewable/unconnected-communities/#:~:text=Guyana%20has%20approximately%20218%20Hinterland,times%20the%20cost%20in%20Georgetown. 12.

World Bank Group: Systematic Country Diagnostic: A Pivotal Moment for Guyana - Realizing The Opportunities, World Bank Group, November 1, 2020. 5.

34.5% of persons over 65 years do not have an income - ILO says in new report, Kaieteur News, December 23, 2022. https://kaieteurnewsonline.com/2022/12/23/34-5-of-persons-over-65-years-do-not-have-an-income-ilo-says-in-new-report/#:~:text=from%20Guyana%20%E2%80%93%20Dr.-,34.5%25%20of%20persons%20over%2065%20years%20do%20not%20have%20an,crisis%20triggered%20by%20the%20pandemic. 8.

$85,000 per child in 2025 as Transportation Grant is introduced, Ministry of Education, Guyana, January 27, 2026. https://education.gov.gy/en/index.php/media2/news-events/10132-85-000-per-child-in-2026-as-transportation-grant-is-introduced. 7.

Artificial Intelligence Acknowledgement: I acknowledge the use of ChatGPT for the purpose of; reviewing the structure of my paper, and researching other countries pension schemes and how it might relate to Guyana's context.

From Commitments to Changes: The Impact of Marine Commitments under the Paris Agreement on Ocean Health

Kalei Elizabeth Tano

ABSTRACT
Climate change has significant impacts on ocean health, from declines in biodiversity to increases in ocean acidification. Unlike many climate policies that fail to address the effects of climate change on the ocean, international governments explicitly included marine-related commitments in their pledges made under the Paris Agreement in 2016. Previous literature has not used the Paris Agreement as a policy shock to causally measure changes in ocean health. Motivated by this literature gap, this capstone project analyzes whether marine-focused commitments led to measurable improvements in ocean health following 2016 using a quasi-experimental difference-in-differences design. The analysis finds that, contrary to the hypothesis, marine-focused commitments did not translate into significant improvements in ocean health after the Paris Agreement, suggesting that commitments alone are insufficient to drive gains in ocean health and highlighting the need for a stronger link between climate and ocean policies.

1. INTRODUCTION

Climate change significantly impacts the ocean, leading to rising water temperatures and biodiversity loss. While international climate change policies such as the Kyoto Protocol or the Paris Agreement have addressed how countries plan to limit their carbon emissions over time, they often fail to comprehensively address how to mitigate the effects of climate change on the ocean. However, the Paris Agreement explicitly included ocean-related outcomes in its negotiations (Gallo et al., 2017). Whether policies such as the Paris Agreement can reduce these negative impacts through marine-focused commitments motivates the present capstone project.

This analysis uses data from "Ocean Commitments under the Paris Agreement" by Gallo et al. (2017) and the Ocean Health Index (OHI). Gallo et al. (2017) developed a metric to assess the breadth and intensity of marine-related commitments that governments submitted under the Paris Agreement. The OHI assesses global ocean health outcomes across 220 countries and 10 goals. The OHI Index (a score showing the average across all goal scores), Coastal Protection, Clean Waters, and Carbon Storage goals are used as outcome variables in the analysis. This project uses a two-way fixed effects difference-in-differences (TWFE DiD) design to estimate the causal relationship between marine-related commitments and ocean health outcomes after the Paris Agreement. Additionally, the analysis conducts an event study to show the treatment effects on each outcome by year, an alternative treatment specification, and a placebo test to measure the robustness of the results.

Despite the connection between climate change and ocean health, few studies have analyzed the OHI using a quasi-experimental design in response to a climate-related policy shock. In light of this research gap, this capstone project analyzes whether marine commitments that governments made under the Paris Agreement in 2016 led to improvements in OHI goals over time. The analysis finds that, for both global and Small Island Developing States (SIDS) outcomes, marine-focused commitments did not lead to significant improvements in the OHI's Index, Carbon Storage, Clean Waters, or Coastal Protection goal scores.

The paper is organized as follows: Section 2 reviews the current literature on climate change and ocean health (2.1), the Paris Agreement (2.2), and the OHI (2.3). Section 2 concludes by outlining the motivation behind this capstone project (2.4). Section 3 describes the data used in the analysis, reviewing the predictor variable (3.1), the outcome variable (3.2), and the summary statistics for all variables of interest (3.3). The empirical methodology, specifically the research design (4.1), the treatment variable definition (4.2), the regression estimations (4.3), and the robustness checks (4.4), are discussed in Section 4. The results are examined in Section 5, explaining the results of the regression analysis for the effect of the MFF on global OHI goal scores (5.1), global Index scores with pre-treatment controls (5.2), SIDS OHI goal scores (5.3), SIDS Index scores with pre-treatment controls (5.4), and the event study results (5.5). Section 6 reviews the robustness checks: the alternative treatment specification (6.1) and the placebo test (6.2). Finally, Section 7 discusses the implications (7.1) and limitations of this analysis (7.2), as well as conclusions and directions for future research (7.3).

2. BACKGROUND AND LITERATURE REVIEW

2.1: Climate Change and Ocean Health

Rising global temperatures have significant impacts on ocean health. As explained by Gallo et al. (2017), the primary stressors from climate change on marine ecosystems are ocean warming, acidification, deoxygenation, and changes in fish productivity. Guggisberg (2019) and Bijma et al. (2013) indicate that atmospheric warming increases ocean temperatures, in turn leading to biodiversity loss, ecosystem disruptions, increased disease prevalence, and escalated coral bleaching. A scenario analysis comparing the current emissions trajectory with a stringent emissions trajectory by Gattuso et al. (2015) indicates that even under a stringent emissions scenario, ocean systems face a range of negative and cumulative impacts. This paper supports the claim that strict climate policies with strong goals and enforcement mechanisms aimed at reducing global warming can benefit ocean health. Many studies, such as the scenario analysis conducted by Gattuso et al. (2015) and the OHI assessment by Halpern et al. (2015), empirically analyze ocean health. However, few studies have sought to demonstrate a causal link between climate policies and ocean health. The present analysis seeks to address this gap by applying a causal design to the relationship between marine commitments and ocean health outcomes.

2.2: The Paris Agreement and Climate Policies

The Paris Agreement, adopted on December 12, 2015 by 195 countries, focuses on mitigating climate change impacts and explicitly recognizes the relationship between the ocean and climate regulation (Gallo et al., 2017). By June 2016, 161 governments had submitted "nationally determined contributions" (NDCs), or commitments that indicate their climate policy priorities. The European Union (EU) submitted one NDC in 2016, covering all 28 EU countries. 70% of the NDCs submitted in 2016 included marine issues. The dominant marine concerns included the impacts of climate change on coasts, ocean warming, and fisheries, revealing these governments' marine policy priorities. 21% of the governments that did not include marine issues in their NDC in 2016 are landlocked

(35 NDCs), while 9% are coastal (14 NDCs). Notably, some of the 14 coastal countries that did not include marine issues in their NDCs have large Exclusive Economic Zones (EEZs), including Australia, the United States of America, and New Zealand. EEZs are areas within and beyond the territorial sea where states have sovereign rights (United Nations, n.d.). Many of the world's poorest people rely on ocean-related factors that affect food, jobs, and revenue. These individuals tend to reside in countries that will be significantly affected by ocean changes resulting from climate change. Including ocean impacts in climate policy strategies is necessary to reduce the harm caused to vulnerable coastal communities, protect biodiversity and marine ecosystems, and support successful climate adaptation and mitigation strategies (Gallo et al., 2017). Analyzing the Paris Agreement as a policy shock and using a difference-in-differences (DiD) design with fixed effects, Ullah et al. (2025) find that the Paris Agreement has supported long-term climate action and resilience. A scenario analysis conducted by Sumaila et al. (2019) finds that implementing the Paris Agreement and achieving warming targets can benefit developing countries that rely on seafood for consumption and economic growth. In addition to the goals outlined by NDCs across governments, policies that reduce greenhouse gas emissions can support ocean health and improve a variety of ocean health dimensions, such as carbon storage, water quality, and coastal protection. These studies indicate that climate change policies, including the Paris Agreement, are beginning to consider ocean health in their policy strategies. By doing so, climate policies can reduce the impacts faced by marine ecosystems, coastal livelihoods, and ocean warming. These assumptions influence the objective of this analysis.

2.3: Overview of the Ocean Health Index

The OHI is a global measure of 220 countries' ocean health outcomes, encompassing 10 separate ocean-related goals that are assessed by an interdisciplinary group of scientists (Ocean Health Index, 2026). These goals include Food Provision, Artisanal Fishing Opportunities, Natural Products, Carbon Storage, Coastal Protection, Livelihoods and Economies, Tourism and Recreation, Sense of Place, Clean Waters, and Biodiversity (Ocean Health Index, 2026). As indicated by the OHI assessment conducted by Halpern et al. (2015), poorer countries with a history of conflict often have low OHI scores, likely due to their limited institutional capacity to address environmental pressures. This assessment reveals low covariance in goal scores across regions, suggesting that each goal score contains unique information about different subsets of ocean health (Halpern et al., 2015). The OHI has rarely been used to measure the impact of policy shocks on ocean health using causal frameworks, possibly because there are few global climate policies, aside from the Paris Agreement, that address ocean health.

2.4: Research Motivation

Given the effects of global warming on ocean health and the Paris Agreement's objectives to stabilize global temperatures, countries that have signed the Paris Agreement may demonstrate higher ocean health scores compared to those that have not. While the aforementioned studies explain the OHI, the Paris Agreement goals, and the effect of rising temperatures on ocean health, there remains a lack of a causal analysis of whether climate policies, such as the Paris Agreement, have led to improvements in ocean health. Previous studies focusing on the OHI have excluded policy shocks as predictor variables, while studies focusing on the Paris Agreement as a policy shock have excluded measures of ocean health. Therefore, this capstone project seeks to fill this research gap. The results can show how climate policies support ocean health or fail to do so, and advocate for stronger enforcement mechanisms that link climate and ocean policies.

3. DATA DESCRIPTION

3.1: Predictor Variable Description – Marine Focus Factor Scores

The predictor variable is the Marine Focus Factor (MFF) scores, constructed by Gallo et al. (2017) and obtained from their paper "Ocean Commitments under the Paris Agreement." The MFF is a metric that captures the intensity and breadth of marine issues that countries include in their NDCs submitted under the Paris Agreement. The MFF is constructed by aggregating the marine word count and marine categories from each NDC, excluding landlocked countries and countries that did not submit an NDC in 2016 (Gallo et al., 2017). The analysis excludes EU countries because the EU submitted a single NDC in 2016, which limits the available MFF data for EU countries.

Due to right-skewness and the presence of zeros in the raw MFF distribution, an inverse hyperbolic sine (IHS) transformation is applied to the raw MFF scores to normalize the distribution and retain countries with MFF scores of zero. Raw MFF scores range from 0 to 23, with higher scores indicating a stronger marine focus, while IHS(MFF) scores range from 0 to 3.84. MFF is time-invariant since countries submitted only one NDC in 2016. Several of the governments with the highest MFF scores come from members of the Small Island Developing States (SIDS), including St. Kitts and Nevis, Maldives, Kiribati, Seychelles, Mauritius, St. Vincent and the Grenadines, and Nauru (Gallo et al., 2017). A table listing the SIDS and their corresponding MFF and IHS (MFF) scores is provided in Table 7 of the Appendix. The analysis hypothesizes that higher MFF scores would prompt countries to revise their marine policies to improve ocean-related outcomes, leading to measurable improvements in OHI goal scores.

Figure 1: Global Distribution of Marine Focus Factor Scores

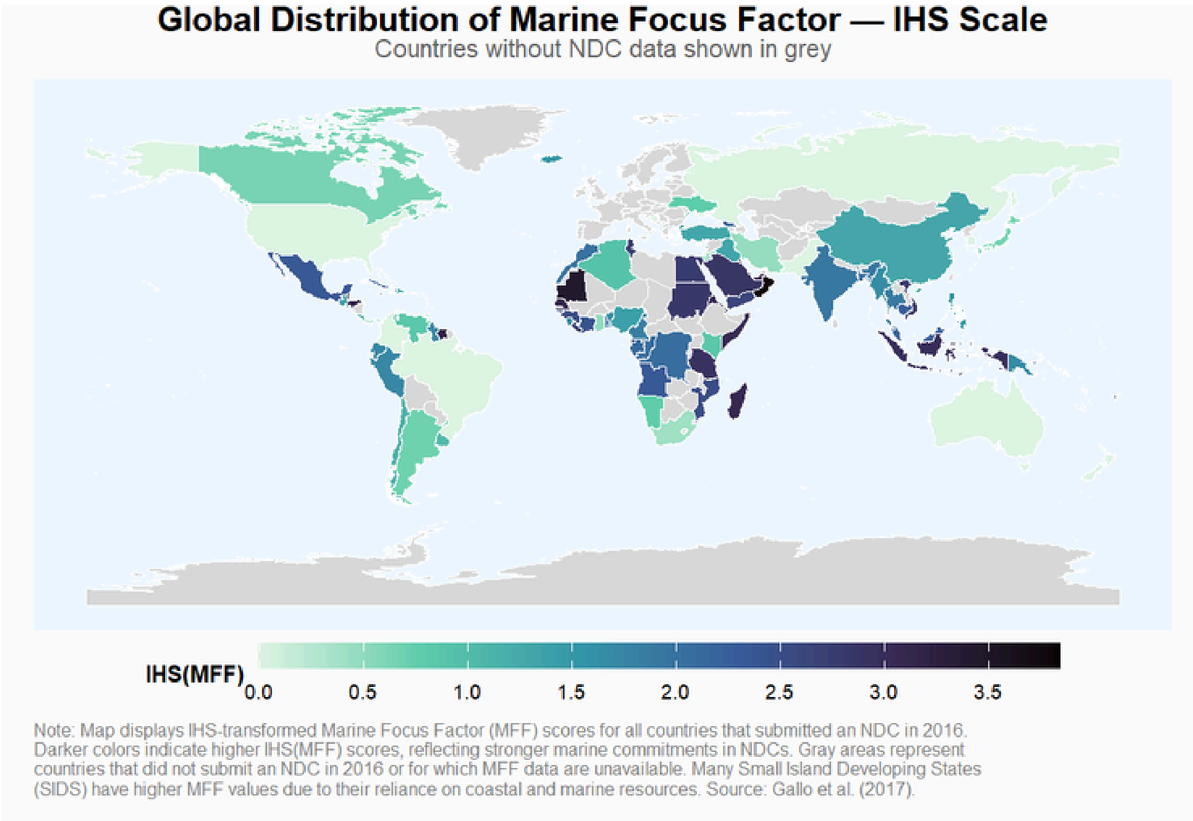

Figure Notes: *The map above shows the global distribution of MFF scores with the IHS transformation, with darker colors indicating higher IHS(MFF) scores. Gray areas represent nations that did not submit an NDC in 2016. Many SIDS have higher MFF values due to their higher dependency on seafood, coastal industries, and artisanal fishing. This figure uses data from Gallo et al. (2017).*

3.2: Outcome Variable Description – Ocean Health Index Goal Scores

The outcome variable is obtained from the Ocean Health Index, a framework measuring each country's ocean health across 10 goals, including but not limited to Coastal Protection, Clean Waters, and Carbon Storage. OHI ranges from 0 to 100, with higher values indicating better ocean health. This analysis uses the Index score – the average OHI score for each country across all 10 goals – as well as Coastal Protection, Clean Waters, and Carbon Storage, as outcome variables. These goals are selected for their potential connections to climate change mitigation strategies, increased coastal policies, and improved water quality and carbon storage.

Figure 2: Ocean Health Index Trends Over Time

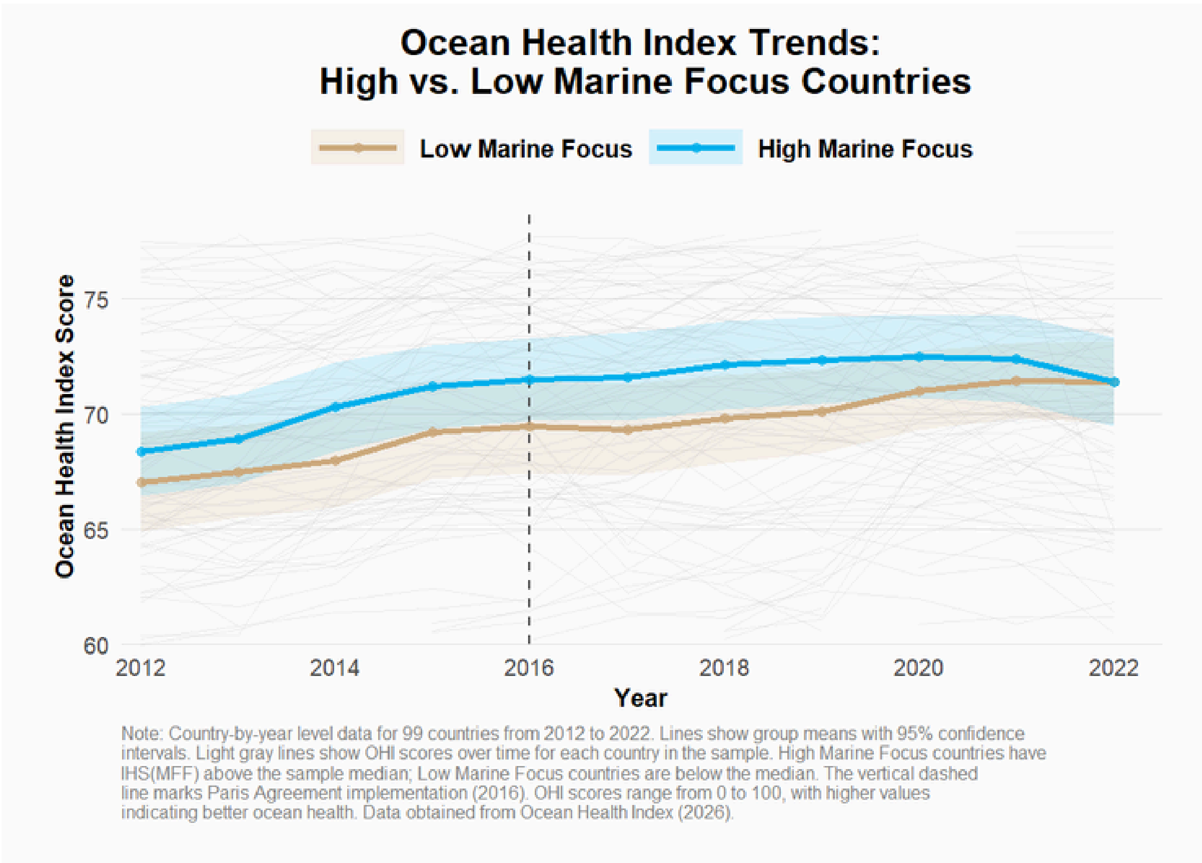

Figure Notes: *The above plot illustrates the average OHI Index scores between 2012 and 2022 for countries with high MFF scores (blue) compared to countries with low MFF scores (tan). Pre-treatment trends (2012–2015) between treated and control groups are similar, suggesting that the parallel trends assumption holds. Light gray lines show each country's OHI Index scores over time. The data for this figure are obtained from Gallo et al. (2017) and Ocean Health Index (2026).*

3.3: Summary Statistics

The following summary statistics show the minimum, maximum, mean, standard deviation, and median values for the main variables of interest for the years 2012–2022. The Marine Focus Factor (MFF) is transformed using the inverse hyperbolic sine (IHS) to reduce skewness in the MFF distribution. The OHI goal scores range from 0 to 100, but only the Carbon Storage and Coastal Protection reach scores of 100. While the total sample has 1,089 observations, 11 observations are excluded for these two goal scores due to missing data, yielding 1,078 observations for these two outcomes.

Table 1: Summary Statistics

Summary Statistics						
Variable	**N**	**Minimum**	**Maximum**	**Mean**	**SD**	**Median**
Marine Focus Factor	1089	0.00	23.21	5.37	5.05	3.62
IHS Marine Focus Factor	1089	0.00	3.84	1.92	1.06	2.00
Ocean Health Index	1089	45.55	87.67	70.29	6.94	70.78
Carbon Storage	1078	4.89	100.00	82.82	15.88	82.29
Coastal Protection	1078	12.24	100.00	77.27	17.30	78.97
Clean Waters	1089	14.29	97.80	65.82	17.85	68.19

Table Notes: *This table presents summary statistics (number of observations, minimum, maximum, mean, standard deviation, and median) for the variables included in this analysis. IHS(MFF) and OHI goal scores are the primary predictor and outcome variables, respectively.*

4. EMPIRICAL METHODOLOGY

This analysis tests whether stronger commitments to marine issues by governments under he Paris Agreement in 2016 led to increases in ocean health outcomes, specifically the OHI Index, Coastal Protection, Clean Waters, and Carbon Storage goal scores. The hypothesis is that higher marine commitments will lead to modest but meaningful increases in these goal scores following 2016.

4.1: Research Design

This analysis employs a quasi-experimental difference-in-differences (DiD) design with two-way fixed effects to estimate the causal impact of marine-focused commitments under the Paris Agreement in 2016 on ocean health. Using a DiD design, the analysis compares changes in OHI scores before and after the Paris Agreement across countries with varying levels of marine commitments, controlling for pre-existing differences between countries and global time trends that affect all countries simultaneously. The unit of analysis is country-year, with 99 countries observed annually from 2012 to 2022. The analysis provides four years of pre-treatment data and six years of post-treatment data. Although data on ocean health are available for 2023 to 2025, the COVID-19 pandemic significantly affected OHI scores after 2022, resulting in lower average scores.

The pandemic disrupted fisheries, coastal economies, and marine governance in ways that could confound post-treatment comparisons, independently of any Paris Agreement effects. Because of this shock, these years are excluded. Country fixed effects are included to control for time-invariant unobserved heterogeneity across countries, while year fixed effects account for global time shocks that affect all countries simultaneously.

4.2: Treatment Variable Definition

The treatment MFF could be defined as either binary or continuous. To determine which specification is more appropriate for this analysis, a cumulative distribution function (CDF) of MFF values was constructed. The CDF shows no discontinuities in the MFF values, indicating that the MFF is not sensitive to arbitrary classifications. Therefore, the treatment is defined as continuous MFF values. However, a binary treatment variable, where 1 = countries with high IHS(MFF) values and 0 = countries with low IHS(MFF) values, is used as a robustness check to assess the sensitivity of the results. Figure 7 in the Appendix shows CDF plots with cutoffs at the 25th, 50th, and 75th percentiles to justify the median treatment cutoff.

4.3: Regression Estimation

The analysis estimates a TWFE DiD regression for each of the four outcome variables – OHI Index, Coastal Protection, Clean Waters, and Carbon Storage – using the continuous IHS(MFF) as the treatment variable. The analysis then applies the same specification on a restricted SIDS subsample to assess whether these nations exhibit a stronger relationship between MFF and OHI scores. Countries within the SIDS subsample are listed in Table 7 of the Appendix. The analysis includes pre-treatment controls for gross domestic product (GDP), the fisheries' value as a share of GDP, and the total percentage of marine protected areas (MPAs). These controls are included because they may confound the relationship between MFF and OHI scores. Countries with higher pre-existing fisheries value or MPA coverage may be more inclined to include marine issues in their NDCs, and they may also have a higher OHI baseline score. Higher GDP may allow countries to expand their climate commitments to marine issues. By interacting these controls with the post-Paris indicator, the analysis isolates the effect of the MFF from the influence of pre-existing economic and marine conditions.

Main Regression Equation

$$OutcomeGoal = \alpha_1 + \beta_1(TreatmentMFF_c * postParis_t) + \beta_2(logGDP*postParis) + \beta_3(fisheries*postParis) + \beta_4(MPA*postParis) + \delta_c + \gamma_t + \epsilon_{ct}$$

Countries with higher MFF scores may have simultaneously implemented other pro-environmental policies following 2016, unrelated to their NDC commitments. If these policies correlate with both MFF and OHI scores, they could inflate the estimates. In addition, countries with stronger governance capacity are likely to have higher MFF and OHI scores than countries with lower governance, potentially confounding the results. Country fixed effects partially mitigate these concerns by controlling for time-invariant differences in governance and policy capacity across countries.

4.4: Robustness Checks

The analysis includes robustness checks to assess the sensitivity of the results. Firstly, the treatment will be adjusted using the binary IHS(MFF) scores rather than a continuous IHS(MFF) treatment to test whether the null findings are sensitive to the treatment definition. Secondly, a placebo test assigns 2013 as a falsified treatment year. If the parallel trends assumptions hold, the placebo interaction should be statistically insignificant across all outcomes. Significant results would indicate pre-existing differences between high and low MFF countries prior to the Paris Agreement. Thirdly, an event study estimates the year-by-year treatment effects using the TWFE interaction estimator to visually inspect pre-trends and reveal whether any treatment effects emerge gradually or immediately after 2016.

5. RESULTS

The regression results for this analysis reveal whether marine commitments within NDCs led to meaningful OHI outcomes, both globally and within SIDS. The hypothesis is that increases in the MFF will lead to higher OHI scores following the Paris Agreement. The main coefficients across all models are small and insignificant, indicating that MFF scores did not lead to substantial improvements in ocean health.

5.1: Effect of Marine Focus Factor on Global OHI Scores

The first regression analyzes the effect of the MFF on global OHI goal scores, using a continuous IHS(MFF) treatment. Overall, the regression results are insignificant; higher marine focus within NDCs did not lead to improvements in global OHI goal scores. Index and Clean Waters have the largest coefficients. A one-unit increase in MFF scores is associated with an average 0.1585-unit increase in Index scores following the Paris Agreement. Similarly, a one-unit increase in MFF is associated with a 0.4338-unit increase in Clean Waters scores following the Paris Agreement, on average. These coefficients have the widest confidence intervals and largest standard errors among the four main outcome coefficients, at 0.4114 and 0.2777, respectively. Coastal Protection shows a negative relationship with MFF, such that a one-unit increase in MFF is associated with a 0.0246-unit decrease in Coastal Protection scores following the Paris Agreement, on average.

The Carbon Storage coefficient indicates that a one-unit increase in MFF scores is associated, on average, with a 0.0222-unit increase in Carbon Storage scores following the Paris Agreement. The 95% confidence intervals for the main coefficient all cross zero, confirming the results are insignificant across all four outcomes. The R^2 shows high model fit, likely due to the inclusion of fixed effects that control for country- and year-level variation.

Table 2: Effect of Marine Focus Factor on Global OHI Scores

The Effect of Marine Focus Factor on Global OHI Goal Scores (Continuous Treatment)

Outcome Variable	(1) OHI Index	(2) Coastal Protection	(3) Clean Waters	(4) Carbon Storage
IHS(MFF) × Post-Paris	0.1585	−0.0246	0.4338	0.0222
	(0.4114)	(0.0570)	(0.2777)	(0.0480)
Pre-Treatment IHS(GDP) × Post	0.1826	−0.0248	0.1175	0.0168
	(0.1673)	(0.0515)	(0.1041)	(0.0172)
Pre-Treatment Fisheries × Post	−5.5619	−0.3222	−3.6893	0.0367
	(7.1070)	(1.2863)	(4.1335)	(1.0787)
Pre-Treatment % MPA × Post	−0.0516*	0.0020	0.0015	−0.0048
	(0.0287)	(0.0123)	(0.0215)	(0.0036)
Num. Obs.	1089	1078	1089	1078
R^2	0.868	0.998	0.991	0.999
Country FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Mean Dep. Var.	70.29	77.27	65.82	82.82
95% CI Lower	−0.6578	−0.1378	−0.1172	−0.0729
95% CI Upper	0.9748	0.0886	0.9848	0.1174

Notes: Country-by-year level regressions for 99 countries observed annually from 2012 to 2022. IHS(MFF) is the inverse hyperbolic sine transformation of the Marine Focus Factor, a continuous measure of marine commitments in each country's NDC submitted under the Paris Agreement in 2016. IHS(MFF) ranges from 0 to 3.84. Post-Paris equals 1 for years 2016–2022. Pre-treatment controls (IHS GDP, Fisheries Value as % of GDP, and % Marine Protected Areas) are pre-2016 country averages interacted with Post-Paris. OHI goal scores range from 0 to 100, with higher values indicating better ocean health. All models include country and year fixed effects. Standard errors clustered at the country level are shown in parentheses. 95% CI refers to the confidence interval for the IHS(MFF) × Post-Paris coefficient. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table Notes: *The above regression table presents results for the regression models estimating the impact of MFF scores following the Paris Agreement in 2016 on the OHI Index, Coastal Protection, Clean Waters, and Carbon Storage goal scores. The regression models indicate that higher MFF scores are not significantly associated with changes in global OHI goal scores.*

5.2: The Effect of Pre-Treatment Controls on Global OHI Index Scores

The second regression table shows how the sequential addition of the pre-treatment controls affects global OHI scores. Without controls, a one-unit increase in MFF scores is associated with an average increase of 0.0927-units in the OHI Index. The addition of pre-treatment GDP and fisheries value does not reveal significant results. The interaction between pre-treatment fisheries value and years following the Paris Agreement yields significant results – a one-unit increase in pre-treatment fisheries value is associated with an average decrease in Index scores of 11.1515-units, significant at the 90% confidence level. This indicates that governments with higher fisheries values before the Paris Agreement showed decreases in fisheries values after 2016, possibly due to stronger fishing regulations. The coefficient for pre-treatment MPAs following 2016 also shows a significant result, revealing that a one-unit increase in the percentage of MPAs before the Paris Agreement is associated with an average decrease in Index scores by 0.0499-units, significant at the 90% confidence level. The 95% confidence interval for the main coefficient crosses zero across all five model specifications, indicating that the MFF coefficient is not statistically distinguishable from zero, regardless of which controls are included in the models.

Table 3: Effect of Pre-Treatment Controls on Global OHI Scores

Continuous IHS(MFF) × Post-Paris — Global OHI Index with Pre-Treatment Controls

Outcome Variable	(1) Baseline	(2) +IHS GDP	(3) +Fisheries	(4) +MPA	(5) All Controls
IHS(MFF) × Post-Paris	0.0927	0.2942	0.1702	−0.0702	0.1585
	(0.3619)	(0.4165)	(0.3566)	(0.3624)	(0.4114)
Pre-Treatment IHS(GDP) × Post		0.1879			0.1826
		(0.1642)			(0.1673)
Pre-Treatment Fisheries × Post			−11.1515*		−5.5619
			(6.5903)		(7.1070)
Pre-Treatment % MPA × Post				−0.0499*	−0.0516*
				(0.0288)	(0.0287)
Num. Obs.	1089	1089	1089	1089	1089
R^2	0.865	0.866	0.866	0.867	0.868
Country FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Mean Dep. Var.	70.29	70.29	70.29	70.29	70.29
95% CI Lower	−0.6254	−0.5324	−0.5374	−0.7895	−0.6578
95% CI Upper	0.8108	1.1208	0.8778	0.6490	0.9748

Notes: Country-by-year level regressions for 99 countries observed annually from 2012 to 2022. Dependent variable is the OHI Index (0–100 scale). Controls are pre-2016 country averages of IHS(GDP), Fisheries Value as % of GDP, and % Marine Protected Areas, each interacted with Post-Paris. All models include country and year fixed effects. Standard errors clustered at the country level are shown in parentheses. 95% CI refers to the confidence interval for the IHS(MFF) × Post-Paris coefficient. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table Notes: *The above table shows the regression results for each control as a separate model. The interaction between pre-treatment percentage of MPAs and years following the Paris Agreement yields significant results retained in the final model, revealing a negative relationship with global OHI Index scores.*

5.3: Effect of Marine Focus Factor on SIDS OHI Scores

The third regression model shows the results for the relationship between MFF scores and OHI goal scores within SIDS. There are no statistically significant results in these models, but the coefficients for the interaction between MFF and years following the 2016 show greater magnitudes. Holding controls constant, a one-unit increase in MFF is associated with a 1.1889-unit increase in Index scores within SIDS, on average. Similarly, a one-unit increase in MFF is associated with a 0.9585-unit increase in Clean Waters scores within SIDS, on average. The weakest coefficients are for the Coastal Protection and Carbon Storage outcomes, as in the global model, and Coastal Protection retains its negative correlation with MFF scores. The pre-treatment fisheries value still has the largest coefficients for the Index and Clean Waters outcomes. Despite the larger coefficient magnitudes in the SIDS models, the 95% confidence intervals remain wide and cross zero across all four outcomes.

Table 4: Effect of Marine Focus Factor on SIDS OHI Scores

The Effect of Marine Focus Factor on SIDS OHI Goal Scores (Continuous Treatment)

Outcome Variable	(1) OHI Index	(2) Coastal Protection	(3) Clean Waters	(4) Carbon Storage
IHS(MFF) × Post-Paris	1.1889	−0.0909	0.9585	0.0555
	(0.7506)	(0.0919)	(0.8221)	(0.1043)
Pre-Treatment IHS(GDP) × Post	0.2498	−0.0313	0.2696	0.0060
	(0.3321)	(0.0359)	(0.2703)	(0.0419)
Pre-Treatment Fisheries × Post	9.8097	−1.5715	18.5411	−1.5960
	(30.2565)	(2.6363)	(22.8419)	(2.3287)
Pre-Treatment % MPA × Post	−0.0797	−0.0172	0.0409	−0.0182
	(0.0701)	(0.0122)	(0.0260)	(0.0117)
Num. Obs.	308	308	308	308
R^2	0.777	1.000	0.977	0.999
Country FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Mean Dep. Var.	73.10	72.92	75.87	89.50
95% CI Lower	−0.3512	−0.2795	−0.7283	−0.1585
95% CI Upper	2.7291	0.0976	2.6452	0.2695

Notes: Country-by-year level regressions for Small Island Developing States (SIDS) observed annually from 2012 to 2022. The sample is restricted to SIDS, which are coastal nations with high dependence on marine resources. IHS(MFF) is the inverse hyperbolic sine transformation of the Marine Focus Factor from Gallo et al. (2017). Pre-treatment controls (IHS GDP, Fisheries Value as % of GDP, and % Marine Protected Areas) are pre-2016 country averages interacted with Post-Paris. OHI goal scores range from 0 to 100. All models include country and year fixed effects. Standard errors clustered at the country level are shown in parentheses. 95% CI refers to the confidence interval for the IHS(MFF) × Post-Paris coefficient. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table Notes: *The above table shows the results for the regression models testing the effect of MFF scores on OHI goal scores within SIDS. There are no statistically significant results in the model. Similar to the global model, Index and Clean Waters goal scores show the largest coefficients.*

5.4: The Effect of Pre-Treatment Controls on SIDS OHI Scores

The fourth regression table shows the effect of the continuous MFF on SIDS OHI Index scores, with the sequential addition of pre-treatment controls. Without controls, a one-unit increase in the MFF is associated with a 1.1660-unit increase in Index scores in the years following the Paris Agreement, significant at the 90% confidence level. The coefficient remains significant at the 90% confidence level and increases in magnitude when pre-treatment controls for GDP and fisheries value are added. The coefficient loses significance and decreases in magnitude when a control for the pre-treatment percentage of MPAs is included. The final model, while insignificant, has the highest magnitude across the main coefficients, but also the highest standard error. The 95% confidence interval for the main MFF coefficient crosses zero across all five models, reflecting the increasing uncertainty.

Table 5: Effect of Pre-Treatment Controls on SIDS OHI Scores

Continuous IHS(MFF) × Post-Paris — SIDS OHI Index with Pre-Treatment Controls

Outcome Variable	(1) Baseline	(2) +IHS GDP	(3) +Fisheries	(4) +MPA	(5) All Controls
IHS(MFF) × Post-Paris	1.1660*	1.1756*	1.1780*	1.1157	1.1889
	(0.6654)	(0.6870)	(0.6829)	(0.7048)	(0.7506)
Pre-Treatment IHS(GDP) × Post		0.1151			0.2498
		(0.3500)			(0.3321)
Pre-Treatment Fisheries × Post			2.2252		9.8097
			(26.4045)		(30.2565)
Pre-Treatment % MPA × Post				−0.0789	−0.0797
				(0.0783)	(0.0701)
Num. Obs.	308	308	308	308	308
R^2	0.772	0.773	0.772	0.776	0.777
Country FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Mean Dep. Var.	73.10	73.10	73.10	73.10	73.10
95% CI Lower	−0.1993	−0.2340	−0.2232	−0.3304	−0.3512
95% CI Upper	2.5313	2.5852	2.5792	2.5619	2.7291

Notes: Country-by-year level regressions for Small Island Developing States (SIDS) observed annually from 2012 to 2022. Dependent variable is the OHI Index (0–100 scale). Pre-treatment controls are pre-2016 country averages of IHS(GDP), Fisheries Value as % of GDP, and % Marine Protected Areas, each interacted with Post-Paris. All models include country and year fixed effects. Standard errors clustered at the country level are shown in parentheses. 95% CI refers to the confidence interval for the IHS(MFF) × Post-Paris coefficient. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table Notes: *The table above shows regression results for each control variable in the SIDS subset. The MFF coefficient increases in magnitude and retains significance between the baseline and the addition of GDP and fisheries, but loses significance and magnitude when adding the pre-treatment percentage of MPAs.*

5.5: Event Study Results

The analysis uses four event studies to show the pre-trends and estimated treatment effects of the MFF on each OHI goal. The event studies use a TWFE interaction estimator that estimates year-by-year interactions between time-to-treatment dummies and the demeaned continuous IHS(MFF) scores, with 2015 as the omitted reference period. The event study includes country and year fixed effects and clusters standard errors at the country level. Across all four outcomes, the pre-trends (years 2012 to 2015) are close to 0 and insignificant, supporting the parallel trends assumption. Following the treatment in 2016, the estimates are small and often close to zero. The confidence intervals across all models cross zero, indicating that there are no statistically meaningful results in the event studies. This supports the null hypothesis that MFF scores have no significant impact on OHI outcomes.

The Index event study shows slightly positive coefficients following 2016, but begins to fall after 2019. The confidence intervals gradually widen, showing increasing uncertainty in later years. The Coastal Protection event study shows a similar pattern, with coefficients gradually increasing before falling after 2019 and widening confidence intervals. The Clean Waters event study shows the largest variability in pre-trends and a significant coefficient for the estimated effect in 2012, indicating that the parallel trends assumption may not hold for this OHI goal. However, the post-treatment coefficients are also centered around zero. Finally, the Carbon Storage event study shows the smallest effect sizes across all models and retains insignificance.

Figure 3: Event Study Results

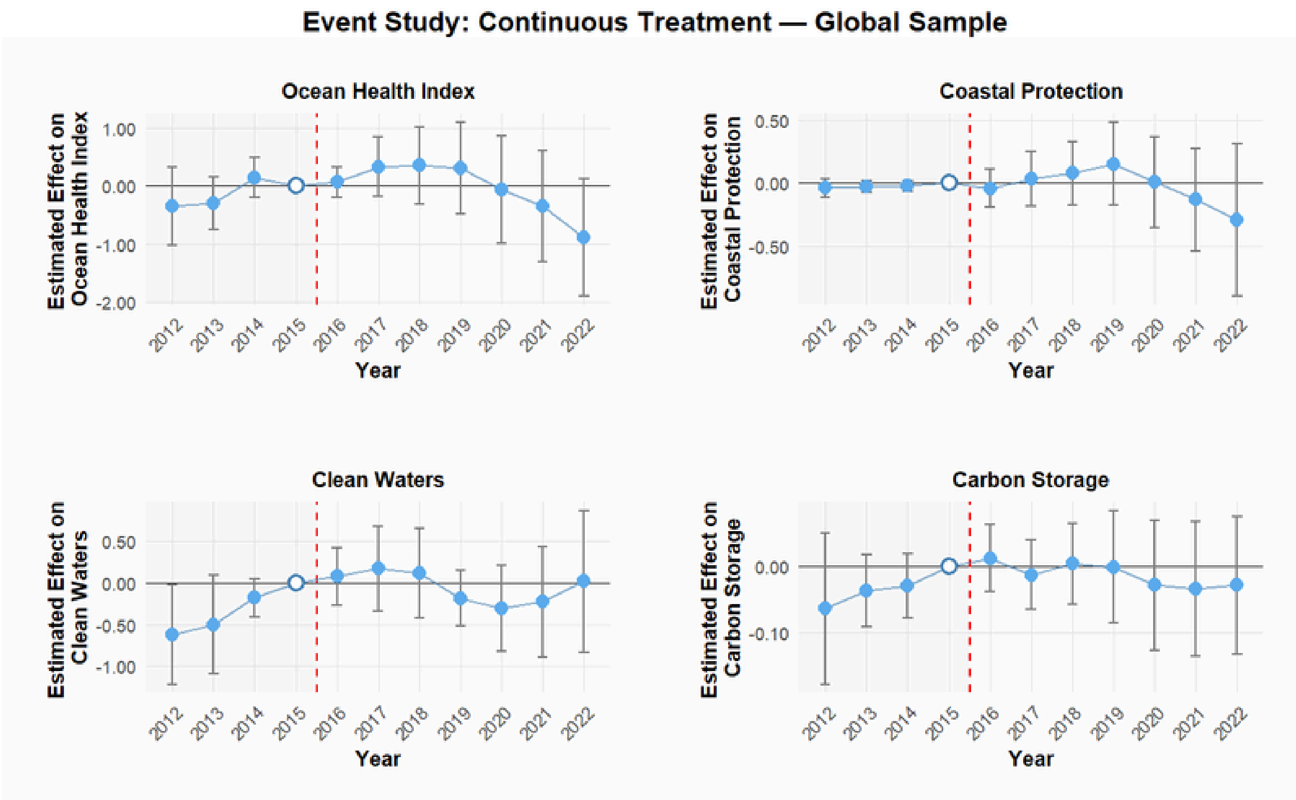

Notes: 2015 is the omitted reference period (open circle, pinned to 0). Error bars show 95% CIs clustered by country.
Treatment: IHS(MFF) demeaned. Darker points indicate 95% significance.

Figure Notes: *The event studies show the treatment effects for the four outcome variables of interest. Each event study reveals insignificant pre-trends, supporting the parallel-trends assumption. Post-treatment effect sizes are small and centered around zero. The Clean Waters event study shows the greatest variability in pre-trends, raising concerns about the assumption of parallel trends.*

6. ROBUSTNESS CHECKS

6.1: Alternative Treatment Specification

To determine if the results are sensitive to the definition of the MFF treatment, a binary treatment variable is used as a robustness check. Using a binary treatment definition, rather than a continuous treatment definition, will allow comparisons between treated and control groups, or between governments with high versus low MFF values. To select the appropriate treatment cutoff, the following cumulative density function (CDF) plots were created. The CDF plots are located in Figure 6 of the Appendix.

The results of the CDF plots show that the most even split between the treated and control groups occurs at the median treatment cutoff. The absence of any natural discontinuities in the distribution of the MFF values further justifies this treatment cutoff. This robustness check examines whether the binary treatment cutoff adjusts the coefficients for the global and SIDS OHI outcomes when pre-treatment controls are added.

Binary Treatment Regression Equation

$$OHI_{ct} = \alpha_1 +| \beta_1(HighMFF_c * postParis_t) + \delta_c + \gamma_t + \epsilon_{ct}$$

The following regression table shows the effect of the binary MFF treatment definition on global OHI Index scores, with pre-treatment controls. Across all models, the interaction between the high MFF and the years following the Paris Agreement remains statistically insignificant, consistent with the main continuous treatment results. The direction of the main coefficient is unstable across the models – it shifts from negative in the baseline model, to positive with the addition of GDP, and back to negative in the final three models. The wide confidence intervals cross zero across all models, confirming that there is no meaningful treatment effect under the binary specification.

Table 6: Effect of the Binary Marine Focus Factor on Global OHI Scores

Binary IHS(MFF) Treatment × Post-Paris — Global OHI Index with Pre-Treatment Controls

Outcome Variable	(1) Baseline	(2) +IHS GDP	(3) +Fisheries	(4) +MPA	(5) All Controls
High MFF × Post-Paris	−0.1451	0.1248	−0.0846	−0.3844	−0.0814
	(0.6701)	(0.6811)	(0.6605)	(0.6671)	(0.6932)
Pre-Treatment IHS(GDP) × Post		0.1413			0.1500
		(0.1519)			(0.1572)
Pre-Treatment Fisheries × Post			−10.1154		−5.4485
			(6.6300)		(7.1658)
Pre-Treatment % MPA × Post				−0.0511*	−0.0545*
				(0.0275)	(0.0285)
Num. Obs.	1089	1089	1089	1089	1089
R^2	0.865	0.866	0.866	0.867	0.868
Country FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Mean Dep. Var.	70.29	70.29	70.29	70.29	70.29
95% CI Lower	−1.4750	−1.2268	−1.3953	−1.7082	−1.4571
95% CI Upper	1.1848	1.4765	1.2260	0.9394	1.2943

Notes: Country-by-year level regressions for 99 countries observed annually from 2012 to 2022. Dependent variable is the OHI Index (0–100 scale). High MFF is a binary indicator equal to 1 for countries with above-median IHS(MFF) values in the global sample. Controls are pre-2016 country averages of IHS(GDP), Fisheries Value as % of GDP, and % Marine Protected Areas, each interacted with Post-Paris. All models include country and year fixed effects. Standard errors clustered at the country level are shown in parentheses. 95% CI refers to the confidence interval for the High MFF × Post-Paris coefficient. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table Notes: *Table 6 confirms that the null findings from the continuous treatment specification are robust to alternative treatment definitions. The results for the main coefficient across all models are insignificant, indicating that the results are not sensitive to arbitrary treatment definitions.*

6.2: Placebo Tests

The second robustness check is a placebo test. The placebo test assigns 2013 as the treatment year to assess whether countries had pre-existing differences in their goal scores. Since the Paris Agreement had not yet come into effect in 2013, there should be no treatment effects on any outcome variable if the parallel trends assumption holds.Insignificant results support the parallel trends assumption, whereas significant results indicate that the parallel trends assumption fails for that outcome variable. The placebo test is run on all four outcome variables.

The Index placebo test is insignificant but has the largest confidence interval among the four outcomes, reflecting greater uncertainty in this estimate. The Coastal Protection and Carbon Storage outcomes are insignificant and have the smallest confidence interval. These three outcomes show that there were no pre-existing trends driving the models' results.

The Clean Waters outcome challenges the parallel trends assumption, as it is statistically significant. This suggests that countries with higher MFF scores had pre-existing differences in their Clean Waters outcomes. While the parallel trends assumption likely does not hold for the Clean Waters goal score and the Index goal score shows large confidence intervals, the Coastal Protection and Carbon Storage goal scores are the most credible findings from this analysis.

Figure 4: Placebo Test Plot

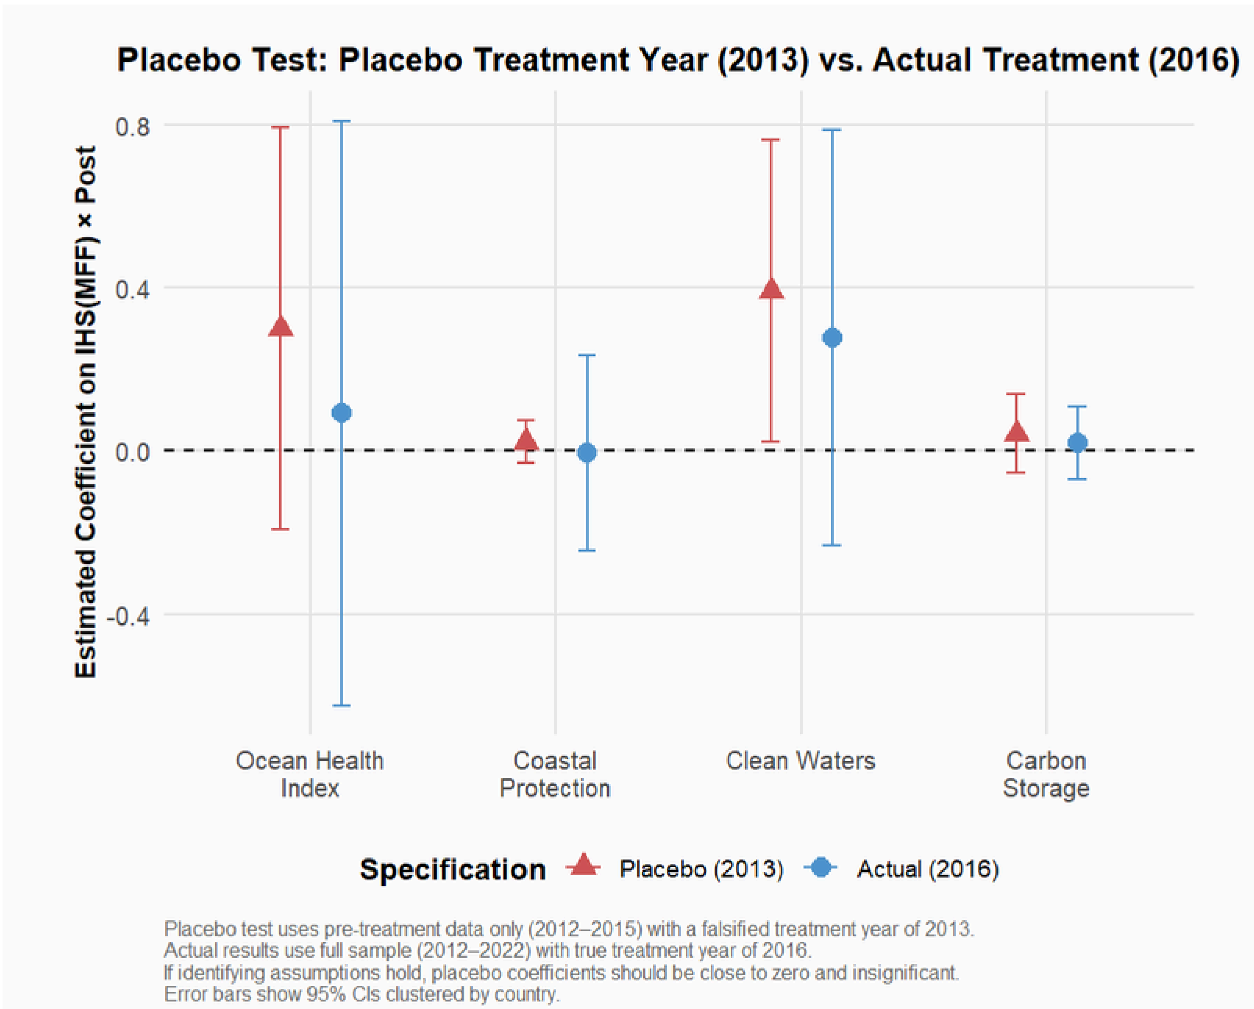

Figure Notes: *The placebo test shows mixed findings, but largely supports the identifying assumptions. The Coastal Protection and Carbon Storage outcomes both show insignificant placebo coefficients. The Clean Waters outcome raises cause for concern because it is statistically significant, indicating that the parallel trends assumption likely does not hold for this outcome.*

7. DISCUSSION

7.1: Implications and Policy Relevance

This analysis hypothesized that stronger commitments to marine issues under the Paris Agreement in 2016 (higher MFF scores) would lead to improvements in ocean health (higher OHI scores) following 2016. However, the results reveal that MFF scores had no statistically significant effect on either global or SIDS OHI goal scores. The analysis fails to reject the null hypothesis that higher MFF scores led to higher OHI scores following the Paris Agreement. Although countries explicitly listed marine-related issues in their commitments under the Paris Agreement, they either failed to implement policies that could improve ocean health or their policies were not effective in doing so within the sample time period. This could indicate that the NDCs are voluntary and non-binding, meaning that countries are not required to follow through with their commitments. The significance of the pre-treatment fisheries value and percentage of MPAs suggests that a government's baseline marine conditions and policies may matter more for ocean health than how strong their marine commitments are.

The null findings are especially concerning for SIDS. SIDS had some of the highest MFF scores out of the different negotiating groups in the analysis sample, but still did not show significant improvements in OHI scores. These nations are particularly vulnerable to ocean degradation due to their dependence on marine resources. Given their dependence on marine resources for food security, economic stability, and livelihoods, the disconnect between commitments and outcomes is worrying for the effectiveness of climate adaptation in these nations.

This analysis fills two existing research gaps. First, it demonstrates that OHI goal scores can be used as outcome variables for a quasi-experimental research design, which has rarely been applied to the OHI. Second, it provides causal evidence on the relationship between marine commitments under the Paris Agreement and ocean health, an area that had not previously been addressed using a causal framework. Despite the strong connection between climate change and ocean health, climate policies often fail to address the impacts on ocean health. A stronger link between climate and ocean policy, whereby policies such as the Paris Agreement enforce measures to mitigate the effects of climate change on the ocean, could be an effective way to support ocean health. Stronger enforcement mechanisms through direct regulation strategies, such as penalties for failing to meet goals or subsidies for strengthening climate policies, could lead to an improved connection between climate and ocean policy.

7.2: Limitations

Gallo et al. (2017) state that the MFF is a measure of intent rather than actual policy change and does not distinguish between marine commitments that support ocean health and those that could harm it, such as seawater desalination. The MFF does not reveal whether a government is focused on improving marine environments. Rather, it reveals whether a government included mentions of marine topics in their NDCs overall, regardless of intent. While these are limitations of the predictor variable, it does not undermine the research question's focus, which asks whether committing to marine issues, regardless of intent, leads to improvements in OHI outcomes.

The OHI outcome variables could present further limitations of the analysis. The OHI incorporates both ecological and human-related dimensions of ocean health. Countries with stronger pre-existing marine governance may have higher baseline OHI scores and show less improvement in OHI following the Paris Agreement. In contrast, countries with low marine governance may exhibit lower OHI baseline scores and fail to act on their marine commitments, resulting in lower outcome OHI scores. Although the Index and Clean Waters coefficients exhibit the highest magnitudes in the global and SIDS models, they also tend to have the largest standard errors. The Clean Waters goal score is the most concerning, as it fails the parallel trends assumption and placebo test. Despite the challenges with the OHI variables, the OHI still provides a unique look at how countries evaluate their ocean health.

Finally, the post-treatment period may be too short to capture the substantial impacts of marine commitments on ocean health. The post-treatment period from 2016 to 2022 may be too short to detect improvements in ocean health, which can take decades to materialize. This is an important caveat to the findings, and longer-term analyses will be needed to further test these research questions as data becomes available.

7.3: Conclusion and Future Research

The statistically insignificant causal relationship between MFF and OHI scores suggests that the link between marine commitments and measurable ocean health outcomes is weak, at least within the sample time frame. Countries may have failed to implement policies that effectively mitigate the effects of climate change on the ocean. International climate agreements, such as the Paris Agreement, need stronger enforcement mechanisms to ensure that countries' commitments translate into genuine improvements in ocean health outcomes.

While the analysis reveals no significant causal relationship between MFF and OHI scores, it guides further research questions. First, future research should focus on further deconstructing NDCs to distinguish between commitments that support and those that harm the marine environment. Second, research should examine whether countries with stronger pre-existing marine governance were more likely to craft higher marine-focused commitments. Third, future work should test whether higher MFF scores predicted domestic ocean policy changes. These research questions can clarify the relationship between marine-focused commitments and ocean health, revealing whether countries are actually implementing more ocean policies when they commit to doing so.

8. REFERENCES

Bijma, J., Pörtner, H. O., Yesson, C., & Rogers, A. D. (2013). Climate change and the oceans– What does the future hold?. Marine pollution bulletin, 74(2), 495-505.

Gallo, N. D., Victor, D. G., & Levin, L. A. (2017). Ocean commitments under the Paris Agreement. Nature Climate Change, 7(11), 833-838.

Gattuso, J. P., Magnan, A., Billé, R., Cheung, W. W., Howes, E. L., Joos, F., ... & Turley, C. (2015). Contrasting futures for ocean and society from different anthropogenic CO2 emissions scenarios. Science, 349(6243), aac4722.

Guggisberg, S. (2019). Funding coastal and marine fisheries projects under the climate change regime. Marine Policy, 107, 103352.

Ha, L. T. (2024). The Function of Green Finance Initiative in Maintaining Ocean Health: Novel Findings From an International Database. Environmental Health Insights, 18, 11786302241304616.

Halpern, B. S., Longo, C., Scarborough, C., Hardy, D., Best, B. D., Doney, S. C., ... & Samhouri, J. F. (2014). Assessing the health of the US West coast with a regional-scale application of the ocean health index. Plos one, 9(6), e98995.

Ocean Health Index. (2026). https://oceanhealthindex.org.

Sumaila, U. R., Tai, T. C., Lam, V. W., Cheung, W. W., Bailey, M., Cisneros-Montemayor, A. M., ... & Gulati, S. S. (2019). Benefits of the Paris Agreement to ocean life, economies, and people. Science advances, 5(2), eaau3855.

Ullah, A., Pinglu, C., Khan, S., & Qian, N. (2025). Transition Towards Climate-Resilient Low-Carbon and Net-Zero Future Pathways: Investigating the Influence of Post-Paris Agreement and G-7 Climate Policies on Climate Actions in G-20 Using Quasi-Natural Experiments. Sustainable Development.

United Nations. (n.d.). United Nations Convention on the Law of the Sea: Part V — Exclusive economic zone. Division for Ocean Affairs and the Law of the Sea. https://www.un.org/depts/los/convention_agreements/texts/unclos/part5.htm

9. APPENDIX

9.1: Figures

Figure 5: Distribution of Marine Focus Factor Scores

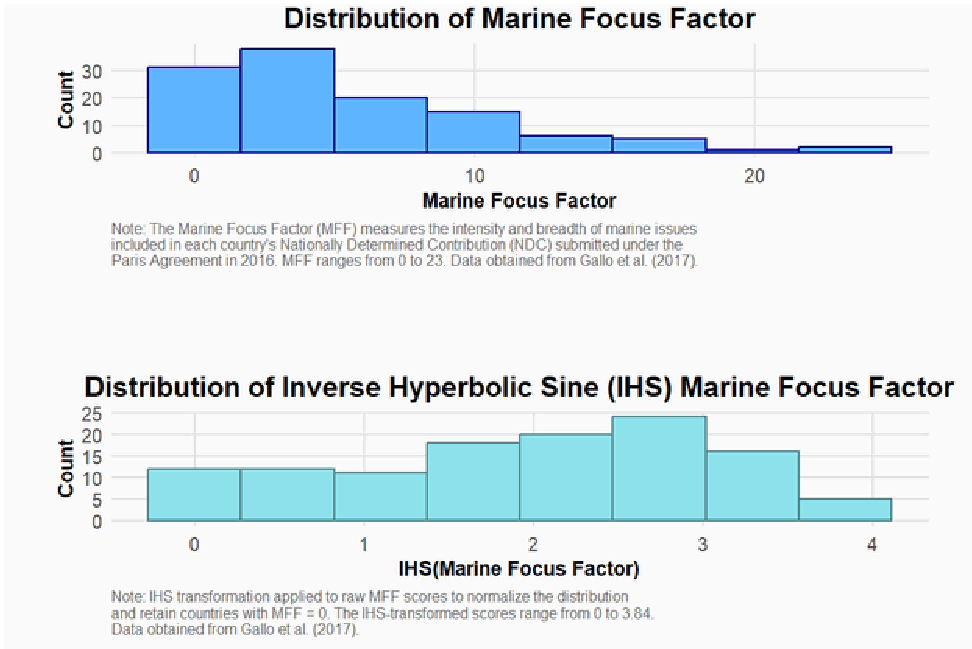

Figure Notes: *The above histogram shows the distribution of raw MFF scores and logarithmic MFF scores. While the raw scores are right-skewed, the log-transformed MFF scores approximate a normal distribution.*

Figure 6: Distribution of Marine Focus Factor Scores by Negotiating Group

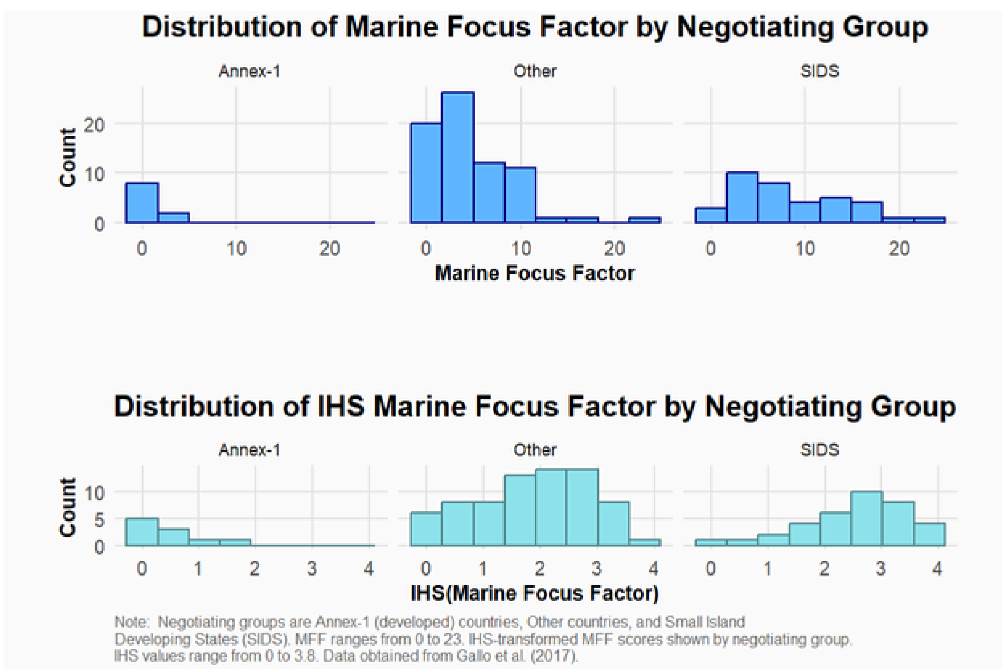

Figure Notes: *The above plots show the distribution of the raw MFF scores and the IHS-transformed MFF scores by negotiating group. The IHS transformation shows a normal distribution for the SIDS negotiating group, where MFF values peak between scores of 2.5 and 3.*

Figure 7: Cumulative Density Function Plot for Treatment Definition

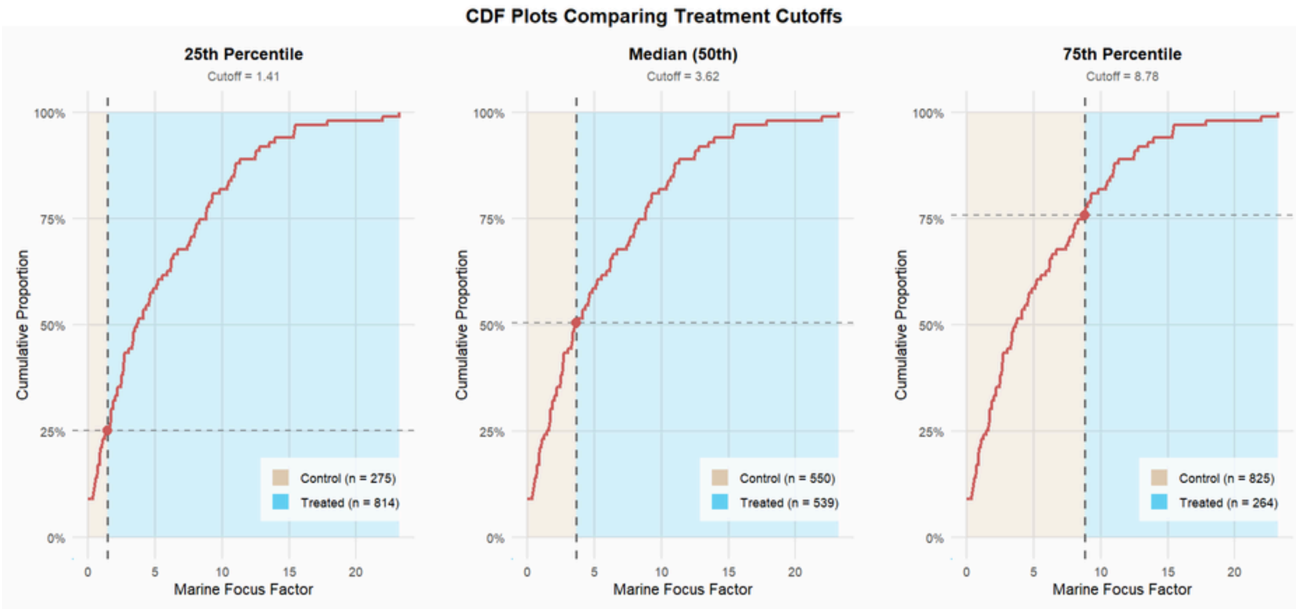

Figure Notes: *The above CDF plots compare how the composition of the treatment and control groups changes across three different binary thresholds – the 25th percentile, the 50th percentile, and the 75th percentile. The median cutoff yields the most balanced split, with 550 observations in the control group and 539 observations in the treated group. Because there is no natural break or discontinuity in the spread of MFF values, the balanced group sizes under the median split provide the strongest justification for using this binary treatment threshold in the analysis.*

9.2: Tables

Table 7: SIDS Subset Summary

SIDS Countries in the Analysis Sample[1]

Sorted by Marine Focus Factor Score (Descending)

	Marine Commitment		Avg. OHI Goal Scores (2012–2022)			
Country	**MFF**	**IHS (MFF)**	**OHI Index**	**Coastal Prot.**	**Clean Waters**	**Carbon Storage**
Maldives	23.21	3.84	72.58	67.04	72.76	69.41
Bahrain	17.84	3.58	72.65	77.96	88.01	97.42
Seychelles	15.42	3.43	81.99	79.28	92.22	99.96
Mauritius	15.32	3.42	76.32	85.34	73.81	99.71
Bahamas	13.91	3.33	77.21	39.83	86.12	84.87
Belize	13.48	3.30	70.85	43.37	85.29	96.55
Dominica	12.79	3.24	72.43	96.78	71.47	100.00
Suriname	12.47	3.22	75.22	71.55	83.01	82.89
Haiti	10.93	3.09	65.18	57.28	44.58	94.80
Solomon Islands	9.28	2.92	74.23	66.60	64.95	74.59
Grenada	9.00	2.89	67.25	56.39	73.89	92.53
Sao Tome and Principe	7.99	2.78	68.13	97.88	61.06	97.88
Antigua and Barbuda	7.65	2.73	81.07	60.52	79.48	98.85
Palau	7.41	2.70	76.45	79.92	91.72	71.70
Barbados	6.18	2.52	72.31	99.32	75.21	100.00
Fiji	6.15	2.52	77.15	84.61	87.97	81.17
Singapore	5.87	2.47	71.04	71.27	62.29	67.78
Marshall Islands	5.49	2.40	80.09	80.90	74.80	100.00
Cuba	4.53	2.22	67.64	12.32	78.34	91.75
Tonga	4.08	2.11	72.70	91.35	84.75	99.97
Vanuatu	3.47	1.96	69.11	76.70	79.25	71.45
Jamaica	3.37	1.93	70.85	75.95	56.43	99.53
Tuvalu	3.32	1.91	76.30	84.51	74.57	100.00
Guyana	2.66	1.70	60.31	31.96	87.64	78.56
Papua New Guinea	2.58	1.68	74.71	73.61	73.23	71.35
Dominican Republic	1.71	1.30	73.66	92.09	57.38	84.35
Trinidad and Tobago	1.09	0.95	71.14	91.43	74.22	99.03
Samoa	0.44	0.42	78.33	95.91	89.96	99.97

[1] MFF scores from Gallo et al. (2017). OHI scores from Ocean Health Index (2026). IHS = Inverse Hyperbolic Sine.

Table Notes: *The table above shows all countries in the SIDS subset, along with their MFF and OHI goal scores. The table is sorted to show countries with the highest MFF scores at the top and the lowest scores at the bottom.*

Table 8: Effect of the Binary Marine Focus Factor on SIDS OHI Scores

Binary IHS(MFF) Treatment × Post-Paris — SIDS OHI Index with Pre-Treatment Controls

Outcome Variable	(1) Baseline	(2) +IHS GDP	(3) +Fisheries	(4) +MPA	(5) All Controls
High MFF × Post-Paris	1.1788	1.1865	1.1913	1.2072	1.3110
	(1.3423)	(1.3675)	(1.4694)	(1.3543)	(1.5080)
Pre-Treatment IHS(GDP) × Post		0.1035			0.2143
		(0.3605)			(0.3356)
Pre-Treatment Fisheries × Post			1.0009		7.0516
			(28.6388)		(32.8896)
Pre-Treatment % MPA × Post				−0.0859	−0.0877
				(0.0766)	(0.0699)
Num. Obs.	308	308	308	308	308
R^2	0.768	0.768	0.768	0.772	0.773
Country FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Mean Dep. Var.	73.10	73.10	73.10	73.10	73.10
95% CI Lower	−1.5754	−1.6195	−1.8236	−1.5715	−1.7832
95% CI Upper	3.9331	3.9925	4.2061	3.9859	4.4052

Notes: Country-by-year level regressions for Small Island Developing States (SIDS) observed annually from 2012 to 2022. Dependent variable is the OHI Index (0–100 scale). High MFF is a binary indicator equal to 1 for SIDS with above-median IHS(MFF) within the SIDS subsample. Controls are pre-2016 country averages of IHS(GDP), Fisheries Value as % of GDP, and % Marine Protected Areas, each interacted with Post-Paris. All models include country and year fixed effects. Standard errors clustered at the country level are shown in parentheses. 95% CI refers to the confidence interval for the High MFF × Post-Paris coefficient. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table Notes: *The above table shows the regression results for the effect of sequentially adding controls on OHI for the SIDS subset, using the binary MFF treatment. There are no statistically significant results in the table, but the main coefficient gains magnitude with the addition of controls. The wide intervals reflect both the small SIDS sample size and the instability of the binary treatment specification within this subsample.*

Renewable Portfolio Standards and the Geography of Utility-Scale Solar Deployment in the United States

Lauren Silver

ABSTRACT
This paper examines how Renewable Portfolio Standards (RPS) shape the geographic distribution of utility-scale solar deployment across counties with strong solar resources and existing grid infrastructure. Using county-level data, the study combines an Exploratory Data Analysis with a Staggered Difference-in-Differences and a Sun-Abraham Event Study. The results show that RPS policies significantly increase solar buildout, but the spatial distribution of new capacity is strongly shaped by underlying solar potential and grid infrastructure. These findings suggest that RPS policies alone may not ensure spatially equitable renewable energy expansion, and that complementary policies targeting infrastructure and resource availability may be necessary.

Keywords: Renewable Portfolio Standards, Utility-Scale Solar, Difference-in-Differences

1 INTRODUCTION

1.1 Policy Context

The transition to renewable energy is a central component of energy policy in the United States. One of the most widely used state-level policies to promote renewable energy generation are Renewable Portfolio Standards (RPS), which require utilities to supply a minimum share of electricity from renewable sources. As of 2025, the majority of U.S. states have adopted some form of RPS policy, making it one of the primary policy tools used to expand renewable energy capacity. While RPS policies generally include renewable energy sources such as wind, solar, and geothermal power, the specific qualifying technologies vary by state, with some states establishing detailed statutory eligibility rules while others leave implementation to regulatory agencies.

While previous literature has examined whether RPS policies increase overall renewable energy generation, much less attention has been paid to where renewable energy infrastructure is built within states. The spatial distribution of renewable energy deployment is important for policy equity. Renewable infrastructure tends to concentrate in areas with favorable natural resources and existing energy infrastructure, which may lead to uneven geographic distribution of investment. If RPS policies disproportionately direct development toward certain counties while leaving others behind, the benefits and costs of the energy transition may be unevenly shared. Understanding how RPS policies shape the geographic allocation of renewable energy development therefore provides important insights for both energy policy design and regional equity.

1.2 Research Question and Hypothesis

This study focused on the following research question: How do Renewable Portfolio Standards shape the geographic distribution of utility-scale solar deployment across counties, and does deployment occur disproportionately in areas with stronger solar resources and transmission infrastructure? Based on the literature the hypothesis is as follows: After a state adopts an RPS, utility-scale solar is disproportionately deployed in counties with higher solar potential and stronger grid infrastructure.

1.3 Main Findings

This study uses a county-level dataset for an Exploratory Data Analysis (EDA), a Staggered Difference-in-Differences (SDiD), and a Sun-Abraham Event Study. The results show that RPS adoption is associated with a statistically significant increase in utility-scale solar deployment. However, the spatial distribution of new solar capacity is strongly shaped by underlying county characteristics. Counties with higher solar potential and stronger transmission infrastructure tend to experience substantially larger increases in solar buildout following RPS adoption. These findings suggest that while RPS policies can be effective at increasing renewable energy development, they may also reinforce existing geographic advantages in resource availability and grid capacity.

The remainder of this paper proceeds as follows. Section 1 reviews the relevant literature on RPS and renewable energy deployment. Section 2 describes the data and Section 3 explains the empirical methodology. Section 4 presents the main results from the EDA, SDiD, and Event Study, including postestimation, mechanisms, and heterogeneity. Section 5 describes the robustness checks used in this study. Section 6 concludes with policy implications.

1.4 Literature Review

1.4.1 Solar Buildout

The current literature around solar buildout is mostly correlational and descriptive, with much of the focus on residential solar. O'Shaughnessy (2019) found that state and local policies affect the number of solar installers that compete in the local, residential PV market. Hence, at the residential solar level, policies have been shown to impact the solar buildout. While there are multiple studies around residential solar, the findings can not be extrapolated to apply to utility-scale solar as they have different locations, purposes, sizes, capacities, and costs.

When it comes to utility-scale solar research, existing research remains relatively limited and largely descriptive. Ding et al (2023) examined the spatial distribution of solar farms across regions of different socio-demographic characteristics and the impact of solar policy on the spatial distribution. The study found that solar policy is an effective tool in promoting utility-scale solar buildout, but their use of Geographical Detector Models and Kernel Density analysis does not establish causal effects. However, this literature primarily focuses on cross-sectional or national-level patterns and does not examine fine-grained, within-state variation over time.

1.4.2 Renewable Portfolio Standards

A smaller body of literature specifically examines Renewable Portfolio Standards (RPS), with mixed findings driven largely by differences in outcomes and empirical approaches. Kroeger and Burgess (2023) found that electric utility plans align with RPS targets, while Wiser et al (2011) found that RPS programs can lead to the development of solar energy and to renewable resource diversification. Shrimali et al (2015) and Yin and Powers (2010) found a positive relationship between RPS adoption and renewable energy electricity generation.

However, other studies find weaker or inconsistent effects. Feldman and Levinson (2023) found that in some cases, RPS programs do not appear to reduce the use of natural gas while increasing the use of wind and solar power. Deschênes et al (2023) found that RPS policies have no significant effect on investments in solar capacity. The mixed findings in the RPS literature are from differing outcome variables and identification methods, rather than pure disagreement about RPS efficacy. Furthermore, even if RPS adoption increases solar buildout on average, studies have not looked at utility-scale solar siting.

1.4.3 Literature Gap

Current studies have looked at the relationship between solar-focused policies, techno-economic factors, and socioeconomic factors to solar buildout. However, how solar policies impact spatial distribution of solar has not been extensively studied. While multiple studies look at the effect of RPS adoption on state-level solar buildout, they do not investigate whether RPS adoption impacts the within-state county-level solar distribution. While solar policies as a whole have been studied, the relationship between RPS adoption and solar buildout remains understudied and inconclusive. The current literature leaves two key gaps. First, a conceptual gap: much of the literature treats solar policy as a broad category without isolating the role of RPS. Second, a methodological gap: existing studies do not analyze county-level variation in utility-scale solar buildout over time. This paper contributes by examining how RPS policies are associated with the within-state spatial allocation of utility-scale solar development across counties.

2 DATA SOURCES

The monadic panel dataset spans from 2000 to 2023 and covers all U.S. counties. An observation in the dataset is a county by year. For more information on data sources, cleaning, transformations, merging, and definitions, refer to Appendix - Table 1 and Appendix - Figure 1. The final dataset was constructed by merging the sources into a balanced county-year panel. County-level solar installation data were first aggregated to the county-year level and linked to county GEOIDs using a county name crosswalk. These observations were then merged with state-year RPS data by state and year so that policy status did not depend on whether a county had solar development in a given year. Static county characteristics, including solar potential and transmission grid infrastructure were aggregated to the county level then merged by GEOID, while socioeconomic controls were added using county-level Decennial Census 2000 data and ACS 5-year estimates. To maintain temporal consistency, 2000 Census values were used for 2000–2005, 2009 ACS values were held for 2006–2009, annual ACS values were used for 2010 onward, and the most recent ACS year was carried forward for later years where needed.

2.1 Dependent Variable Utility-Scale Solar Buildout.

The utility-scale solar buildout dataset is from the U.S. Geological Survey from the U.S. Department of the Interior (Fujita, K.S.). It includes the locations and array boundaries of U.S. photovoltaic (PV) facilities with solar capacity of 1 megawatt or more, spanning from 1985 to 2024 across all U.S. states. It has been transformed into Inverse Hyperbolic Sine of Utility-Scale Solar MW Added, to accommodate zero values and reduce skewness, allowing for a more stable and interpretable estimation while preserving observations with no solar buildout.

2.2 Independent Variables

2.2.1 Treatment Variable

RPS Targets. Therenewable portfolio standards dataset comes from Energy Markets & Planning Berkeley Lab (Barbose, Galen). It covers 2000-2050 renewable portfolio standards targets. The dataset is updated as of 2024. RPS targets are defined as the mandated percentage of electricity sales that must be generated from renewable sources in a given state-year. In the main specification, RPS is measured as a binary indicator equal to 1 if a state has an RPS policy in effect in a given year. As a robustness check, a continuous measure of RPS stringency, defined as the required renewable share (in percentage points), is also used.

2.2.2 Explanatory Variables

U.S. Electric Power Transmission Lines. The dataset for U.S. electric power transmission lines comes from Esri U.S. Federal Datasets (U.S. Climate Resilience Toolkit), updated as of 2025, and includes transmission features from 1996 to 2024. Each observation represents an individual transmission line segment. In this study, these data are aggregated to the county level to construct a measure of grid infrastructure, defined as the total length of transmission lines (in kilometers) within each county-year. This variable has been transformed into log +1 (Appendix - Figure 1 for justification and visuals).

Solar Potential (Global Horizontal Irradiation). Solar potential has been included using global horizontal irradiation (GHI) which is the sum of direct and diffuse irradiation components received by a horizontal surface and is measured in kilowatt hours per square meter. GHI includes air temperature, wind and snow, atmospheric pollution and thickness, dust, and other geological factors. This dataset comes from the World Bank - Global Solar Atlas ("Global PV Potential Study."). The data is a time-invariant, or a long-term yearly average of daily and yearly totals. In this study, GHI is assigned at the county level and used as a measure of underlying solar resource availability. Solar Potential has been transformed into log, see Appendix - Figure 1 for justification and visuals.

2.3 Control Variables

Population and Median Income. County-level population and median household income data are obtained from the U.S. Census Bureau. For 2009–2024, values come from the American Community Survey (ACS). Due to data availability, earlier years are constructed using Decennial Census 2000 values for 2000–2005 and ACS 2009 values for 2006–2009. Population is defined as the total number of residents in a county, and median household income is measured as the median income of households in the past 12 months, reported in inflation-adjusted dollars. In the analysis, both variables are included in log form.

3 EMPIRICAL METHODOLOGY

3.1 Staggered Difference-in-Differences

To estimate the effects of RPS adoption on utility-scale solar buildout for higher solar potential or higher grid infrastructure counties, this study uses a SDiD design. Appendix - Figure 2 shows all the treated states that are included in the SDiD. Model 1 identifies the average treatment effect of RPS adoption by comparing changes in solar deployment in treated counties to counties in states that have not yet adopted an RPS or never adopted one. Model 2 shows if after RPS adoption, if utility-scale solar is disproportionately allocated in counties with higher solar potential. Similarly, Model 3 determines if after RPS adoption, is utility-scale solar buildout disproportionately in counties with higher grid infrastructure. This study employs these three main estimating equations where s indexes states, c indexes counties, and t indexes years:

$$\text{Model 1: } Solar_{sct} = \beta\,(RPS_{st}) + X_{ct}'\,\theta + \mu_c + \delta_t + \varepsilon_{sct}$$
$$\text{Model 2: } Solarsct = \beta\,(RPS_{st} \times \text{IHS } SolarPotential_c) + X_{ct}'\,\theta + \mu_c + \lambda_{st} + \varepsilon_{sct}$$
$$\text{Model 3: } Solarsct = \beta\,(RPS_{st} \times \text{IHS } Grid_c) + X_{ct}'\,\theta + \mu_c + \lambda_{st} + \varepsilon_{sct}$$

The outcome variable (denoted as $Solar_{sct}$) is county-level utility-scale solar capacity additions, measured in megawatts and transformed to IHS. RPS_{st} denotes the treatment variable, which is a binary variable for RPS adoption equal to 1 beginning in the year a state adopts an RPS and remaining 1 thereafter. Model 1 β is interpreted as the average change in solar buildout when a state adopts an RPS. For Model 2, the treatment variable is interacted with a key explanatory variable, Solar Potential, denoted as $SolarPotential_c$. Model 2 β is interpreted as how the effect of RPS adoption changes with solar potential. While Model 3 is the treatment variable interacted with Grid Infrastructure, written as $Grid_c$. For Model 3, β is interpreted as how the RPS effect changes with transmission infrastructure. Solar Potential, Grid Infrastructure, and Solar Buildout are at the county level, while RPS adoption is at the state level.

For all three models, County Fixed Effects (FE) (denoted as μ_c) have been included to control for time-invariant county characteristics such as land availability, geography, historical solar buildout, etc. For Model 1 additionally includes Year FE (δ_t) to account for national shocks. State-by-year fixed effects are not included in Model 1 due to perfect collinearity with the RPS indicator. For Models 2 and 3, State-by-Year FE (written as λ_{st}) have been included which absorb all state-level variation over time, including RPS adoption. As a result, the interaction coefficients are identified from within-state, cross-county differences in solar potential and grid infrastructure following adoption.

County-level controls have been added, denoted as $X_{ct}'\theta$, to control for time-varying county-level logged population and median income. Log population controls for scale effects, meaning larger counties may host more projects. Log median income accounts for local economic capacity, which could influence solar land markets or political support. Standard errors are clustered at the state level to account for serial correlation in policy adoption and unobserved shocks within states.

3.1.1 Identification Concerns

Despite these controls, some identification concerns remain. For example, transmission expansion may be endogenous if new lines are built in anticipation of solar buildout. Local permitting policies could also coincide with RPS adoption. In addition, RPS stringency may vary in ways not captured by a binary policy measure. In order to address these concerns, these robustness checks are employed: (1) replacing the binary RPS variable with continuous RPS targets and (2) using log + 1 MW Added as an alternative dependent variable. Furthermore, to check mechanisms, this study estimates the model using the number of projects to test whether RPS primarily accelerates project timing or increases total siting. To check heterogeneity this study interacts RPS with tertiles of solar potential and grid infrastructure to test nonlinear allocation effects. This specification addresses key threats to identification. Firstly, selection bias occurring from permanent differences between high and low solar counties has been addressed through the County FE. Simultaneity concerns are mitigated by using policy timing variation that is plausibly exogenous to short-run county-level solar deployment trends.

3.2 Sun-Abraham Event Study

The key identifying assumption is the parallel trends assumption: absent RPS adoption, treated and untreated counties would have followed similar trends in utility-scale solar buildout. For the interaction models, the assumption is that counties with different levels of solar potential or grid infrastructure would have followed parallel trends within states in the absence of RPS adoption. This assumption could be violated if high-solar counties were already experiencing faster growth independent of RPS adoption or if states adopted RPS policies in response to accelerating solar growth. As a result, the estimates should be interpreted as policy-associated effects rather than strictly causal impacts.

To address these concerns, this study estimates a dynamic event-study specification using the Sun and Abraham (2021). This event study equation replaced RPS_st with leads and lags of time relative to adoption, which allows for visual inspection of pre-treatment coefficients. Flat and statistically insignificant pre-period coefficients provide evidence consistent with parallel trends. The Sun–Abraham event-study specification is estimated as follows, where 1{t−FirstRPS_s=k} denotes an event-time indicator equal to one when a county in state s is k years relative to RPS adoption:

Model 1: $$\text{IHS}(MW\_added_{ct}) = \sum_{k \neq -1} \Big(\beta_k + \delta_k \log(SolarPotential_c)\Big)\, \mathbf{1}\{t - \text{FirstRPS}_s = k\} + X'_{ct}\theta + \mu_c + \lambda_t + \varepsilon_{ct}.$$

Model 2: $$\text{IHS}(MW\_added_{ct}) = \sum_{k \neq -1} \Big(\beta_k + \gamma_k \log(GridKm_c)\Big)\, \mathbf{1}\{t - \text{FirstRPS}_s = k\} + X'_{ct}\theta + \mu_c + \lambda_t + \varepsilon_{ct}.$$

Model 3: $$\text{IHS}(MW\_added_{cst}) = \sum_{k \neq -1} \beta_k \cdot \mathbf{1}\{t - \text{FirstRPS}_s = k\} + X'_{ct}\theta + \mu_c + \lambda_t + \varepsilon_{cst}.$$

4 RESULTS

4.1 Exploratory Data Analysis

This paper focuses first on an EDA to motivate the SDiD and Event Study. The EDA includes the spatial distribution of utility-scale solar facilities overlaid on the county-level average solar potential and grid infrastructure, solar buildout by RPS adoption cohort, and a snapshot of the summary statistics organized by states with and without an RPS.

4.1.1 Solar Potential

Figure 1 shows Utility-Scale Solar Facilities and Solar Potential at the county level. In the North-East and South-West increased solar buildout is spatially located. This pattern is interesting as the highest solar potential is in the South-West and Florida. Unexpectedly, in the North-East there is solar buildout even though the solar potential is low. There is also spatial solar facilities clustering in the East where there is moderate-high solar potential, which is an expected finding.

Figure 1. Solar Potential and Utility-Scale Solar Facilities

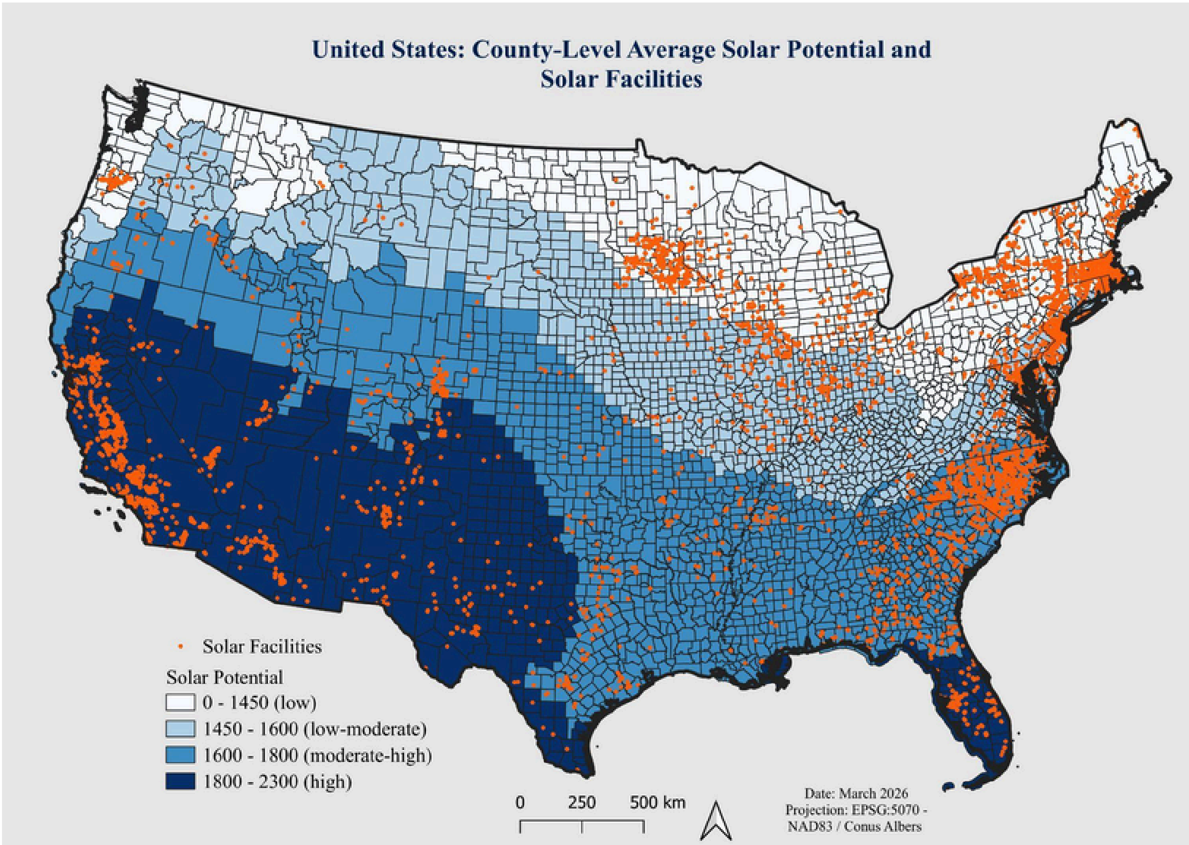

Notes: *Figure 1 includes county-level solar potential (sum of direct and diffuse irradiation components received by a horizontal surface and is measured in kilowatt hours per square meter) and solar facilities (U.S. PV facilities with solar capacity of 1 megawatt or more). The solar potential data comes from the World Bank - Global Solar Atlas. The solar potential data is time invariant. Utility-Scale Solar data comes from the U.S. Geological Survey from the U.S. Department of the Interior. The Utility-Scale Solar data is updated as of 2024.*

4.1.2 Grid Infrastructure

Figure 2 shows that solar facilities tend to be spatially located in counties with high grid infrastructure. Many solar facilities are located in the Southwest and Northeast, where there is also greater grid infrastructure. This finding, compared to Figure 2, is unsurprising as solar facilities are likely to be built out where there is grid infrastructure. Hence, Figure 2 shows that solar facilities tend to be spatially located in counties where there is higher grid infrastructure.

Figure 2. Grid Infrastructure and Utility-Scale Solar Facilities

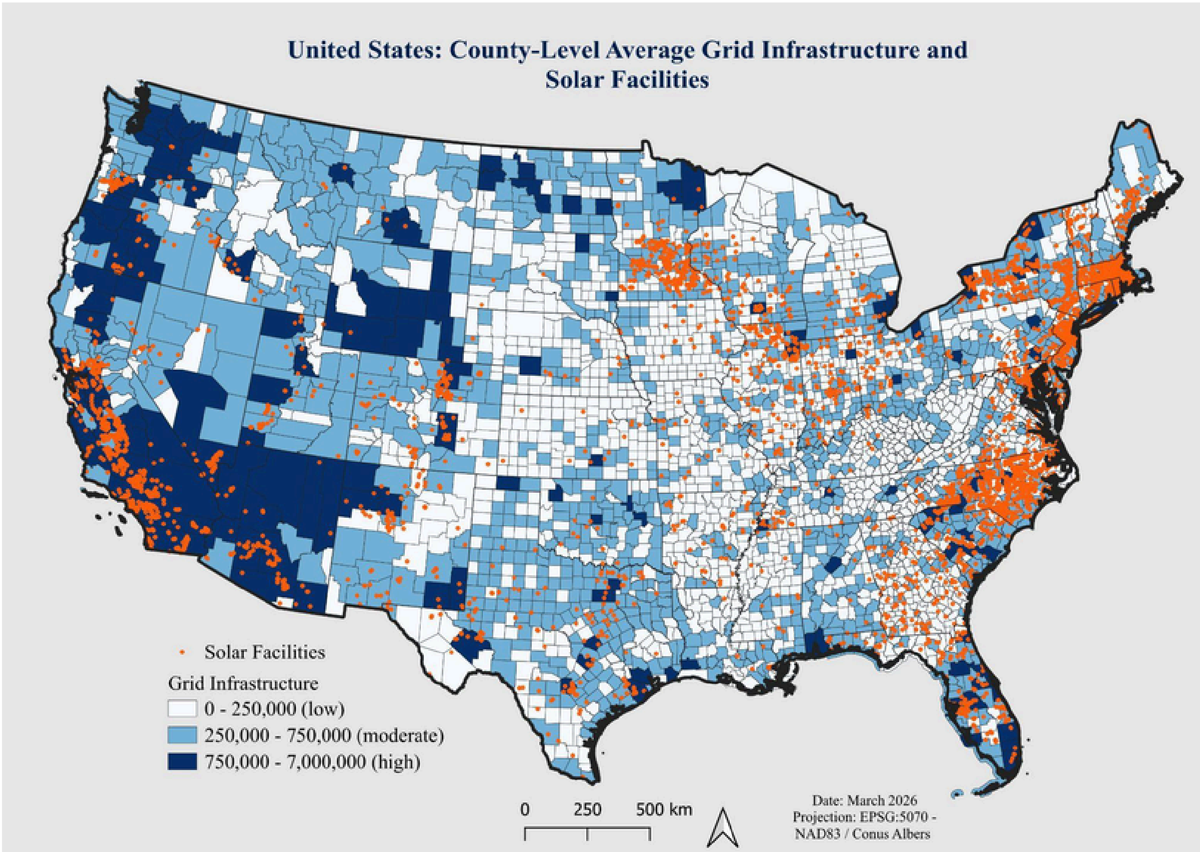

Notes: *Figure 2 includes county-level grid infrastructure (electric power transmission feature) and Utility-Scale Solar (U.S. PV facilities with solar capacity of 1 megawatt or more). The grid infrastructure data comes from the Esri U.S. Federal. The grid infrastructure data is updated as of 2024. Utility-Scale Solar data comes from the U.S. Geological Survey from the U.S. Department of the Interior. The Utility-Scale Solar data is updated as of 2024.*

4.1.3 RPS Adoption

Figure 3 shows utility-scale solar buildout by RPS adoption cohort, while Figure 4 maps the geographic distribution of these cohorts across U.S. states. The main finding from this part of the EDA is that counties in states that never adopted an RPS exhibit the lowest levels of solar buildout among the four groups. This pattern, combined with the spatial variation in policy timing shown in Figure 4, further motivates the SDiD and event study approaches. Another notable pattern is that early and mid adopters both experienced a peak in solar buildout around 2021, potentially reflecting broader energy market dynamics following the COVID-19 shock. Lastly, across all cohorts, solar buildout remained near zero until around 2006–2007, consistent with the fact that only 11 states had adopted an RPS by 2005. This suggests that pre-treatment trends were broadly similar across groups, a pattern further supported by the summary statistics in Table 1.

Figure 3. Solar Buildout by RPS Adoption Cohort

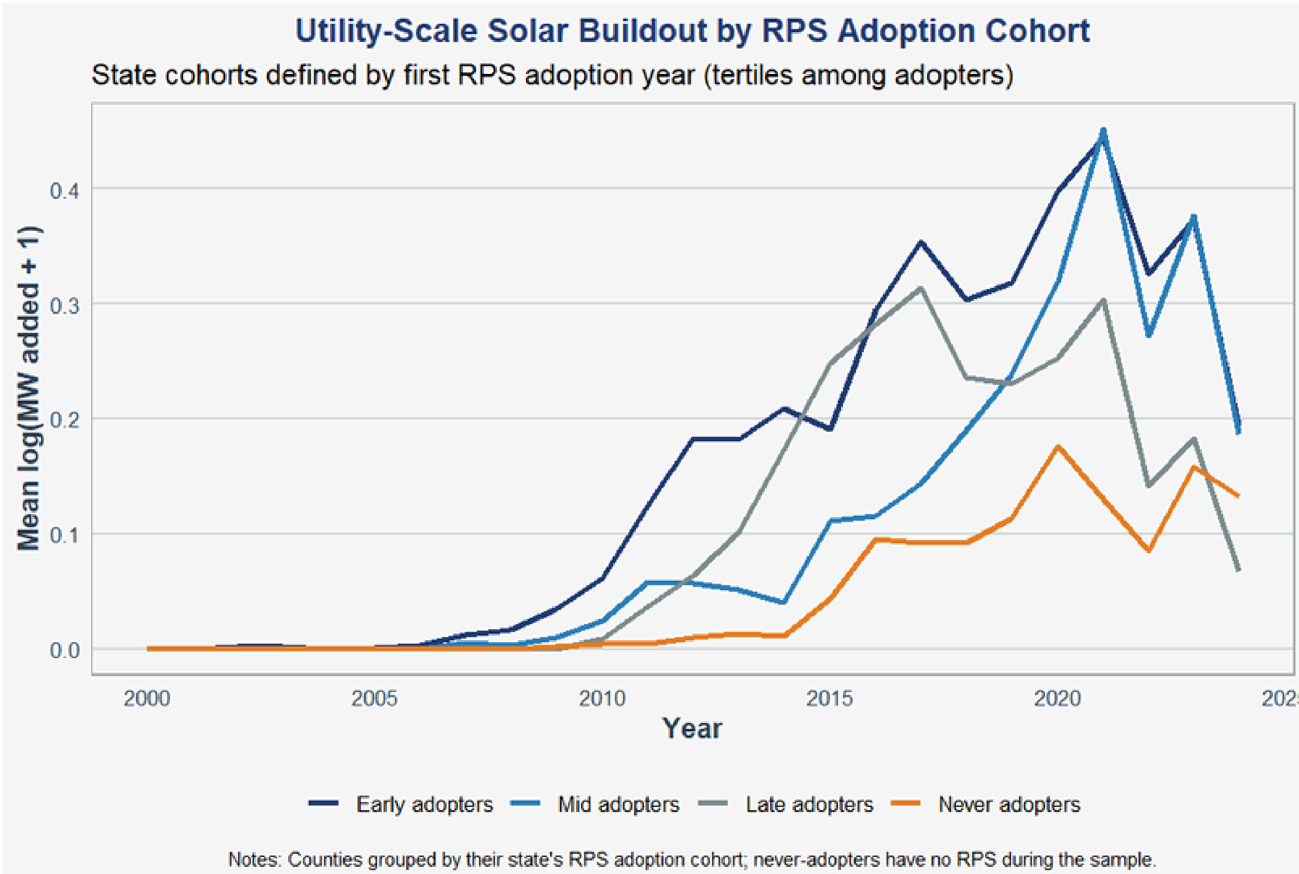

Notes: *Figure 3 includes Utility-Scale Solar by RPS Adoption cohort. Utility-Scale Solar data comes from the U.S. Geological Survey from the U.S. Department of the Interior. Utility-Scale Solar is defined as solar facilities with capacity of 1 megawatt or more. The Utility-Scale Solar data is updated as of 2024. The RPS data comes from the Energy Markets & Planning Berkeley Lab. The RPS data is updated as of 2024.*

Figure 4. RPS Adoption Cohort

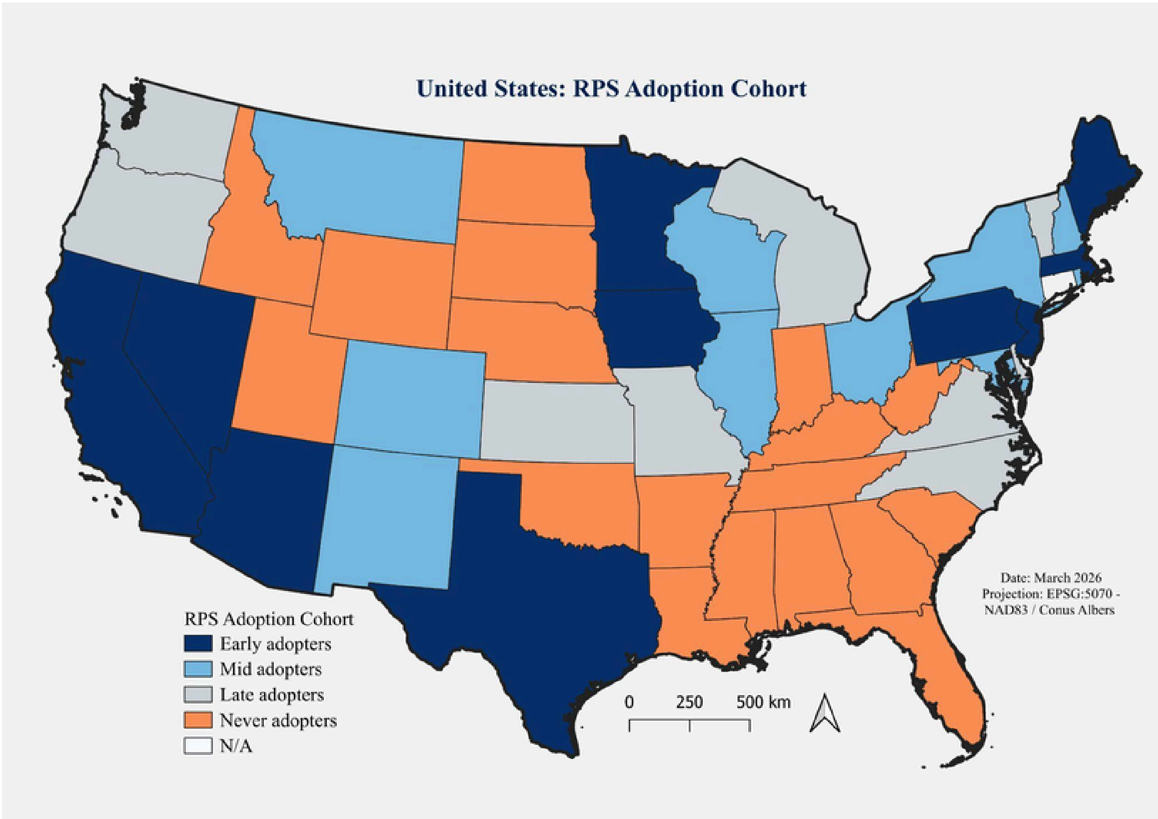

Notes: *Figure 4 maps RPS adoption cohorts by state based on the timing of first Renewable Portfolio Standard adoption. States are grouped into early, mid, and late adopters using tertiles of adoption year, with non-adopting states categorized as "Never adopters." The RPS data comes from the Energy Markets & Planning division at Lawrence Berkeley National Laboratory and is updated as of 2024.*

4.1.4 Summary Statistics

For Table 1, the Control group includes states that are never treated, while the Treated groups includes states that have adopted an RPS. The Treated groups refer to states that adopt an RPS Pre and post refer to years relative to adoption. Notably, the average Solar MW added between the Control group and the Treated (Pre) group is similar (see Appendix - Table 2 for the t-test of Mean Differences Across Groups), compared to the Treated (Post) group, showing that before adoption states see similar pretrends of solar buildout. Interestingly, the average Solar MW added is slightly higher for the Control group than the Treated (Pre) group. The average solar MW added can be slightly higher in the Control group because solar deployment has been driven by national incentives, declining technology costs, and favorable solar resources, allowing some non-RPS states to experience solar growth even without adopting an RPS policy. Another notable difference is seen between the Treated (Pre) and Treated (Post) Solar IHS Utility-Scale Solar MW Added, as the Post group has about fourteen times more Solar MW Added compared to the Pre group.

Table 1. Summary Statistics by RPS Status

Summary Statistics by Treatment Status

	Control			Treated (Pre)			Treated (Post)		
	N	Mean	SD	N	Mean	SD	N	Mean	SD
IHS Utility-Scale Solar MW Added	32088	0.052	0.453	14429	0.016	0.230	27715	0.225	0.856
Log Solar Potential (GHI)	31824	7.387	0.101	14429	7.344	0.101	27715	7.368	0.129
Log Transmission Line Length (km + 1)	32088	4.342	1.650	14429	4.261	1.743	27715	4.445	1.775

Notes: *Table 1 includes summary statistics for the Control, Treated (Pre), and Treated (Post) groups. Control group means the state never adopted an RPS. Treated counties are those in states that adopt an RPS. Pre and post refer to years relative to adoption. Utility-Scale Solar data comes from the U.S. Geological Survey from the U.S. Department of the Interior. Utility-Scale Solar is defined as solar facilities with capacity of 1 megawatt or more and is updated as of 2024. The RPS data comes from the Energy Markets & Planning Berkeley Lab and is updated as of 2024. The solar potential data comes from the World Bank - Global Solar Atlas, is time invariant, and is defined as the sum of direct and diffuse irradiation components received by a horizontal surface and is measured in kilowatt hours per square meter. The grid infrastructure data comes from the Esri U.S. Federal and is updated as of 2024 and is defined as electric power transmission features. The sample includes all U.S. counties from 2000 to 2023. RPS status is state-year.*

In conclusion, this EDA shows that in states with an RPS there tends to be an increase in utility-scale solar buildout, and this is more prominent in counties with higher solar potential and with greater grid infrastructure. In order to investigate the question the study now turns towards a SDiD multivariate regression analysis then a Sun-Abraham Event Study.

4.2 Staggered Difference-in-Differences Multivariate Regression Analysis

Table 2 is the SDiD model with three main models: Model 1 identifies the average treatment effect of RPS adoption; Model 2 shows RPS adoption x Solar Potential; while Model 3 shows RPS adoption x Grid Infrastructure. From Model 1 we see that RPS adoption is associated with a modest increase in IHS Utility-Scale Solar additions. Specifically, RPS adoption is associated with approximately an 11% increase in IHS-transformed utility-scale solar capacity added in a county, holding observed controls constant and accounting for county and year fixed effects.

Models 2 and 3 examine spatial heterogeneity. The interaction between RPS adoption and solar potential is positive and statistically significant, indicating that counties with greater solar resources tend to experience substantially larger increases in solar deployment following RPS adoption. Specifically, after RPS adoption, a 1% increase in solar potential is associated with a 0.83% increase in solar buildout. Similarly, the interaction between RPS adoption and grid infrastructure is positive and significant, suggesting that counties with more transmission infrastructure tend to capture a larger share of solar development. After RPS adoption, a 1% increase in transmission grid infrastructure is associated with a 0.045% increase in solar deployment. Solar potential and grid infrastructure are time-invariant at the county level and are therefore absorbed by county fixed effects; only their interactions with RPS adoption are identified. Together, these findings suggest that states who adopt an RPS tend to see an increase in solar buildout, which is especially concentrated in counties with favorable solar resource endowments and existing grid infrastructure.

Table 2. RPS Adoption and Spatial Allocation of Solar Buildout

RPS Adoption and Spatial Allocation of Solar Buildout

	Model 1	Model 2	Model 3
Dependent variable		IHS Solar MW Added	
RPS adopted (1 = post)	0.109*	-5.973*	-0.082
	(0.064)	(3.158)	(0.065)
RPS × Log solar potential		0.829*	
		(0.435)	
RPS × Log grid infrastructure			0.045**
			(0.018)
Log population	0.525***	0.508***	0.509***
	(0.146)	(0.139)	(0.143)
Log median income	-0.288**	-0.308***	-0.267**
	(0.109)	(0.108)	(0.102)
Observations	73890	73890	73890
R^2	0.222	0.224	0.224
County FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Controls	Pop, Income	Pop, Income, (RPS×Solar)	Pop, Income, (RPS×Grid)
Mean of DV	0.110	0.110	0.110

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Standard errors clustered at the state level.

Notes: Table 2 includes a multivariate regression for Inverse Hyperbolic Sine Utility-Scale Solar MW Added. Utility-Scale Solar data comes from the U.S. Geological Survey from the U.S. Department of the Interior. Utility-Scale Solar is defined as solar facilities with capacity of 1 megawatt or more. The Utility-Scale Solar data is updated as of 2024. The RPS data comes from the Energy Markets & Planning Berkeley Lab. The RPS data is updated as of 2024. The solar potential data comes from the World Bank - Global Solar Atlas. The solar potential data is time invariant. The grid infrastructure data comes from the Esri U.S. Federal. The grid infrastructure data is updated as of 2024. The log population and log median income data come from the United States Census Bureau. The socioeconomic data spans from 2000-2024. The sample includes all U.S. counties from 2000 to 2023. RPS status is state-year. Solar potential and grid infrastructure main effects are absorbed by county fixed effects; only interaction terms are identified.

4.2.1 Post-Estimation

To understand the findings from 4.2, a post-estimation has been conducted. Figure 5 shows that counties with high solar potential experience more than double the RPS-induced increase in solar buildout compared to counties with low solar potential. This suggests that RPS policies tend to drive development where solar resources are strongest, rather than spreading solar evenly across all counties. Regarding grid infrastructure, RPS adoption also leads to larger increases in solar deployment in counties with stronger transmission infrastructure. Counties with more existing grid capacity see a noticeably larger response to RPS policies. These post-estimation results indicate that the effect of RPS adoption on solar buildout is substantially larger in counties with higher solar potential and greater transmission infrastructure, suggesting that renewable mandates tend to concentrate development in locations with more favorable resource and grid conditions.

Figure 5. Estimated Effect of RPS Adoption by County Characteristics

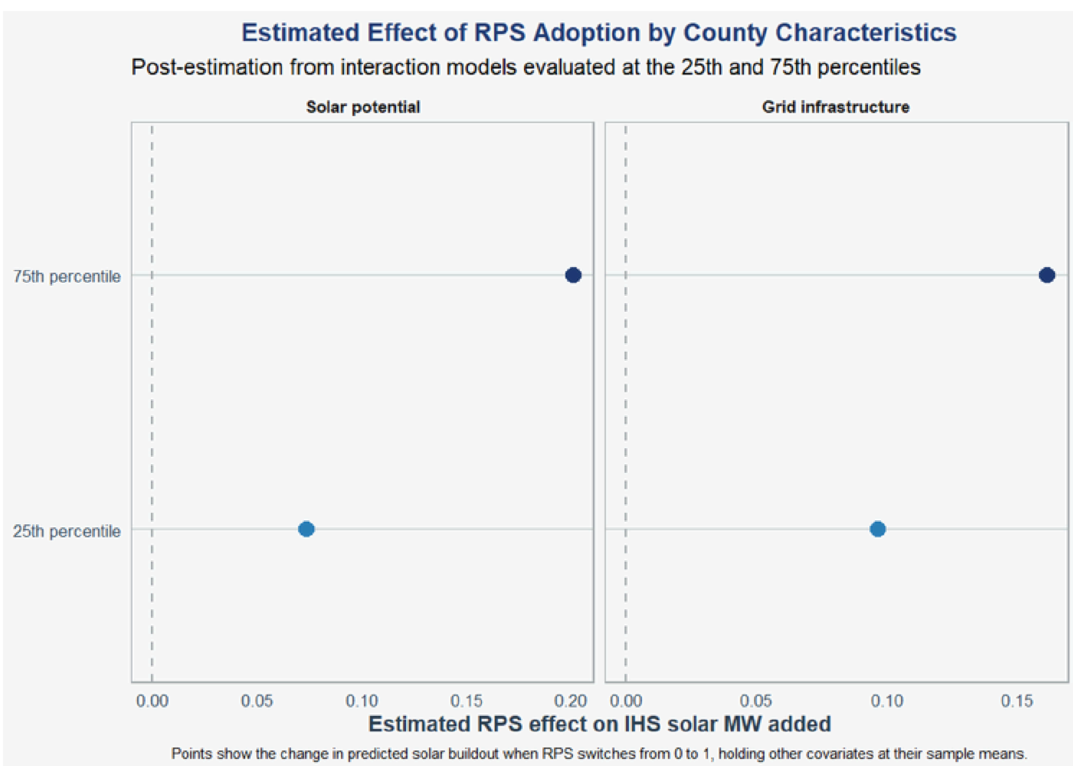

Notes: Figure 5 presents post-estimation marginal effects from the interaction models estimating the effect of state RPS adoption on utility-scale solar buildout. Points show the estimated change in predicted solar buildout when RPS adoption switches from 0 to 1, evaluated at the 25th and 75th percentiles of county solar potential (left panel) and transmission grid infrastructure (right panel), while holding other covariates at their sample means. The dependent variable is the inverse hyperbolic sine of utility-scale solar megawatts added in a county-year. All models include county and year fixed effects and control for log population, log median household income, solar potential, and grid infrastructure, with standard errors clustered at the state level. Utility-scale solar data come from the U.S. Geological Survey (U.S. Department of the Interior) and include photovoltaic facilities with capacity of 1 megawatt or greater; the dataset is updated through 2024. RPS policy data come from the Berkeley Lab Energy Markets and Policy dataset, updated through 2024. Solar potential data come from the Global Solar Atlas (World Bank), and transmission infrastructure data come from Esri U.S. Federal electric power transmission features. The analysis sample includes all U.S. counties from 2000–2023, and RPS treatment status varies at the state-year level.

4.2.2 Mechanism Test

To further investigate the mechanisms underlying the results in 4.2, this study estimates a model using an alternative dependent variable: the number of utility-scale solar projects (full regression table in Appendix - Table 3). This allows for a distinction between extensive margin effects (the siting of new projects) and intensive margin effects (increasing capacity at existing sites). Model 1 finds a positive and statistically significant relationship between RPS adoption and the number of projects. Model 3 similarly finds a positive and statistically significant interaction between RPS adoption and grid infrastructure, suggesting that areas with stronger transmission capacity experience greater increases in project siting. In contrast, while Model 2 estimates a positive interaction between RPS adoption and solar potential, the coefficient is not statistically significant. Overall, these results suggest that RPS policies primarily operate through the extensive margin by increasing the number of solar projects, rather than solely expanding capacity at existing sites. This reinforces the interpretation that RPS policies influence the geographic distribution of solar development by shaping where new projects are located.

4.2.3 Heterogeneity Test

To further examine how RPS policies affect the distribution of solar deployment, a heterogeneity test was conducted (full regression table in Appendix - Table 4). The results show no statistically significant evidence that RPS policies allocate solar development toward counties with greater solar potential, as neither the medium nor high solar interaction terms are statistically significant. While the interaction from 4.2 suggests a positive relationship, the tertile-based specification does not detect statistically significant differences across discrete categories. In contrast, the grid infrastructure heterogeneity test shows positive and statistically significant effects for both medium and high grid counties. These findings suggest that RPS policies may amplify pre-existing geographic advantages, directing new solar development toward areas with better infrastructure rather than promoting uniform deployment across regions.

Overall, the results from the SDiD show that RPS adoption is associated with a statistically significant increase in utility-scale solar deployment at the county level. The mechanism test indicates that this effect operates primarily through the extensive margin, with RPS policies increasing the number of solar projects rather than solely expanding capacity at existing sites. The heterogeneity result further shows that these effects are not evenly distributed, with counties that have stronger grid infrastructure experiencing larger increases in solar buildout. Together, these findings suggest that while RPS policies effectively promote renewable energy development, they tend to concentrate new investment in areas that are already more favorable for solar siting.

4.3 Sun-Abraham Event-Study

To further investigate the results from 4.2, this study uses an Event-Study, shown by Figure 6. The main finding from the Event Study is that RPS adoption is associated with an increase in utility-scale solar buildout.

The event-study estimates show general support for the parallel trends assumption. In Model 1, pre-treatment coefficients are close to zero and do not exhibit systematic trends prior to RPS adoption, while post-treatment effects show a gradual increase in solar buildout, consistent with the time required for project development and construction. In Model 2 of the Event Study, which introduces heterogeneity by solar potential, counties with high solar potential exhibit positive coefficients prior to RPS adoption. However, these differences appear stable over time rather than trending, suggesting baseline level differences rather than violations of the parallel trends assumption. Model 3, which examines heterogeneity by grid infrastructure, shows the pre-period coefficients are close to zero, again suggesting reasonable parallel trends. Overall, while baseline differences exist across counties, there is no clear evidence of differential pre-trends, supporting the validity of the empirical design.

Figure 6 . Baseline Sun-Abraham Event Study: RPS Adoption and Solar Buildout

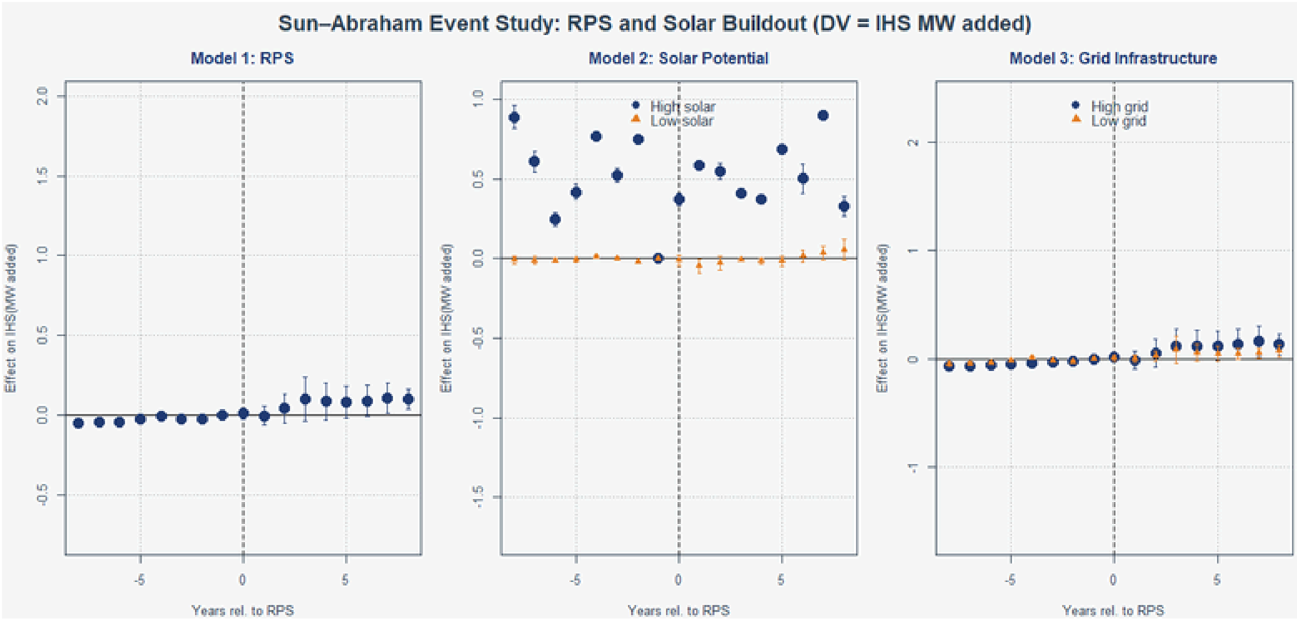

Notes: Figure 6 presents Sun–Abraham event study estimates of the effect of state RPS adoption on utility-scale solar buildout. The dependent variable is the inverse hyperbolic sine of utility-scale solar megawatts added in a county-year. Model 1 shows the baseline dynamic effect of RPS adoption. Model 2 estimates heterogeneous effects by solar potential (high vs. low solar counties), and Model 3 estimates heterogeneous effects by grid infrastructure (high vs. low grid counties). All models include county and year fixed effects, and standard errors are clustered at the state level. Utility-scale solar data come from the U.S. Geological Survey (U.S. Department of the Interior) and include photovoltaic facilities with capacity of 1 megawatt or greater; the dataset is updated through 2024. RPS policy data come from the Berkeley Lab Energy Markets and Policy dataset, updated through 2024. The analysis sample includes all U.S. counties from 2000–2023, and RPS treatment status varies at the state-year level.

5 ROBUSTNESS CHECKS

To test the robustness of the results found in Section 4.2, the following Robustness Checks were employed: (1) using Continuous RPS as a different treatment variable and (2) using Log + 1 MW Added as the Dependent Variable. As seen in Appendix - Table 5, when changing the treatment variable to a Continuous RPS treatment, we see a positive, statistically significant effect across all three models, meaning that RPS adoption tends to increase solar buildout, and this happens more in counties with greater solar potential and grid infrastructure. Hence, the results of this robustness check complement the findings from Section 4.2, as for both we see a positive, statistically significant effect across all three models.

Furthermore, from Appendix - Table 6, we see that RPS adoption tends not to uniformly increase solar buildout across all counties, but rather reallocates deployment to technically advantaged counties (higher solar potential and higher grid infrastructure). Comparatively, Model 1 suggests that an RPS adoption alone does not result in statistically significant more solar buildout. Hence, this robustness check further supports the findings from Section 4.2 as the effect of RPS on Solar Buildout is statistically significant and positive when interacted with both Solar Potential and Grid Infrastructure.

6 CONCLUSION

To conclude, this study assessed the following research question: How do Renewable Portfolio Standards shape where utility-scale solar is deployed across counties, and whether deployment occurs in locations with strong solar resources and grid infrastructure? This study used a monadic panel dataset spanning from 2000 to 2023 and covering all United States counties, to do an EDA, a SDiD, and a Sun-Abraham Event Study. The EDA found that RPS adoption has a positive association with Solar Buildout. The SDiD found (1) RPS adoption is associated with an increase in Utility-Scale Solar Buildout, (2) in states that have adopted an RPS, counties with high solar potential tend to see greater solar buildout, and (3) in states that have adopted an RPS, counties with greater grid infrastructure tend to see greater solar buildout. The Event Study shows (1) no clear evidence of differential pre-trends prior to RPS adoption and a gradual increase in solar buildout beginning after adoption, (2) high solar potential counties tend to see positive coefficients prior to RPS adoption, and (3) high grid infrastructure counties tend to see pre-period coefficients close to zero. To further substantiate the results, a mechanism test, a heterogeneity test, and robustness tests were employed.

Taken together, the findings suggest that RPS policies are associated with renewable energy deployment, but that the geographic distribution of that deployment is strongly shaped by underlying resource and infrastructure conditions. In practice, this means that while RPS policies can stimulate additional solar development, new projects tend to be more likely to occur in counties that already have high solar potential and sufficient transmission infrastructure to support large-scale generation. Additional research could investigate how local land-use regulations, permitting processes, and community opposition influence the spatial distribution of solar development. Even when solar potential is high, regulatory or political barriers at the county or municipal level may limit project siting.

This paper concludes with the following policy recommendation: in order to effectively build out renewable energy equitably, states should adopt targeted policies alongside RPS. In particular, policies that support grid expansion and infrastructure development in underserved regions could help ensure that renewable energy deployment is not limited to areas with existing advantages. Additionally, targeted financial incentives or programs aimed at lower-income or lower-capacity counties could help reduce barriers to participation and promote a more balanced geographic distribution of renewable energy development.

7 REFERENCES

Barbose, Galen. U.S. State Renewables Portfolio and Clean Electricity Standards: 2025 Data Update. Berkeley, CA: Lawrence Berkeley National Laboratory, August 2025. https://emp.lbl.gov/publications/us-state-renewables-portfolio-clean-0

Deschênes, Olivier, et al. 2023. "Causal Effects of Renewable Portfolio Standards on Renewable Investments and Generation: The Role of Heterogeneity and Dynamics." Resource and Energy Economics 75: 101393. https://doi.org/10.1016/j.reseneeco.2023.101393.

Ding, Yueting, et al. 2023. "Spatial Disparity of Utility-Scale Solar Energy and the Role of Solar Policy in the United States." Energy Strategy Reviews 48: 101179. https://doi.org/10.1016/j.esr.2023.101179.

Feldman, R., & Levinson, A. (2023). Renewable Portfolio Standards. The Energy Journal, 44(5), 1-20. https://doi.org/10.5547/01956574.44.4.rfel

Fujita, K.S., Ancona, Z.H., Kramer, L.A., Straka, M., Gautreau, T.E., Garrity, C.P., Robson, D., Diffendorfer, J.E., and Hoen, B., 2023, United States Large-Scale Solar Photovoltaic Database (v3.0, April, 2025): U.S. Geological Survey and Lawrence Berkeley National Laboratory data release, https://doi.org/10.5066/P9IA3TUS.

"Global PV Potential Study." Global Solar Atlas. https://globalsolaratlas.info/global-pv-potential-study

Kroeger, Grace D., and Matthew G. Burgess. 2023. "Electric Utility Plans Are Consistent with Renewable Portfolio Standards and Clean Energy Standards in Most US States." Climatic Change 177 (1): 1–?. https://doi.org/10.1007/s10584-023-03645-7.

O'Shaughnessy, Eric. 2019. State and Local Policy Impacts on the Residential Solar PV Installation Industry. https://docs.nrel.gov/docs/fy19osti/72149.pdf

Shrimali, G., Chan, G., Jenner, S., Groba, F., Indvik, J., 2015. Evaluating renewable portfolio standards for in-state renewable deployment: Accounting for policy heterogeneity. EEEP 4 (2). https://www.iaee.org/eeep/eeepannex/jenner_etal_OnlineAppendix_copyedited.pdf

U.S. Census Bureau, American Community Survey 5-Year Estimates, Table B01003: "Total Population", data.census.gov https://data.census.gov/table/ACSDT5Y2023.B01003 12

U.S. Census Bureau. American Community Survey 5-Year Estimates, Table S1903: "Median Income in the Past 12 Months (in 2024 Inflation-Adjusted Dollars)". data.census.gov. https://data.census.gov/table/ACSST5Y2023.S1903
U.S. Census Bureau. Census 2000 Summary File 3. Washington, DC: U.S. Census Bureau, 2002.

U.S. Climate Resilience Toolkit. U.S. Electric Power Transmission Lines (Archive). Climate Mapping for Resilience and Adaptation (CMRA) Portal. https://resilience.climate.gov/datasets/d4090758322c4d32a4cd002ffaa0aa12_0/about

Wiser, Ryan H., et al. 2011. "Supporting Solar Power in Renewables Portfolio Standards: Experience from the United States." Energy Policy 39, no. 7 (July): 3894–3905. https://doi.org/10.1016/j.enpol.2010.11.025.

Yin, H., Powers, N., 2010. Do state renewable portfolio standards promote in-state renewable generation? Energy Policy 38 (2), 1140–1149. https://doi.org/10.1016/j.enpol.2009.10.067

8 APPENDIX

Appendix - Table 1. Variable Table

Variable Name and Source	*Variable Definition*	*Variable Cleaning and Transformations*
Utility-Scale Solar Buildout U.S. Geological Survey from the U.S. Department of the Interior	Locations and array boundaries of U.S. photovoltaic (PV) facilities with solar capacity of 1 megawatt or more.	County-level solar data were imported and cleaned using standardized state, county, and year identifiers. Observations were restricted to the 2000–2025 period. The data were then aggregated to the county–year level by summing the capacity of all solar installations to construct total capacity added (mw_added, measured in MW AC) and counting the number of projects (n_plants) in each county-year. This variable has been transformed to Inverse Hyperbolic Sine.
RPS Targets Energy Markets & Planning Berkeley Lab	Renewable portfolio standards targets.	State-level RPS data were reshaped to a state–year panel. Because the original values were stored as either numeric proportions or percentage strings, all entries were parsed and standardized to percentage points. Two measures were created: a continuous RPS target (rps_percent) and a binary indicator (rps_binary) equal to 1 when the RPS target is greater than zero and 0 otherwise. The dataset was restricted to the 2000–2025 period.
U.S. Electric Power Transmission Lines. U.S. Electric Power Transmission Lines from Esri U.S. Federal Datasets	A single electric power transmission feature.	County-level grid infrastructure was constructed from transmission line shapefiles. Transmission lines were filtered to include only in-service lines with voltages of 115 kV or greater, then spatially intersected with county boundaries and aggregated to the county level. For each county, total transmission line length (total_line_km), maximum voltage (max_voltage), a binary indicator for grid presence (grid_present), and the log-transformed transmission line length +1 (log_grid_km) were calculated.
Solar Potential (Global Horizontal Irradiation). World Bank - Global Solar Atlas	Sum of direct and diffuse irradiation components received by a horizontal surface and is measured in kilowatt hours per square meter.	County-level solar potential was measured using average Global Horizontal Irradiance (ghi_mean). The data were imported from a spatial dataset, cleaned, and matched to counties using GEOID identifiers. The resulting variable represents the mean solar irradiance for each county and is time-invariant. This variable has been transformed to log.
Median Income (Household) United States Census Bureau - American Community Survey (ACS) and 2000 Decennial Census	The median income of households in the past 12 months, adjusted to inflation-adjusted dollars.	County-level median household income was obtained from the U.S. Census Bureau. For years 2000–2005, median household income was measured using the 2000 Decennial Census. For years 2006–2009, the 2009 ACS 5-year estimate was assigned to each year, and for 2010–2023, annual ACS 5-year estimates were used. For years after the most recent ACS release, the latest available ACS estimate was carried forward. The variable was log-transformed for use in the analysis. In the final sample, a small number of county-equivalent jurisdictions (n = 58; primarily Connecticut counties and independent cities in Virginia and the District of Columbia) were dropped because they could not be consistently matched to Census county GEOIDs required for merging socioeconomic controls.
Population Total United States Census Bureau - American Community Survey (ACS) and 2000 Decennial Census	The total number of residents in a geographic area.	County-level population was obtained from the U.S. Census Bureau. For years 2000–2005, population was measured using the 2000 Decennial Census. For years 2006–2009, the 2009 ACS 5-year estimate was assigned to each year, and for 2010–2023, annual ACS 5-year estimates were used. For years after the most recent ACS release, the latest available ACS estimate was carried forward. The variable was log-transformed for use in the analysis. The dataset was accessed via the Census API through tidycensus in R. In the final sample, a small number of county-equivalent jurisdictions (n = 58; primarily Connecticut counties and independent cities in Virginia and the District of Columbia) were dropped because they could not be consistently matched to Census county GEOIDs required for merging socioeconomic controls.

Notes: *This table describes all variables used in the analysis, including definitions, sources, and data transformations. All variables are measured at the county-year level unless otherwise specified.*

Appendix - Table 2. t-tests of Mean Differences Across Groups

t-tests of Mean Differences Across Groups

Variable	Comparison	Difference in Means	p-value	
IHS Utility-Scale Solar MW Added	Control vs Treated (Pre)	-0.035	0.000***	***
Log Solar Potential (GHI)	Control vs Treated (Pre)	-0.042	0.000***	***
Log Transmission Line Length (km + 1)	Control vs Treated (Pre)	-0.081	0.000***	***
IHS Utility-Scale Solar MW Added	Control vs Treated (Post)	0.174	0.000***	***
Log Solar Potential (GHI)	Control vs Treated (Post)	-0.019	0.000***	***
Log Transmission Line Length (km + 1)	Control vs Treated (Post)	0.103	0.000***	***
IHS Utility-Scale Solar MW Added	Treated (Pre) vs Treated (Post)	0.209	0.000***	***
Log Solar Potential (GHI)	Treated (Pre) vs Treated (Post)	0.023	0.000***	***
Log Transmission Line Length (km + 1)	Treated (Pre) vs Treated (Post)	0.184	0.000***	***

Notes: *Appendix - Table 2 includes t-tests of Mean Differences Across Groups for Utility-Scale Solar added, RPS, solar potential, grid infrastructure, log population, and log median household income for the Control, Treated (Pre), and Treated (Post) groups. Utility-Scale Solar data comes from the U.S. Geological Survey from the U.S. Department of the Interior. Utility-Scale Solar is defined as solar facilities with capacity of 1 megawatt or more. The Utility-Scale Solar data is updated as of 2024. The RPS data comes from the Energy Markets & Planning Berkeley Lab. The RPS data is updated as of 2024. The solar potential data comes from the World Bank - Global Solar Atlas. The solar potential data is time invariant and is defined as the sum of direct and diffuse irradiation components received by a horizontal surface and is measured in kilowatt hours per square meter. The grid infrastructure data comes from the Esri U.S. Federal. The grid infrastructure data is updated as of 2024 and is defined as electric power transmission features. The log population and log median income data come from the United States Census Bureau. The socioeconomic data spans from 2000-2024. The sample includes all U.S. counties from 2000 to 2023. RPS status is state-year. The main finding from this Table are as follows: the IHS MW Added means difference between Control vs Treated (Pre) is very minimal (-0.035), while between Control vs Treated (Post) is substantial (0.174), and lastly between Treated (Pre) vs Treated (Post) is substantial (0.209).*

Appendix - Table 3. Mechanism Test

Mechanism Test: Does RPS Increase the Number of Solar Projects?

	Model 1	Model 2	Model 3
Dependent variable		Number of utility-scale solar projects	
RPS adopted (1 = post)	0.081*		
	(0.043)		
RPS × Log solar potential		0.999	
		(0.627)	
RPS × Log grid infrastructure			0.051**
			(0.019)
Log population	0.309***	0.224**	0.192**
	(0.107)	(0.089)	(0.084)
Log median income	-0.182**	-0.152***	-0.131**
	(0.074)	(0.055)	(0.051)
Observations	73890	73890	74119
R^2	0.274	0.400	0.401
County FE	Yes	Yes	Yes
State-Year FE	No	Yes	Yes
Controls	Pop, Income, Solar, Grid	Pop, Income, Solar + (RPS×Solar)	Pop, Income, Grid + (RPS×Grid)
Mean of DV	0.073	0.073	0.073

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Standard errors clustered at the state level.

Notes: *Appendix - Table 3 presents mechanism results by examining whether RPS adoption increases the number of utility-scale solar projects at the county-year level. Solar project data come from the U.S. Geological Survey within the U.S. Department of the Interior and include solar facilities with capacity of 1 megawatt or greater. The dataset is updated as of 2024. RPS policy data come from Lawrence Berkeley National Laboratory's Energy Markets & Planning group and are updated as of 2024. Solar potential data come from the World Bank's Global Solar Atlas and are time invariant, measuring global horizontal irradiation in kilowatt-hours per square meter. Grid infrastructure data come from the Esri U.S. Federal electric power transmission dataset and are updated as of 2024. Population and median household income data are obtained from the U.S. Census Bureau. Socioeconomic data span 2000–2024. The sample includes all U.S. counties from 2000 to 2023. RPS adoption varies at the state-year level. Standard errors are clustered at the state level. All models include county fixed effects, and Models 2 and 3 additionally include state-by-year fixed effects.*

Appendix - Table 4. Heterogeneity Test

Heterogeneity Test: Do RPS policies allocate solar differently across counties?

	Solar Potential	Grid Infrastructure
Dependent variable	IHS Solar MW Added	
RPS × Medium Solar	-0.005	
	(0.053)	
RPS × High Solar	0.174	
	(0.138)	
RPS × Medium Grid		0.114**
		(0.053)
RPS × High Grid		0.210***
		(0.066)
Log population	0.269**	0.232**
	(0.105)	(0.106)
Log median income	-0.181**	-0.160**
	(0.071)	(0.065)
Observations	73890	74119
R^2	0.314	0.315
County FE	Yes	Yes
State-Year FE	Yes	Yes
Controls	Pop, Income + (RPS×Solar Tertile)	Pop, Income + (RPS×Grid Tertile)
Mean of DV	0.110	0.110

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Solar potential and grid infrastructure are divided into tertiles. Standard errors clustered at the state level.

Notes: *Appendix - Table 4 presents heterogeneity test results examining whether RPS adoption allocates utility-scale solar capacity differently across counties based on solar potential and grid infrastructure. The dependent variable is the inverse hyperbolic sine of utility-scale solar megawatts added at the county-year level. Utility-scale solar data come from the U.S. Geological Survey within the U.S. Department of the Interior and include solar facilities with capacity of 1 megawatt or greater. The dataset is updated as of 2024. RPS policy data come from Lawrence Berkeley National Laboratory's Energy Markets & Planning group and are updated as of 2024. Solar potential data come from the World Bank's Global Solar Atlas and measure global horizontal irradiation in kilowatt-hours per square meter. Grid infrastructure data come from the Esri U.S. Federal electric power transmission dataset. Population and median household income data come from the U.S. Census Bureau and span 2000–2024. The sample includes all U.S. counties from 2000 to 2023. Solar potential and grid infrastructure variables are divided into county-level tertiles. All models include county and state-year fixed effects, and standard errors are clustered at the state level.*

Appendix - Table 5. Robustness Check 1: Continuous RPS

Robustness Check 1: Continuous RPS Stringency and Spatial Allocation of Solar Buildout

	Model 1	Model 2	Model 3
Dependent variable		IHS Solar MW Added	
RPS requirement (10 percentage points)	0.118***		
	(0.034)		
RPS × Log solar potential		0.016***	
		(0.005)	
RPS × Log grid infrastructure			0.031***
			(0.008)
Log population	0.519***	0.517***	0.502***
	(0.140)	(0.139)	(0.134)
Log median income	-0.319**	-0.320**	-0.307**
	(0.120)	(0.120)	(0.119)
Observations	73890	73890	73890
R^2	0.227	0.227	0.231
County FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Controls	Pop, Income, Solar, Grid	Pop, Income, Grid + (RPS×Solar)	Pop, Income, Solar + (RPS×Grid)
Mean of DV	0.110	0.110	0.110

* p < 0.1, ** p < 0.05, *** p < 0.01
Standard errors clustered at the state level.

Notes: *Appendix - Table 5 presents Robustness Check 1, RPS policy is measured as a continuous variable representing the RPS requirement (in 10 percentage point increments), and the dependent variable is the inverse hyperbolic sine of utility-scale solar megawatts added at the county-year level. Utility-scale solar data come from the U.S. Geological Survey within the U.S. Department of the Interior and include solar facilities with capacity of 1 megawatt or greater. RPS policy data come from Lawrence Berkeley National Laboratory's Energy Markets & Planning group. Solar potential data come from the World Bank's Global Solar Atlas and measure global horizontal irradiation in kilowatt-hours per square meter. Grid infrastructure data come from the Esri U.S. Federal electric power transmission dataset. Population and median household income data come from the U.S. Census Bureau and span 2000–2024. The sample includes all U.S. counties from 2000 to 2023. All models include county fixed effects; Models 2 and 3 additionally include state-by-year fixed effects. Standard errors are clustered at the state level.*

Appendix - Table 6. Robustness Check 2: Different Dependent Variable - Log + 1

Robustness Check 2: RPS Adoption and Spatial Allocation of Solar Buildout

	Model 1	Model 2	Model 3
Dependent variable		Log Solar MW Added	
RPS adopted (1 = post)	0.087		
	(0.052)		
RPS × Log solar potential		1.463**	
		(0.560)	
RPS × Log grid infrastructure			0.050***
			(0.017)
Log population	0.441***	0.225**	0.193**
	(0.124)	(0.091)	(0.088)
Log median income	-0.236**	-0.151**	-0.129**
	(0.090)	(0.057)	(0.055)
Observations	73890	73890	74119
R^2	0.219	0.308	0.308
County FE	Yes	Yes	Yes
State-Year FE	No	Yes	Yes
Controls	Pop, Income, Solar, Grid	Pop, Income + (RPS×Solar)	Pop, Income + (RPS×Grid)
Mean of DV	0.090	0.090	0.090

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
Standard errors clustered at the state level.

Notes: *Appendix - Table 6 presents Robustness Check 2, RPS policy is measured using a binary indicator for RPS adoption (1 = post adoption), and the dependent variable is the log of utility-scale solar megawatts added. Utility-scale solar data come from the U.S. Geological Survey within the U.S. Department of the Interior and include solar facilities with capacity of 1 megawatt or greater. RPS policy data come from Lawrence Berkeley National Laboratory's Energy Markets & Planning group. Solar potential data come from the World Bank's Global Solar Atlas and measure global horizontal irradiation in kilowatt-hours per square meter. Grid infrastructure data come from the Esri U.S. Federal electric power transmission dataset. Population and median household income data come from the U.S. Census Bureau and span 2000–2024. The sample includes all U.S. counties from 2000 to 2023. All models include county fixed effects; Models 2 and 3 additionally include state-by-year fixed effects. Standard errors are clustered at the state level.*

Appendix - Figure 1. Variables Before and After Transformation

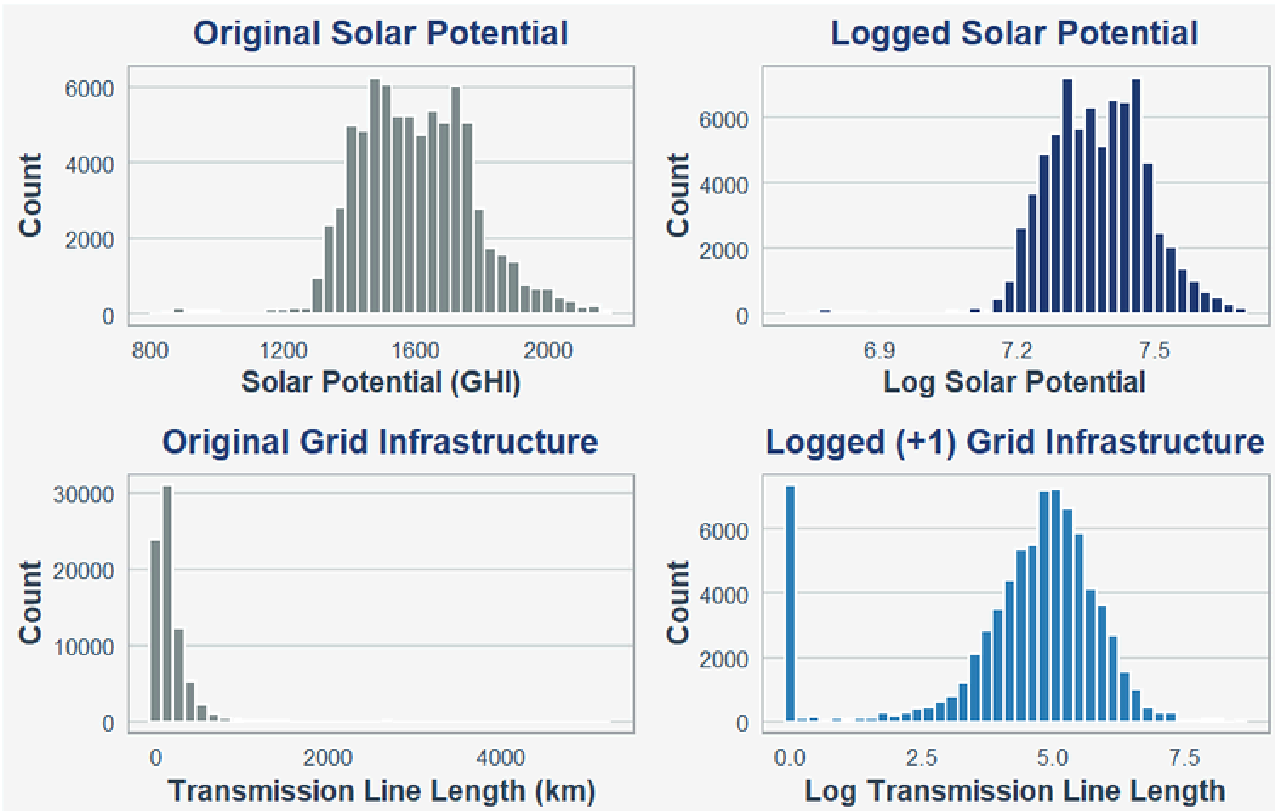

Notes: *Appendix - Figure 1 shows the distribution of solar potential and grid infrastructure before and after transformation. Solar potential is measured using Global Horizontal Irradiance (GHI) from the World Bank Global Solar Atlas and represents the sum of direct and diffuse solar irradiation received by a horizontal surface (kWh/m²). Grid infrastructure is measured as the total length of electric power transmission lines (km) from Esri U.S. Federal electric power transmission features. The sample includes all U.S. counties from 2000–2023.*

Solar potential is strictly positive and is therefore log-transformed directly, while grid infrastructure includes zero values and is transformed using a log(x + 1) specification to retain observations with zero transmission line length. The log transformations reduce skewness and produce more approximately symmetric distributions. Because these variables do not contain large negative values and are not the primary dependent variable of interest, standard log transformations are appropriate and provide a straightforward interpretation of coefficients as semi-elasticities. The inverse hyperbolic sine (IHS) transformation is typically used for dependent variables that include zero values and require elasticity-like interpretation. Since these variables are used as explanatory covariates rather than outcomes, the log and log(x + 1) transformations provide a simpler and more interpretable specification.

Appendix - Figure 2. RPS Adoption by State

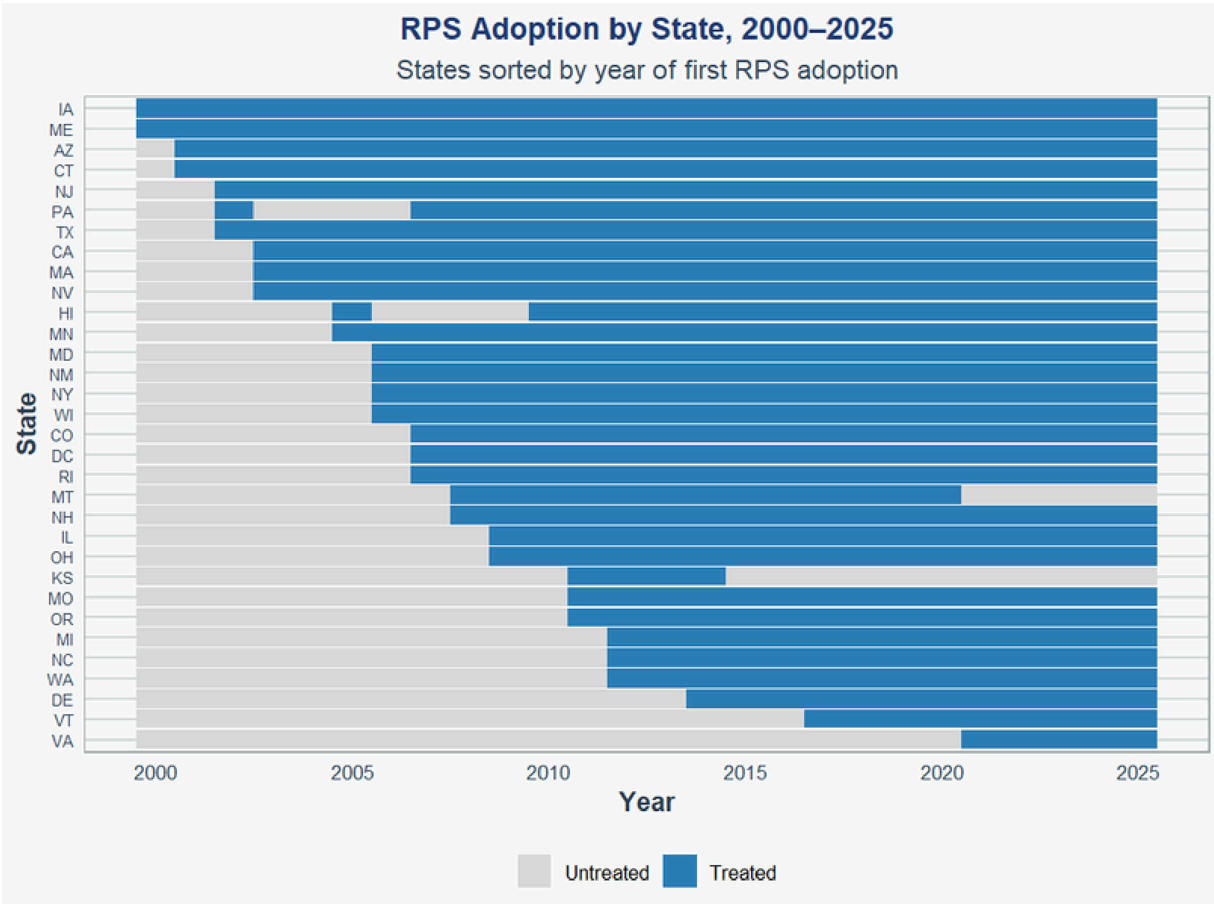

Notes: *Appendix - Figure 2 includes RPS adoption by state from 2000 to 2025. The RPS data comes from the Energy Markets & Planning Berkeley Lab. The RPS data is updated as of 2024. The figure shows if the state does or does not have an RPS in that year. As seen, two states Iowa and Maine are always adopters, meaning that they adopted an RPS prior to 2000.*

Appendix - Figure 2 shows the states that eventually adopted an RPS in the years 2000-2025, all other states not included in the figure never adopted an RPS. This figure shows that only two states had adopted an RPS prior to 2000. Furthermore, this figure shows the 30 states used in the SDiD model as treated units (all RPS states minus the two always treated states).