



Fall 2025 Issue, Volume V

Ohio SB52 and Renewable Energy

Tourism and Labor Markets

Trade Centrality Index

Impact of Medicaid Expansion on Crime Rates



CASE
WESTERN
RESERVE
UNIVERSITY

journalofeconomics@case.edu
cwrujoe.org
[linkedin.com/company/cwru-joe](https://www.linkedin.com/company/cwru-joe)

CWRU Economics Department
Weatherhead School of Management at
Case Western Reserve University
Peter B. Lewis Building
11119 Bellflower Road
Cleveland, OH 44106

Table of Contents

Staff	iii
Letter from the Editors	iv
Research and Editorial Division	1
Tourism and Labor Markets: Insights from Utah’s Parks <i>Anika Kaur, Noah Leibowitz</i>	3
Brain Drain in Nigeria and the United Kingdom’s Migration Paradox <i>Michelle Gantumur, Shamiso Ruwende, Yingfei Xiong</i>	8
Rural-Urban Divergence in the United States: County-Level Evidence from 2014-2024 <i>Ethan Przytulski, Calais Michaelsson</i>	12
Student Debt and Early-Career Earnings <i>Emerson Au, Justin Sweeney, Harish Rajagopalan</i>	18
Who Steps Forward? Demographic Patterns in Living vs. Deceased Kidney Donation <i>Siya Motwani, Lidia Rodriguez</i>	22
Housing Market Responses to New Transit: An Exploratory Analysis of Silver Line Phase II in Northern Virginia <i>Katelyn Kim, Elizabeth Odife</i>	26
A Descriptive Analysis of the Post-COVID-19 Labor Market for Recent Graduates <i>Alexander Klautky</i>	34
Regional Analysis Division	45
The Cost of Saying No: Ohio SB 52 and Renewable Energy Development <i>Zoe Epperson, Elizabeth Krotine</i>	47
Housing Pressure Points: Trends in Homelessness Across Ohio’s Urban Counties <i>Yeva Butovska, David Smith</i>	52
What are the Effects of Game-Day Rainfall on Major League Baseball Home Game Attendance? <i>Erin Butler, Matteo Aron, Paulo Aguiar</i>	55
Index Team	58
Trade Centrality Index <i>Fiona Arado, Bo Cao, Daniel Chase, Isma’il Seddon</i>	61

External Submissions	64
Tax Portfolio Composition and Revenue Volatility: A Comprehensive Analysis of State Tax Systems, 2005-2024 <i>Ryan Hydrick, Peter Constant</i>	67
An Analysis of U.S. GDP and its Effect on the EUR/USD Exchange Rate <i>Ethan Juhasz</i>	77

Editorial Board

Joanna Chiu, *President and Editor-in-Chief*
Ammar Sulemanjee, *Vice President of the Research and Editorial Division*
Elvin Stowell, *Vice President of the Regional Analysis Division*
Kyan Howe, *Vice President of the Outreach Division*
Sidharth Jindal, *Index Lead Investigator*
Alexander Giordano, *Section Lead*
Athena Mandal, *Section Lead*
Allison Su, *Section Lead*
Dr. Sining Wang, *Faculty Advisor*

Index Team

Fiona Arado
Daniel Chase

Bo Cao
Isma'il Seddon

Outreach Division

Isabella Lopez
Naomi Moneme

Steven Mendez
Aanchal Nair

Writers and Researchers

Emerson Au
Anika Kaur
Alexander Klautky
Calais Michaelsson
Elizabeth Odife
Harish Rajagopalan
Shamiso Ruwende
Yingfei Xiong

Michelle Gantumur
Katelyn Kim
Noah Leibowitz
Siya Motwani
Ethan Przytulski
Lidia Rodriguez
Justin Sweeney

Regional Analysis Writers

Paulo Aguiar
Erin Butler
Zoe Epperson
David Smith

Matteo Aron
Eve Butovska
Elizabeth Krotine

Peer Review Board

Sahasra Dendi
David Wu

Damaris Makuh

Design Team

Athena Mandal
Téa Tamburo

Michelle Gantumur
Yingfei Xiong

Alumni Advisors

Brooke Hathhorn
Vaishnavi Kumar
Katherine Merritt

Dear CWRU Journal of Economics Reader,

On behalf of the Editorial Board, I would like to thank you for taking the time to read the fifth edition of our Journal. Now, at the start of its third year, the Journal continues to grow from its original mission: to introduce students to the world of economic research. With a team of 40 passionate and committed students, the Journal has advanced in rigor and quality in ways unparalleled in past years.

Economics research plays a vital role in understanding the world we live in and guiding effective decision-making. Whether evaluating labor market trends, assessing public policies, or analyzing financial systems, economic research equips students with the tools to interpret data, test hypotheses, and uncover correlations that shape real outcomes. At a time when misinformation and uncertainty are increasingly common, economics research offers clarity grounded in evidence and analysis. The papers in this volume reflect that purpose, demonstrating how undergraduate scholarship can meaningfully contribute to broader conversations in economics.

Over the past year, we have witnessed remarkable growth within our teams. Members have strengthened their analytical abilities, sharpened their writing and quantitative skills, and developed a deeper appreciation for the discipline of economics. Many students who joined with little research experience now confidently engage with economic literature. Teams trained in econometrics are eager to apply their coursework to their research projects, consistently pushing to advance the research methods of their work. This internal development has been one of our greatest accomplishments, as the Journal has become not only a publication platform but also a training ground for emerging scholars.

This semester, we were also especially excited to debut the End-of-Year Journal of Economics Conference, where members of the Journal and external submission authors had the opportunity to share their semester-long work with the Case Western Reserve University community. This marks the beginning of an exciting journey. As the first research conference hosted by a Weatherhead School of Management club, we are proud to help expand access to undergraduate economics research and to foster a culture of curiosity across campus. We hope that this conference will inspire students to recognize the feasibility, accessibility, and value of conducting economics research at Case Western Reserve University.

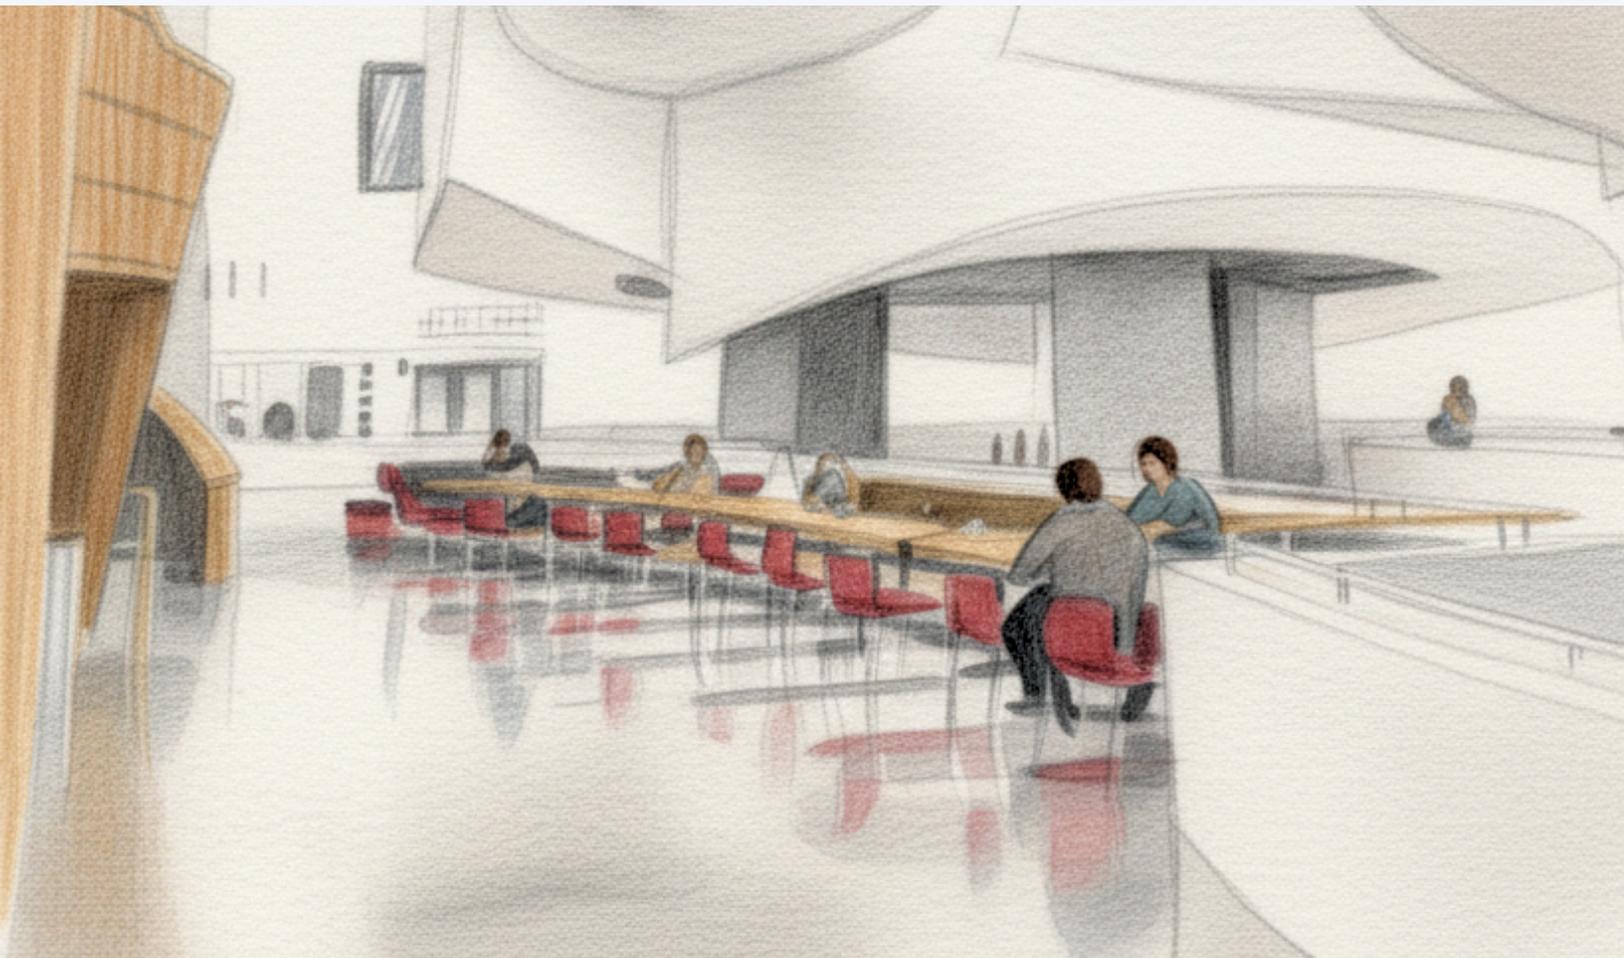
It is with pride and gratitude that we present the fifth volume of the CWRU Journal of Economics. We are grateful for the opportunity to support students in their pursuit of economics research and are eager to see our community continue to grow.

Sincerely,

Joanna Chiu
President & Editor-in-Chief



Research and Editorial



Tourism and Labor Markets: Insights from Utah's Parks

Anika Kaur, Noah Leibowitz

In February 2025, President Donald Trump implemented a series of government layoffs, resulting in the termination of approximately 1,000 National Park Service employees. Many states, particularly Utah, rely on national parks to support local economic activity, raising concerns about potential negative economic consequences. This paper examines the economic spillover effects on local economies by analyzing unemployment rates and park visitation from January 2023 to September 2025. The primary goal is to assess whether these layoffs negatively impacted employment in Utah counties near national parks. We employed exploratory data analysis (EDA) to examine trends in both unemployment and park visitation over time. Results indicate that while park visitation remained largely stable, following typical seasonal patterns, county-level unemployment increased after the layoffs, suggesting a measurable negative impact on local employment.

Introduction

TOURISM plays a vital role in local economies across the United States. In Utah, a large portion of tourism revenue is generated from national park visitation. In February 2025, the Trump Administration imposed a hiring freeze and staffing reductions within the National Parks Service (Sottile & Hurley, 2025). There has been a growing concern that the reduction in employment within the parks would reduce the park's quality and maintenance which would in turn lower visitation. A reduction in visitation can pose a wider risk to local economies that are highly dependent on park-related tourism.

Previous research has shown that through the multiplier effect, national park visitation has a measurable impact on local communities. Increased visitation will stimulate economic activity, particularly in the hospitality sector, which further supports job creation and GDP growth. In a case study from 1970-2017 conducted on the economic impacts of US national parks, a multiplier effect was observed, though it took around a year to see a roughly 4% increase in local employment as a result of increased visitation (Szabó & Ujhelyi, 2022). It was discovered that when park visitation rose, employment rose by a larger amount around a year later. This complements another study in Zambia which found that national park tourism accounted for 40% of local income and around half of GDP growth (Chidakel, Child, & Muyengwa, 2021). Together, these studies emphasize the dependence of local communities on national park tourism.

However, much of the existing research examines total national park visitation over a long time horizon. Our research will focus on short-run changes, from January 2023 to August 2025, to park visitation and unemployment levels for a specific region of parks. Each park's local community may respond differently due to changes in visitation. We chose five national parks in Utah (Zion, Bryce, Capitol Reef, Canyonlands, and Arches) as Utah's local economies are largely dependent on park visitation. Additionally, by choosing parks in a similar region, we can control for seasonal shifts in visitation that may influence the indirect changes to local unemployment. Furthermore, our research will open insights to how local labor markets in Utah counties respond to changes in park visitation. Ultimately, this would allow for future research on the effects of the policy within a collection of national parks.

By examining the county-level unemployment rates, we can determine patterns of relative monthly changes in visitation. Therefore, we hypothesize that relative visitation counts are negatively correlated with unemployment rates.

Literature Review

Much of the existing literature on the economic impact of national parks has focused on the long-run multiplier effects of park visitation and spending on surrounding local economies. This is because the indirect effects of a shock can take time for an impact to be noticeable. National parks stimulate regional economic growth, and changes in park visitation produce spillover effects that can take up to a year to be noticeable (Szabó & Ujhelyi, 2022). Similarly, visitor spending generates substantial income and employment gains (Chidakel et al., 2021). However, previous studies used a survey-based methodology to emphasize microeconomic effects, instead of broad economic changes. This demonstrates that national parks are important drivers of local economic activity. However, most studies focus on long-term effects rather than the short-term shocks, which could be associated with mass layoffs.

The theoretical contributions from (Domański & Gwosdz, 2010) help explain how geographic scale and economic leakages influence the strength of multiplier effects, providing a useful framework for variation in counties in Utah's park-dependent economies. Moreover, evidence from studies on public-sector layoffs, shows that localized job losses can have wider indirect consequences on local labor markets (Gathmann, Helm, & Schönberg, 2014). Using a similar

approach, our study seeks to identify short-run (monthly) spillover effects following the February 2025 national park layoffs. In doing so, we expand on the existing literature by shifting the focus from long-run trends and visitor spending to the immediate economic repercussions of local employment shocks.

Data

We collected monthly recreational visitation counts for the five national parks from the National Park Service (NPS) Visitor Use Statistics Database. County-level monthly unemployment data were collected from the Utah Department of Workforce Services. We gathered labor force and unemployment counts to calculate the unemployment rate, expressed as a percentage, in the primary county that each park is located in. Zion National Park is primarily located in Washington County, Bryce is primarily located in Garfield County, Capitol Reef is primarily located in Wayne County, Canyonlands is primarily located in San Juan County, and Arches is primarily located in Grand County. Each park’s visitation counts were aggregated with the county-level unemployment data to capture local economic conditions. Observations were collected from January 2023 through August 2025 to capture short-run fluctuations while limiting the influence of external factors such as the COVID-19 pandemic, which had significantly affected park tourism in earlier years.

While both datasets are publicly available, they don’t capture all of the variables that can affect the relationship between visitation and unemployment. Such variables include seasonal patterns and natural events. Seasonal patterns were controlled for by choosing parks that are located in Utah and thus have a similar climate. However, natural events are localized and therefore are harder to control in each park. Despite these limitations, our dataset allows for time series analysis of unemployment and visitation patterns.

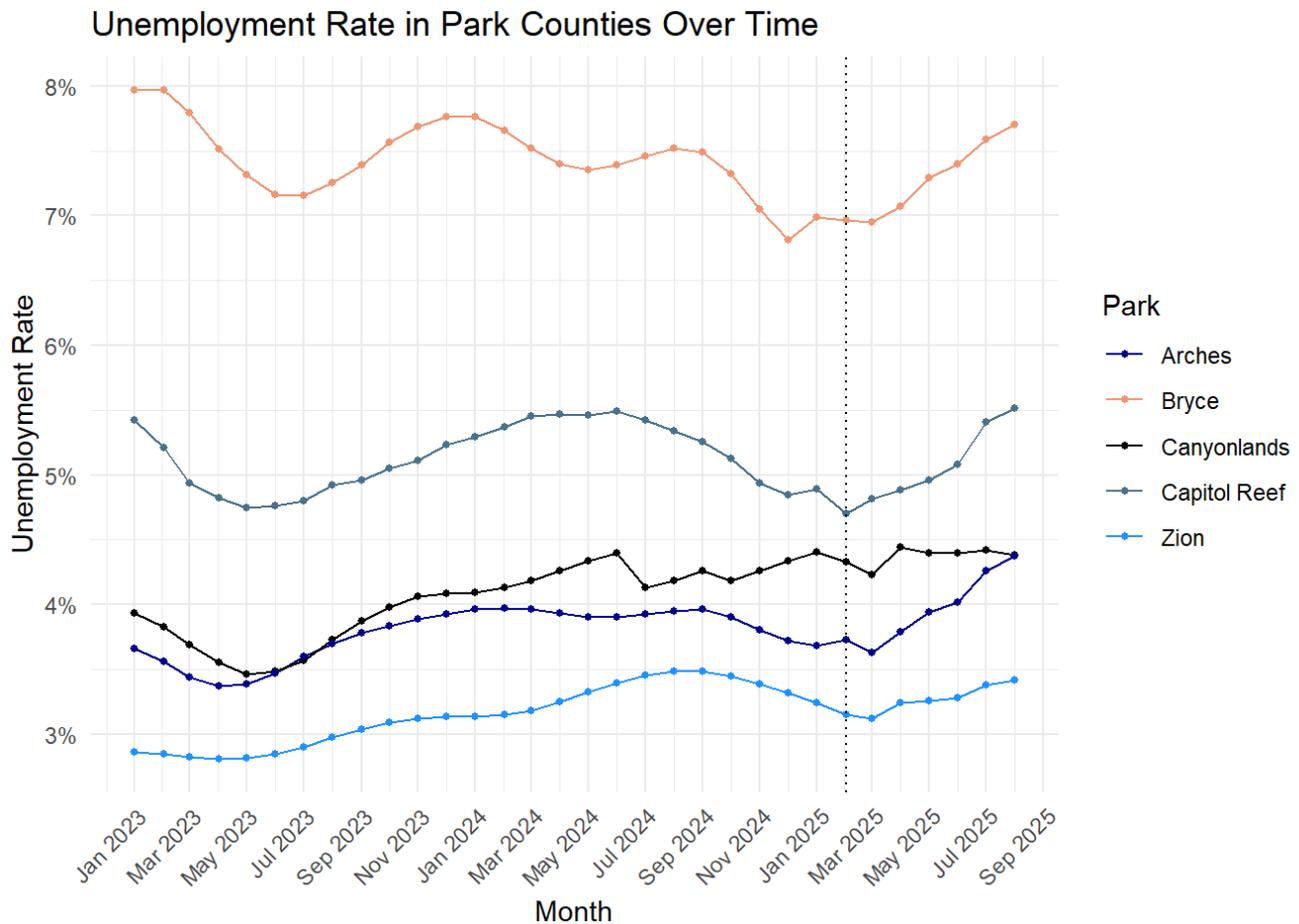


Figure 1: Unemployment rates from January 2023 to August 2025 for the primary county of each park (Utah Department of Workforce Services)

Research Method

To evaluate the relationship between national park visitation and local unemployment, we utilized time series analysis. Specifically, we examined the month-to-month percent change in visitation and monthly county unemployment rates for each respective park and county. We conducted a time series analysis of the percent change in visitation and unemployment rates to observe any seasonal shifts in visitation and unemployment trends. We then constructed a scatter plot for all five parks to visualize any relationship between the variables.

Our analysis is not causal, as there are several factors that influence the relationship between park visitation and unemployment, such as seasonal cycles, natural disasters, and broader macroeconomic fluctuations in Utah. For a more refined analysis, all of these variables should be considered to determine a causal relationship. Despite these limitations, our analysis can be used for future causal studies on the regional effects of park visitation and tourism.

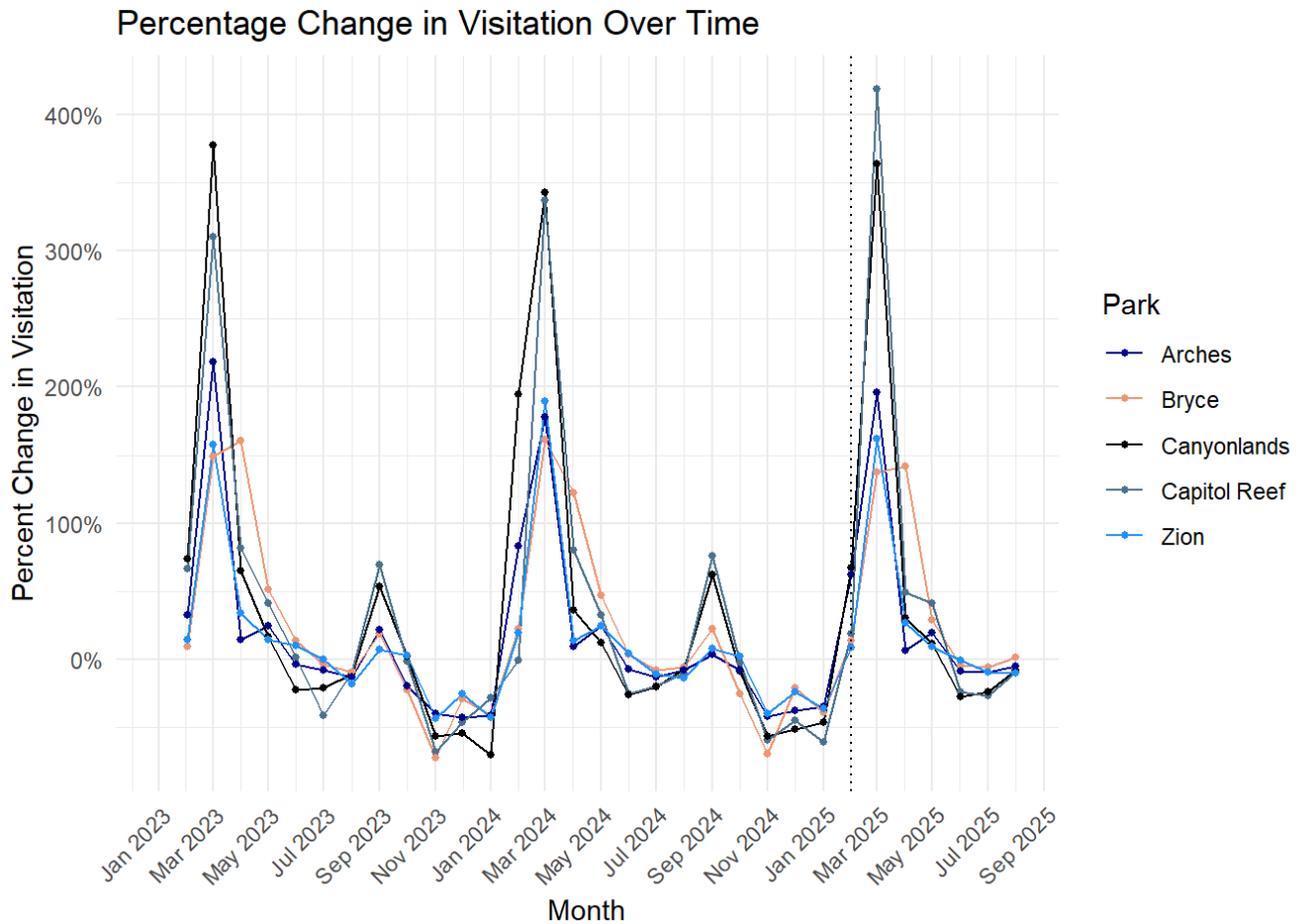


Figure 2: Month-to-month change in visitation from January 2023 to August 2025 for the five parks (National Park Service Visitor Use Statistics Database - Monthly Public Use Report)

Analysis

To visualize trends in county-level unemployment, we conducted a time series analysis, as depicted in Figure 1. The dashed vertical line in February 2025 shows when the mass layoffs of park staff were enacted by the Trump Administration. The unemployment rates for each county do show an upward trend leading up to February 2025. However, after the policy was put in place, Bryce, Capitol Reef, Arches, and Zion saw a steeper rise in unemployment. One possible explanation for the increase in unemployment is that these counties rely heavily on the hospitality and tourism sectors, which may be sensitive to seasonal shifts or current macroeconomic trends. Furthermore, external shocks in other industries can contribute to a rise in overall unemployment. While this trend coincides with the layoffs, our analysis does not establish causality. Regardless, unemployment has risen across Utah counties since the layoffs occurred.

We then analyzed monthly visitation patterns to assess if the policy had an impact on park visitation, as shown in Figure 2. Specifically, we examined whether lower employment in parks decreased the quality of the parks and lowered visitation. Our analysis shows a clear cyclical pattern across all parks. Visitation tends to rise and fall together and peaks in the springtime. Interestingly, the percent change in visitation for each park was the same or even higher in 2025 after the policy went into effect. This shows that using visitation as a proxy for unemployment within the parks was not a good predictor of the total county-level unemployment rate. Additionally, despite the staffing cuts, visitors still attended the parks. This finding, along with the rise in unemployment, suggests that other macroeconomic factors may be playing a stronger role than park-specific employment dynamics.

Overall, our exploratory analysis suggests that while unemployment rose following the February 2025 layoffs, visitation patterns remained largely unchanged, thus leading to a rejection of our hypothesis. Based on the cyclical nature of visitation counts, it is unlikely that any correlation between unemployment and visitation would be seen. The cyclical pattern likely reflects visitors' insensitivity to short-term staffing changes, as they continue traveling for recreation despite reduced park services. Our results point to potential spillover effects leading to a rise in local unemployment in Utah counties, but do not provide evidence of a direct or indirect causal link between the staffing cuts within the parks and a contraction of local labor markets.

Shortcomings

While our study provides insight into the short-run effects of national park visitation on tourism and local economic activity in Utah, there are several limitations. First, the February 2025 policy is recent, and a noticeable economic effect can take over a year to show (Szabó & Ujhely, 2022). Our use of county-level unemployment made our analysis more difficult. Future research should focus on unemployment changes in the hospitality sector. While we couldn't find monthly hospitality unemployment data, such an analysis would allow for an analysis of the impact tourism has on the communities surrounding each park.

Additionally, our analysis only focuses on one region. Future studies can examine a wider range of parks to determine if certain parks were affected differently than others, while controlling for variables such as seasonal variation and climatic differences. Furthermore, incorporating several macroeconomic variables and controls can allow for a more refined model to examine the direct and indirect effects that the policy may have on the parks in the future.

Finally, relying on recreational visitation counts as a proxy for staffing counts in the parks, due to limited public availability in data provided by the NPS, limits the ability to link the direct and indirect effects. Further studies should examine any relationships between staffing counts, visitation, and park quality through measures such as visitor satisfaction.

Conclusion

This study highlights the importance of short-term dynamics of tourism on local labor markets. While our analysis does not establish causality, we found that shifts in unemployment, whether resulting from the 2025 national parks staffing cuts or broader macroeconomic activity, are more noticeable than a reduction in tourism. These findings offer important insights to park administrators and local policymakers to manage the potential negative effects the policy can have on the parks in the future.

References

- Chidakel, A., Child, B., & Muyengwa, S. (2021). Evaluating the economics of park-tourism from the ground-up: Leakage, multiplier effects, and the enabling environment at south luangwa national park, zambia. *Ecological Economics*, 182, 106960. doi: 10.1016/j.ecolecon.2021.106960
- Domański, B., & Gwosdz, K. (2010). *Multiplier effects in local and regional development*. Manuscript, Institute of Geography and Spatial Management, Jagiellonian University, Cracow, Poland. (Manuscript received May 12, revised June 7, 2010; Accessed: 2025-11-13)
- Gathmann, C., Helm, I., & Schönberg, U. (2014). Spillover effects in local labor markets: Evidence from mass layoffs. In *Beiträge zur jahrestagung des vereins für socialpolitik 2014: Evidenzbasierte wirtschaftspolitik - session: Local labour markets, no. a13-v4*. (Provided in cooperation with ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel und Hamburg; Accessed: 2025-11-13)

- Sottile, Z., & Hurley, G. (2025, February 24). Fired national park rangers say services will suffer after 1,000 jobs eliminated. *CNN*. Retrieved from <https://www.cnn.com/2025/02/24/us/national-park-rangers-layoffs> (Accessed: 2025-11-13)
- Szabó, A., & Ujhelyi, G. (2022). Economic impacts of the us national park system. *University of Houston Working Paper Series*. Retrieved from <https://uh.edu/~aszabo2/nps11.pdf> (Accessed: 2025-11-13)

Brain Drain in Nigeria and the United Kingdom's Migration Paradox

Michelle Gantumur, Shamiso Ruwende, Yingfei Xiong

Introduction

BRAIN drain—the departure of educated professionals from one country to another—has become a topic of interest in politics. Human capital flight and the associated brain drain losses in Sub-Saharan Africa are immense, with Nigeria in particular witnessing pronounced outflows of highly skilled professionals to developed countries like the UK due to escalating “push factors” such as deteriorating economic conditions and institutional fragility, and “pull factors” such as wage differentials and level of opportunity (Adelakun et al., (2025); Osayemwenwre et al., (2023)). In 2022, the physician-to-patient ratio in Nigeria was 0.4 to 1,000, while in the United Kingdom, the ratio was 3.2 to 1,000. The departure of skilled professionals has debilitated the country’s healthcare sector and diminished the quality and output of Nigerian institutions, resulting in an overall reduced productive capacity that hinders the country’s economic growth (Osayemwenwre et al., (2023)).

The UK’s attempts to curb total migration have consistently yielded outcomes indicative of the migration paradox, also known as the balloon effect, in which restricting one form of migration leads to growth in others. As a result, the composition of migration changes without reducing overall levels (Rienzo and Vargas-Silva, (2015)).

Following the UK’s withdrawal from the European Union in 2020, which ended free movement, and the introduction of the Points-Based Migration system in 2021, migration levels rose significantly, with the predominant share of these flows comprising highly skilled non-EU migrants. This was driven by increased international recruitment and labor market demand for non-EU skilled labor to fill the shortage gaps in key sectors previously occupied by EU workers (EPRS, (2025)). The skills and qualification-based framework of the Points-Based System reflects dual priorities. The system made entry harder, particularly for EU migrants who no longer received preferential treatment. However, it aimed to attract skilled migrants and optimize their contributions to the economy by expanding visa routes and lowering the minimum salary and education thresholds, making skilled non-EU migration more accessible (EPRS, (2025)).

This study aims to examine the implications of the UK’s post-Brexit migration restrictions on brain drain trends and how they may exacerbate the depletion of human capital and the disparities in its distribution in developing countries such as Nigeria. Unlike other studies that assess the overall impact on the UK’s migration patterns, this analysis focuses on an individual sending country and evaluates specific outcome variables associated with brain drain. The UK government published a migration White Paper in May 2025 outlining forthcoming measures to control and limit migration, which involve increased qualification requirements. As past trends indicate, the proportion of skilled migrants is likely to increase. The findings of our analysis aid in understanding and predicting the trajectory of Nigeria’s economic development under these complex conditions.

Literature Review

Akafa et al. (2023) conducted a cross-sectional survey study investigating the push and pull factors involved in migration from Nigeria to a destination country. 400 participants, sampled by convenience and snowball sampling from the 6th Edition Cardiovascular Symposium, received a self-administered questionnaire that included multiple-choice questions about push and pull factors that drive health professionals away from Nigeria. The major push factors are low remuneration (71.2%), insecurity (62.7%), difficult working environments (55.9%), limited career opportunities (53.6%), lack of equipment and technology for health services (52.5%), and unfavorable government policy (52.5%). The major pull factors were high remuneration potential (76.6%), opportunity of career growth (70.8%), educational opportunities for self (59.3%), educational opportunities for children (55.9%), and high-level equipment for technology (54.9%). 99.3% of respondents also answered “Yes” for “Have you ever considered emigrating out of Nigeria to live and/or work elsewhere?” The United Kingdom was also the most favored destination country, where 50.5% participants chose on multiple choice for countries of consideration for emigration (Akafa, Okeke, and Ore, (2023)).

Rienzo and Vargas-Silva (2015) evaluated how the UK’s pre-Brexit net-migration target measures, which imposed restrictions on migration from outside the European Economic Area (EEA), affected the size and composition of skilled-worker migration flows between 2007 and 2013. The study used Labour Force Survey data to analyze changes among recent migrant workers (RMW). Following the introduction of these restrictions, Non-EEA RMW fell from 154,000 in 2011 to 94,000 in 2013, while EEA skilled RMW rose from 51,000 to 78,000. Although overall migrant

numbers declined, the proportion of highly educated migrants rose by 50-60% over this period. While this study was conducted before Brexit, it provides a useful framework for predicting and analyzing the effects of post-Brexit policies by identifying the fundamental dynamics that influence migration patterns (Rienzo and Vargas-Silva, (2015)).

Research Method

Our method tracks select variables linked to brain drain over a six-year time window to analyze how the implementation of the Points-Based Migration System relates to skilled migration trends and the economic conditions faced by sending countries such as Nigeria. The results from this analysis do not imply a causal relationship. The objective of this method is to identify whether the critical trends that could inform potential causal analysis can be observed in the data.

Data and Analysis

Our analysis incorporates three data sets, the first of which is physician counts sourced from the World Health Organization’s Health Workforce Statistics Database. Our second measure of brain drain was derived from the Fragile State Index’s Brain Brain Index to capture fluctuations and the extent of human capital flight. The third measure of net migration is based on data from the World Bank. Together, these three measures illustrate how post-Brexit policies may induce further skilled migration and, consequently, worsen the fragility of already strained sectors such as healthcare.

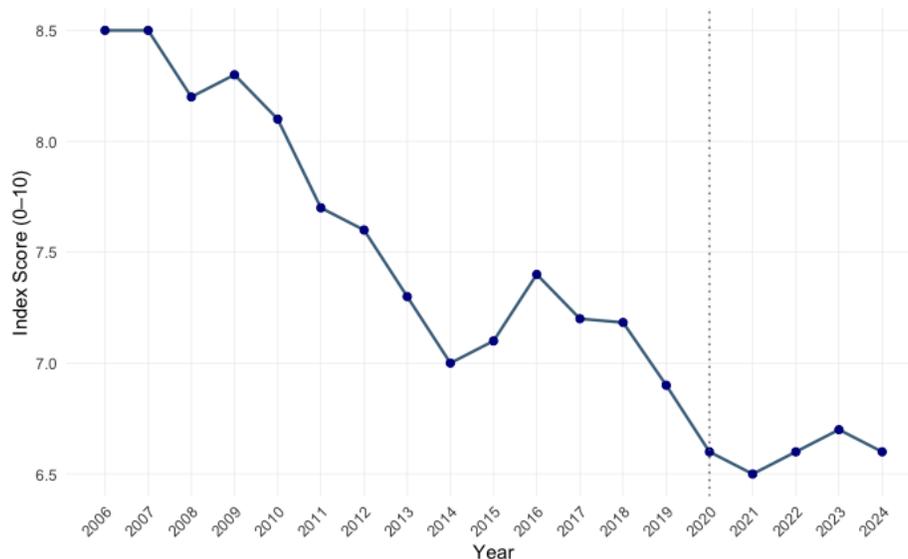


Figure 1: Nigeria – Human Flight (Brain Drain) Index Over Time

As shown in Figure 1, the brain drain index indicates that the period following the UK’s exit from the EU and the introduction of related migration reforms did not correspond with any significant rise in human capital flight as initially hypothesized, reaching its lowest point in 2021, with only slight fluctuations thereafter. The index does not differentiate between-skilled migration categories, encompassing various degrees, qualifications, and experience levels. While high levels of attrition among doctors is a significant issue, these healthcare-sector-specific trends are not reflected in these aggregate brain drain indicators since skilled migration appears to be stable and decreasing overall.

As shown in Figure 2, the increase in the number of doctors post-2020, following the dramatic decline observed from 2016 to 2018, contradicts the expected reduction associated with the heightened demand for non-EU labor and international recruitment post-Brexit.

There is no significant correlation between the change in doctors and net migration, as shown in Figure 3. A potential explanation for this could be the fact that the UK eased restrictions for certain skilled professionals and shortage occupations by expanding visa routes and reducing entry requirements. This could act as an external pull factor that the targeted sectors of the sending country respond to, rather than their entire labor force. For this reason, it may be argued that the sector-specific movement of doctors cannot be detected in general migration data, let alone net migration, which reflects all types of migration.

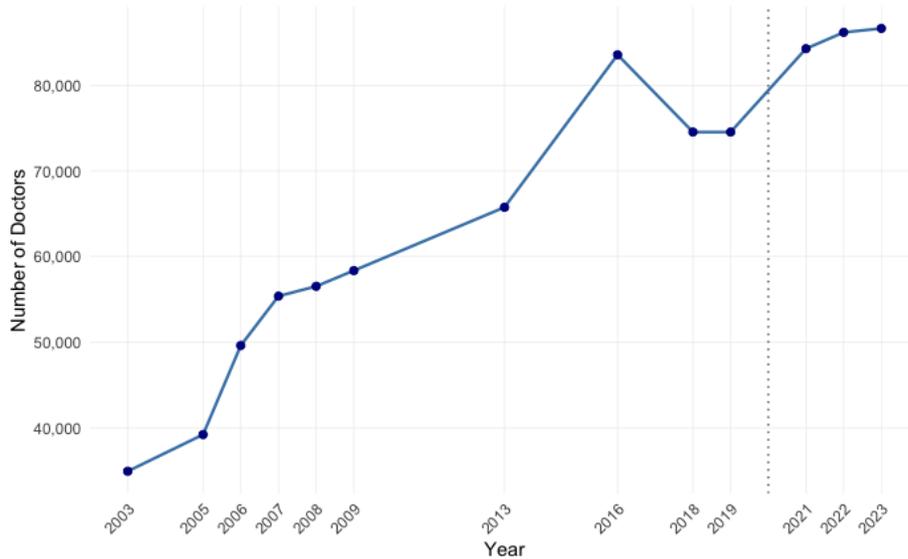


Figure 2: Nigeria – Healthcare Doctors Over Time

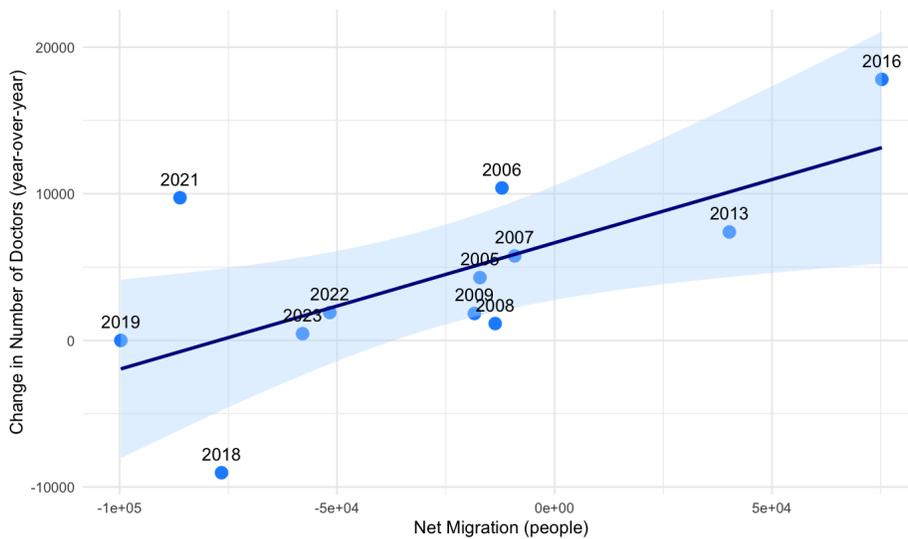


Figure 3: Change in Number of Doctors and Net Migration in Nigeria

Overall, our analysis suggests that the disproportionate effects on migration prompted by the Points-Based Migration System and the corresponding implications for the healthcare sector and skilled workforce cannot be clearly identified in Nigeria’s country-level data.

Shortcomings

One limitation of our study is that it does not demonstrate causation, and any observed relationships may be attributed to a multitude of factors. Additionally, there was limited availability of migrant flow data, and our study relied on net migration data as a secondary alternative. Because Nigeria is a developing nation, significant gaps existed in the data, with certain years unaccounted for. The national-level data sources commonly used in studies of brain drain proved inadequate for illustrating the correlations and trends associated with the specific policy our study sought to identify.

Conclusion

While the number of non-EU skilled migrants has surpassed pre-Brexit levels as a result of the Points-Based Migration System, this broader trend does not appear to have had a notable effect on the number of doctors and brain drain in Nigeria. The inconclusive analysis results suggest that the impact of brain drain, particularly physician brain drain, are insignificant and does not place severe pressures on Nigeria's economic development. To retain its skilled labor force, Nigerian policymakers should prioritize strengthening the reliability and quality of its institutions, improving governance, as well as establishing pathways and incentives that encourage return migration, diaspora investment, and other contributions.

References

- Adelakun, O. J., et al. (2025). Economic and geopolitical push factors and migration tendencies in nigeria: The intervening role of institutional quality. *Scientific African*, 30, e02955. doi: 10.1016/j.sciaf.2025.e02955
- Akafa, T. A., Okeke, A., & Oreh, A. (2023). Push and pull factors of emigration among physicians in nigeria. *International Journal of Emerging Multidisciplinaries: Biomedical and Clinical Research*, 1(2), 1–12. doi: 10.54938/ijemdbmcr.2023.01.2.251
- EPRS. (2025). *The brexit paradox: How leaving the eu led to more migration*. European Union Briefing. Retrieved from <https://www.europarl.europa.eu/thinktank> (Accessed: 2025-11-22)
- Osayemwenwre, O. B., et al. (2023). Nigeria pattern of migration: Issues on human sustainable development. *Journal of Techno-Social*, 15(1), 1–15. doi: 10.30880/jts.2023.15.01.001
- Rienzo, C., & Vargas-Silva, C. (2015). Targeting migration with limited control: The case of the uk and the eu. *IZA Journal of European Labor Studies*, 4(1), 9. doi: 10.1186/s40174-015-0039-9

Rural-Urban Divergence in the United States: County-Level Evidence from 2014-2024

Ethan Przytulski, Calais Michaelsson

The aim of this research is to determine how the economic gap between rural and urban areas is changing over time. We argue that there is a divergence between these regions with growth in urban areas outpacing their rural counterparts. To do this, we estimated a linear panel regression on county-by-quarter data with the dependent variable in logs of total quarterly wages. Additionally, we used the USDA's Rural-Urban Continuum Code (RUCC) to divide regions into 9 categories based on their rurality (with 9 representing the most rural areas). We found that rural areas had lower baseline wages compared to urban ones. More importantly, these rural areas also demonstrated lower growth rates, with RUCC 9 areas having 1% lower quarterly growth compared to the most urban areas. A difference in growth was observed between RUCC 1 areas and every other category except 2, demonstrating that there is minimal difference in growth between mostly urban areas. This difference in growth compounds over time and leads to vastly different economic outcomes. This information is vital to policymakers and allows them to put into practice more targeted interventions to reduce economic disparities.

Introduction

ECONOMIC inequality remains a pressing issue both politically and socially. A survey by The Pew Research Center found that 61% of Americans believe there is too much economic inequality in the US, and 42% see this inequality as a key issue for the federal government (Kochhar, 2020). Despite this public concern, economic research in this area looks at the US as a whole, instead of the sub-national divisions. Similarly, research that only compares states can lead to comparisons of entire states and overlook common trends between similar areas across states. Our research aims to go deeper than the state level and look at the economic divide between rural and urban communities across the nation. Our research also goes beyond simply looking at urban-rural differences at a single moment and is instead more focused on similarities/differences between their levels of economic growth.

In addition, we move beyond single instance comparisons of rural and urban outcomes and focus instead on how their levels of economic growth differ over time. To make this comparison, we used county-level wage data from the Bureau of Labor Statistics' Quarterly Census of Employment and Wages (QCEW) and the USDA's Rural-Urban Continuum Codes (RUCC) to track the rural-urban divide over time. This allows us to evaluate whether the economic gap between rural and urban areas is converging or diverging. By showing which areas are advancing or falling behind, the analysis enables policymakers to craft more informed and effective policy responses.

Literature Review

There is relatively little research that treats rural economic development as its own area of study. Most work uses rural communities mainly as a comparison point to urban areas (Kilkenny, 2010). Even with this limited focus, existing research shows that rural communities face several unique barriers to growth, including weaker agglomeration effects, out-migration into cities, and dependence on natural resource industries. These issues do not affect urban areas in the same way, and they can contribute to slower long-run growth in rural counties.

At the same time, recent national data shows that the widening economic gap between rural and urban areas is not happening because rural communities are getting poorer. Instead, the main driver is that urban areas are growing much faster (U.S. Department of Commerce, 2023). So even if both areas are improving, the gap between them expands because cities are pulling away. This pattern directly relates to the central question of whether the rural-urban economic divide is increasing.

More detailed regional work supports this trend. Evidence of divergence across most of the Plains states has been documented, even when focusing on a specific geographic region (Mannion & Zougris, 2009). Other research has found that the relationship between inequality and growth is positive only in the 5% most densely populated counties (Fee, 2025). In the rest of the country, this relationship is either negative or nonexistent. In other words, the places with the strongest connection between inequality and growth are large urban counties, not rural ones. This is consistent with the idea that urban counties are the main engines of economic growth.

Another key finding from the literature is that there is a high degree of heterogeneity across rural areas where many underperform, while a couple of cases show strong economic performance (Ketels, Miller, & Bryden, 2004). While prior research shows that rural and urban areas often grow at different rates, there is still a gap in the literature. Most studies either compare the two groups at the national level or focus on specific regions. Very few

papers look at how the rural-urban economic divide changes across all smaller levels of rurality. To accomplish this, we are using the Rural-Urban Continuum Codes (RUCC), which assign all U.S. counties to one of nine categories based on population size and proximity to metropolitan areas.

This paper contributes to this gap by examining whether the economic divergence between rural and urban areas persists across the entire rural-urban continuum. This matters because understanding where the divide is strongest and whether it is driven by slow rural growth, fast urban growth, or both has direct implications for economic development policy.

Methodology

We estimate a linear panel model in which the dependent variable is the log of total quarterly wages, regressed on a linear time trend, a county’s rural–urban classification (RUCC), population, and the interaction between RUCC and time. This specification is designed to capture not only average wage growth over calendar quarters but also systematic differences in the evolution of wages across the rural–urban spectrum. The coefficient on the time variable reflects the baseline rate at which wages change across quarters for the reference RUCC category, while the RUCC terms capture persistent differences in wage levels associated with rurality or urbanicity at the baseline period. The population term adjusts for the fact that counties differ meaningfully in scale, ensuring that the model separates structural wage differences from the simple effect of having more or fewer residents. The interaction between RUCC and time is central to the interpretation: it allows each RUCC category to follow its own wage-growth trajectory rather than forcing rural and urban counties to share a common slope, thereby testing whether the wage gap between these areas widens, narrows, or remains stable over time.

$$\log(\text{Total Quarterly Wages}) = \beta_0 + \beta_1 \text{YearQtr} + \beta_2 \text{RUCC} + \beta_3 \text{Population} + \beta_4(\text{YearQtr} \times \text{RUCC}) + \varepsilon$$

Interpreting the coefficients on the log scale emphasizes proportional rather than absolute differences. In this framework, a coefficient represents the percentage difference in wages associated with a one-unit change in the covariate, holding other factors constant. This form is particularly appropriate for wage data, which tend to vary by orders of magnitude across counties and exhibit strong skewness. Working in logs stabilizes the variance, reduces the influence of extremely large counties, and ensures that effects are interpretable as relative changes that can be compared across areas with very different economic bases. Moreover, because wages evolve multiplicatively rather than additively, the log transformation aligns the model with the underlying economic process governing earnings growth.

The inclusion of the $\text{RUCC} \times \text{time}$ interaction follows naturally from this perspective: if rural and urban counties face different long-run structural pressures—such as differences in industrial composition, labor demand, or exposure to economic shocks—then their wage trajectories are unlikely to be parallel. The interaction term allows the model to reflect these distinct dynamic paths and reveals whether counties become more similar or more divergent in wage levels as time progresses. Taken together, the specification captures both the baseline movement of wages through time and the way those movements differ across the rural–urban continuum, providing a descriptive but informative view of long-run wage evolution at the county level.

Data

This paper utilized the Bureau of Labor Statistics (BLS) Quarterly Census of Employment and Wages (QCEW), a long-standing and comprehensive dataset of employer-reported wage and employment information across industries and occupations. First published in 1975, the QCEW underpins numerous economic programs and statistical outputs, and is regularly used in conjunction with the Department of Commerce and the Bureau of Economic Analysis to assess local economic growth and labor-market activity.

To build our dataset, we used the BLS API to retrieve county-level records by specifying county FIPS codes (five-digit geographic identifiers), year, and quarter. A simple R function, `qcewGetAreaData`, dynamically inserted these parameters into the API endpoint and returned the corresponding CSV file. A looped list structure automated the retrieval of these files, dividing them into two groups for memory management.

Using parallel processing to improve efficiency, the full dataset was compiled over roughly sixteen hours, resulting in approximately 260,000 observations spanning 2014 Q1 through 2024 Q3 (the extent of the APIs linked data). The long runtime reflects the limitations of the BLS API, the large number of U.S. counties, and the structure of the underlying data: files are stored by either year or quarter, without embedded quarter-level detail in annual files. While the BLS also provides ZIP archives containing cumulative and quarterly breakdowns for each county-year, these are not directly accessible through the streamlined API workflow used here.

After combining all retrieved files into a single dataset, the raw QCEW data was filtered to retain only cumulative county-level totals—excluding entries disaggregated by industry or ownership type (private, public, or federal). All relevant fields were then cleaned and converted from character formats to appropriate numeric types. To streamline the dataset for analysis, numerous “code” columns—BLS indicators denoting whether county-level changes exceeded or lagged behind national averages—were removed.

Finally, the cleaned QCEW dataset was merged with the Rural–Urban Continuum Codes (RUCC) published by the USDA Economic Research Service. RUCC codes classify every U.S. county on a scale from 1 (most urban) to 9 (most rural), incorporating factors such as population density, adjacency to metro areas, and infrastructure characteristics. This allowed the data set to pair economic results with consistent measures of rurality in the county.

Analysis

The estimated model examines how total quarterly wages evolve over time across counties with varying levels of rurality, as measured by the USDA Rural-Urban Continuum Codes (RUCC). The dependent variable is the natural log of total quarterly wages, which stabilizes variance and allows coefficients to be interpreted as percentage changes. The key independent variables include a continuous quarter index (`year_qtr_num`), the categorical RUCC measure, and an interaction between the two. As a result of these variables, all results are compared against the most urban counties (RUCC 1) and use Alabama as a baseline. This framework enables isolation of how rurality shapes both the level and trajectory of wage generation over the ten-year period from 2014 to 2024.

The results show a strong positive and statistically significant effect of time on county-level earnings. The coefficient on `year_qtr_num` implies that total quarterly wages grow at roughly 5% per quarter when averaged nationally. This reflects not only wage increases but also changes in employment levels, inflation, and broader macroeconomic expansion over the decade. When comparing across RUCC categories, the model indicates substantial baseline gaps between urban and rural counties. Most rural categories (especially RUCC levels 4 through 9) have significantly lower log wage totals than urban counties even after controlling for state differences and time. These results are consistent with longstanding patterns of urban areas hosting more diversified, dense, and higher-output labor markets.

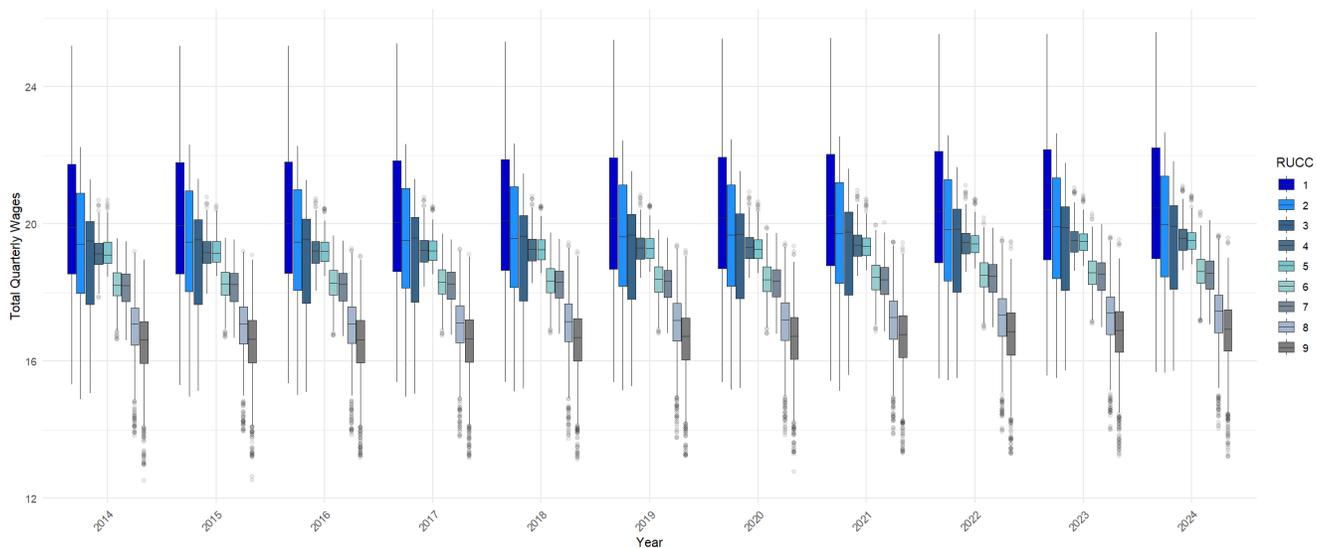


Figure 1: Total Quarterly Wages Aggregated by RUCC Code

The interaction terms between `year_qtr_num` and RUCC categories reveal an important dynamic: the growth rates themselves differ between urban and rural counties. For nearly all rural RUCC codes (4–9), the interaction coefficients are negative and statistically significant. This indicates that rural counties not only begin with lower total wage levels but also experience slower growth over time. The implication is that the economic gap between urban and rural regions has widened over the decade studied. For example, the most rural counties (RUCC 9) exhibit growth rates over one percentage point lower per quarter than the most urban counties, a difference that compounds substantially over ten years. Only the RUCC 2 category shows no meaningful deviation from urban growth patterns, suggesting that mildly rural counties keep pace more effectively than deeply rural counties.

High-income states such as California, Massachusetts, New York, New Jersey, and Washington show large positive coefficients, reflecting their consistently stronger wage and employment bases. Conversely, several rural or low-income

Dependent variable: log(Total Quarterly Wages)					
	Model (1)	Model (2)	Model (3)	Model (4)	Model (5)
RUCC Code = 2	-0.6743*** (0.0135)	-0.6743*** (0.0134)	-0.1397*** (0.0118)	-0.1397*** (0.0118)	0.5182 (7.3340)
RUCC Code = 3	-1.1658*** (0.0136)	-1.1658*** (0.0135)	-0.4438*** (0.0121)	-0.4438*** (0.0121)	3.7558 (7.4092)
RUCC Code = 4	-0.9292*** (0.0162)	-0.9292*** (0.0161)	-0.1751*** (0.0143)	-0.1750*** (0.0143)	8.5565 (8.8362)
RUCC Code = 5	-0.9159*** (0.0233)	-0.9159*** (0.0232)	-0.1384*** (0.0203)	-0.1384*** (0.0203)	16.2564 (12.6949)
RUCC Code = 6	-1.8643*** (0.0134)	-1.8643*** (0.0133)	-1.0380*** (0.0121)	-1.0379*** (0.0121)	12.6122* (7.2848)
RUCC Code = 7	-1.9608*** (0.0150)	-1.9608*** (0.0150)	-1.1220*** (0.0135)	-1.1220*** (0.0135)	15.9767* (8.2007)
RUCC Code = 8	-3.1305*** (0.0126)	-3.1305*** (0.0125)	-2.2689*** (0.0115)	-2.2688*** (0.0115)	11.4637* (6.8702)
RUCC Code = 9	-3.6238*** (0.0120)	-3.6238*** (0.0119)	-2.7512*** (0.0111)	-2.7511*** (0.0111)	14.2022** (6.5499)
Year/Qtr		0.0421*** (0.0011)	0.0408*** (0.0009)	0.0401*** (0.0010)	0.0457*** (0.0025)
Population			0.000002*** (0.0000)	-0.00001** (0.00001)	0.000002*** (0.0000)
Year/Qtr:Population				0.0000** (0.0000)	
RUCC Code(2):Year/Qtr					-0.0003 (0.0036)
RUCC Code(3):Year/Qtr					-0.0021 (0.0037)
RUCC Code(4):Year/Qtr					-0.0043 (0.0044)
RUCC Code(5):Year/Qtr					-0.0081 (0.0063)
RUCC Code(6):Year/Qtr					-0.0068* (0.0036)
RUCC Code(7):Year/Qtr					-0.0085** (0.0041)
RUCC Code(8):Year/Qtr					-0.0068** (0.0034)
RUCC Code(9):Year/Qtr					-0.0084*** (0.0032)
Constant	20.2588*** (0.0091)	-64.6594*** (2.1426)	-63.0007*** (1.8470)	-61.6537*** (1.9345)	-72.8398*** (4.9821)
Observations	131,898	131,898	131,898	131,898	131,898
R ²	0.5149	0.5206	0.6437	0.6438	0.6438
Residual Std. Error	1.2305	1.2232	1.0545	1.0545	1.0545

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 1: Regression Results

states—such as Kansas, Nebraska, Mississippi, Montana, and the Dakotas—show negative fixed-effect estimates, capturing structural differences in labor markets that persist beyond RUCC classification. The fact that RUCC effects remain strong and significant even after accounting for these state-level differences reinforces the conclusion that rurality itself plays a distinct role in shaping economic performance.

Overall, the regression indicates both static inequality—where rural counties start with much lower total wage output—and dynamic inequality, where wage growth rates diverge over time. Urban counties not only maintain but expand their economic advantage relative to rural areas. These findings support the hypothesis of widening rural–urban economic divergence in the United States and highlight the importance of spatial economic structure in shaping long-term wage outcomes.

Shortcomings

Several limitations should be acknowledged when interpreting these findings. First, the underlying QCEW data may contain structural inconsistencies, especially in small or sparsely populated counties. Very low-population counties often have few employers, and for confidentiality reasons the BLS suppresses or perturbs certain values; in other cases, county governments or employers may under-report or misclassify wages. These issues may introduce noise or downward bias in rural wage estimates, potentially exaggerating rural–urban differences. Additionally, because the data were pulled via a large batch API process, any dropped requests, improperly formatted CSVs, or incorrect concatenations could contribute subtle inconsistencies across years or counties—particularly given the long runtime and large number of files retrieved.

The regression model itself, while structurally sound, also presents some methodological constraints. Because states contain both highly urban and deeply rural counties, fixed effects absorb only statewide shocks, not county-level demographic differences. As a result, the coefficients on rurality (RUCC) reflect wage differences net of state-level factors but still confound population scale with economic performance. Furthermore, modeling total quarterly wages rather than per-capita or per-establishment measures inherently favors larger counties, where even modest wage levels can produce enormous total wage sums due to sheer employment volume. This choice of dependent variable captures total economic output but may not be the best metric for assessing comparative economic well-being or wage growth across urban and rural communities.

Finally, the interaction between quarter index and rurality assumes a linear and uniform growth pattern over time, which may oversimplify the complex business cycles and industry-specific shocks that affected counties differently between 2014 and 2024. Specifically, the YearQtr coefficient of 0.051 gives cause for concern when compared to the calculated compounded annual growth and the published annual growth of total wages. We were unable to discern what caused this anomaly and were unable to address it in our regression despite multiple iterations and changes.

Conclusion

This paper moves the conversation on economic inequality beyond national and state averages to the county level, where the rural–urban divide is lived and felt. By pairing county QCEW wage data with the USDA’s RUCC classification and tracking changes quarter by quarter, we show a clear, structured relationship between rurality and both the level and trajectory of wage activity. Urban counties consistently outpace their rural counterparts, and those differences are not only persistent but tend to widen over time. These patterns echo well-established mechanisms—agglomeration, market thickness, infrastructure, and talent concentration—while recognizing that rural development is heterogeneous and that some rural places succeed on their own terms.

Why these results matter is practical as much as it is academic. A county-level view turns a diffuse debate about inequality into a concrete map of where gaps are largest and how they evolve. That, in turn, helps policymakers, businesses, and civic leaders target place-based investments—broadband, transportation, childcare, workforce training, industry diversification—where they can bend growth trajectories rather than simply chase outcomes. It offers a way to monitor whether disparities are narrowing or widening and to design interventions that reflect the distinct realities of rural and urban labor markets. For firms, the findings inform location strategy and talent planning; for communities, they provide evidence to make the case for resources and partnerships that can unlock local potential.

At the same time, readers should treat these findings as descriptive rather than causal. RUCC codes are a useful proxy for rurality but encompass factors that overlap with industry mix, infrastructure, and policy environments. Our focus on total wages highlights economic scale but can conflate county size with well-being; complementary measures—per capita wages, employment rates, or establishment counts—would enrich the picture. QCEW data are comprehensive yet imperfect in small or sparsely populated counties, and a single linear time trend cannot capture the full complexity of shocks and recoveries. Future work can strengthen inference by incorporating county and

calendar-quarter fixed effects, exploring nonlinear time dynamics and structural breaks, and testing targeted policies in event-study or difference-in-differences frameworks.

Taken together, the contribution here is a tractable, nation-wide lens on the rural–urban wage landscape—one that quantifies where and how the divide manifests and provides a baseline for action. If the goal is widely shared growth, the path forward lies in translating this map into tailored investments and credible tests of what works, so counties at every point on the rural–urban continuum can participate more fully in the gains of a globalized economy.

References

- Ketels, C. H. M., Miller, K., & Bryden, R. T. (2004). Competitiveness in Rural U.S. Regions: Learning and Research Agenda.
- Kilkenny, M. (2010). Urban/Regional Economics and Rural Development. *Journal of Regional Science*, 50(1), 449–470. Retrieved 2025-11-18, from <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9787.2009.00661.x> (_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-9787.2009.00661.x>) doi: 10.1111/j.1467-9787.2009.00661.x
- Kochhar, J. M. H., Ruth Igielnik and Rakesh. (2020, January). *Most Americans Say There Is Too Much Economic Inequality in the U.S., but Fewer Than Half Call It a Top Priority*. Retrieved 2025-11-20, from <https://www.pewresearch.org/social-trends/2020/01/09/most-americans-say-there-is-too-much-economic-inequality-in-the-u-s-but-fewer-than-half-call-it-a-top-priority/>
- Mannion, E., & Zougris, K. (2009, January). Metro-nonmetro Economic Growth and Convergence in the Plains States - Rethinking the Rural-Urban Relationship in a Global Economy. *Online Journal of Rural Research & Policy*, 4(1). Retrieved 2025-11-18, from <https://newprairiepress.org/ojrrp/vol4/iss1/2> doi: 10.4148/ojrrp.v4i1.50
- U.S. Department of Commerce. (2023, June). *Geographic Inequality on the Rise in the U.S. | U.S. Department of Commerce*. Retrieved 2025-11-18, from <https://www.commerce.gov/news/blog/2023/06/geographic-inequality-rise-us>

Student Debt and Early-Career Earnings

Emerson Au, Justin Sweeney, Harish Rajagopalan

Repayment of student loans usually starts when graduates begin low-paying jobs, which creates concerns that repayment obligations might lower early-career wages. This study investigates whether changes in borrowers' yearly incomes are related to the beginning of student loan repayment. We track a cohort of individuals born between 1980 and 1984 from 1997 to 2007 using the National Longitudinal Survey of Youth 1997 (NLSY97). We use a staggered difference-in-differences event-study approach to estimate wage dynamics around the first year in which a borrower's outstanding loan total decreases. Pre-treatment estimates violate the parallel trends assumption necessary for a causal interpretation by demonstrating consistent disparities in salary trajectories between earlier and later repayers. Post-treatment effects are negligible and statistically indistinguishable from zero. We find no conclusive evidence that salaries change at the start of repayment within these bounds, and our research is further limited by missing controls on hours worked, occupation, sector, and socioeconomic background.

Introduction

THE way US graduates enter the labor market is influenced by their student loan debt. Repayment usually starts when employees start low-paying, unstable jobs in their early careers. Previous research indicates that debt can impact labor supply, sector choice, and job search, but there is little data on whether salaries change after repayment starts, particularly among comparable borrowers.

This paper provides descriptive evidence on wage patterns around the start of repayment from the NLSY97 panel. We track annual wages across early adulthood and examine how they change in the years before and after repayment starts. By using an event-study framework that compares earlier and later repayers, we assess whether there is a difference in wage trajectories after the beginning of repayment.

Literature Review

Recent studies consider links between student loan debt, early-career job search, sector choice, and income, but results differ across disciplines and identification strategies. Ji (2020) develops a dynamic general-equilibrium search model calibrated to NLSY97 data and shows that, because of repayment obligations, indebted graduates search less intensively and accept lower-paying first job offers. In his calibration, borrowers search about 0.9 weeks less and earn roughly \$1,479 less in their first year than comparable non-borrowers. Reduced-form OLS estimates indicate that an additional \$10,000 of debt is associated with 1.41–1.57 fewer weeks of search and 2.7–4% lower early earnings. Ji also runs counterfactuals that highlight the role of policy design: income-based repayment (IBR) insures job-search risk, raises average wage income, nearly eliminates default, increases college attendance, and generates welfare gains of about 1.8% through higher lifetime income and consumption. Overall, the paper provides evidence for a liquidity-constraint mechanism in which initial wages and search effort are both reduced by student debt.

Complementing this mechanism, Rothstein and Rouse (2007) present causal evidence that debt alters occupational choice. Analyzing a no-loans financial-aid policy reform at an elite university through difference-in-differences and instrumental variables methods, they show that increases in debt push students into the higher-salary private sector, and away from public service. In particular, they find that each additional \$10,000 of debt decreases the probability of entering a government, nonprofit, or education career by 5–6 percentage points (from a $\sim 17\%$ baseline), and is associated with an approximately \$2,000 higher annual starting salary. In a contrasting study, Daniels and Smythe (2019) utilize the National Longitudinal Survey of Youth cohort 1997 (NLSY97), within a difference-in-differences framework, to find that debt holders earn +8% more income, work +6% more hours, have +1% higher wages, and have a +5 percentage point higher probability of full employment than non-debt holders. Their findings indicate that the majority of income differences are due to labor supply and not a wage premium, suggesting a behavioral adjustment, where debt holders work more hours to fulfill their obligations, despite small differences in hourly earnings.

Policy designs elsewhere around the world help to support these observations. For example, Beyer, Hastings, Neilson, and Zimmerman (2015) examine the loan-cap system in Chile which ties allowable borrowing to potential future earnings by degree level and characteristics about the student. The loan caps shift resources to programs with higher expected returns which can affect the rates of graduation, repayment, and default. While the Chilean analysis does not provide estimates about individual wages as the studies in the U.S. do, the Chilean evidence still points to how design features of the programs can create incentives that drive both borrowing and labor market outcomes based upon expected earnings. Across these studies, a consistent narrative is built: debt shifts

early-career behavior—through search intensity (Ji), sector sorting (Rothstein & Rouse) and labor supply (Daniels Jr. and Smythe)—potentially augmented by policy rules (IBR; loan caps). Our contribution is to extend this research using NLSY97 data from 1997–2007, restricting attention to individuals who ever hold student loans, and applying a staggered difference-in-differences/event-study design to examine wage trajectories around the start of loan repayment. We focus on log annual wages and ask whether there is any detectable change in wages when repayment begins, relative to borrowers who have not yet started repayment. Given data limitations on hours worked, sector, and detailed socioeconomic background, we do not attempt to separate wage effects from labor supply or sectoral mechanisms; instead, we provide descriptive evidence on whether wages themselves move at the onset of repayment.

Data

We use data from the U.S. Bureau of Labor Statistics (NLSY97), a sample of 8,984 people born from 1980–1984 and surveyed annually or biennially since 1997. The panel structure allows us to observe wage trends before and after the repayment of loans has started. We focus our study on the years from 1997–2007 to avoid labor market disruptions from the 2008 financial crisis and because survey participation becomes sparser in later years, with more survey years missing. This time window captures respondents in their early labor market years.

Our analysis is restricted only to those who have received a student loan. Non-borrowers differ systematically from borrowers in ways that affect both wages and repayment behavior. Including them would create a third, structurally different comparison group and would obscure the actual effect of starting repayment.

A limitation of the dataset is that relatively few people have fully repaid their loans within our observation window. As a result, our analysis focuses on the effects of starting repayment rather than the effects of completing repayment. We take the natural log of yearly wages to reduce skewness and heteroskedasticity in the distribution and to interpret estimated effects as approximate percent changes.

Methodology

We estimate the effect of beginning student loan repayment using a staggered difference-in-differences (DiD) approach designed by Callaway and Sant’Anna(2021). A traditional two-period DiD would be inappropriate because individuals begin repayment in different years, so there is no well-defined “before” and “after” period. The staggered DiD framework addresses this by treating each repayment year as its own treatment cohort.

Treatment is defined as the first year in which an individual’s outstanding loan balance decreases relative to the previous year, indicating the start of the repayment process. For example, all individuals who begin repayment in 2004 form the 2004 group, those who started repayment in 2005 form the 2005 group, and so on. For each treated cohort, the control group consists of borrowers who have not yet started repayment in that year (and are therefore not yet treated) or never start repayment in the sample period. By restricting the treatment and control group to only borrowers, we ensure that comparisons are made between individuals with similar repayment responsibilities and financial constraints.

The estimator computes the Average Treatment Effect on the Treated (ATT) for each separate cohort and each post-treatment year. We then aggregate each cohort’s ATTs using a simple average to obtain an estimate that reflects the average effect of beginning repayment across all treatment cohorts. Parallel trend assumptions are evaluated in the Analysis section.

Results and Analysis

Table 1 reports the event-study estimates from the staggered DiD design. Each event-time coefficient measures the difference in $\ln(\text{wage})$ between borrowers who have begun repaying their loans at event time t and borrowers who have not yet started repayment in that same year, averaged across repayment cohorts.

The pre-treatment coefficients (Tm5–Tm1) reveal notable differences between the two groups before repayment starts. In particular, Tm4 is large and highly significant, and the pre-treatment average (Pre_avg) is strongly negative and statistically different from zero. These results indicate that borrowers who repay their loans earlier follow different wage trajectories than those who repay their loans later. Nearly all the post-treatment coefficients (Tp0–Tp4) are statistically insignificant, meaning there is no consistent pattern of wage changes after the repayment process has begun. Although Tp3 is significantly different from zero, the violation of the parallel trends assumption in the pre-treatment period prevents this from being interpreted as a causal effect of repayment.

Figure 1 visualizes these event-study estimates. Each point corresponds to an ATT estimate for a given year relative to the start of repayment, and the vertical bars show 95% confidence intervals. The horizontal axis measures

Table 1: Event-Study Estimates of Aggregate Differences in $\ln(\text{wage})$

Event Time	Coefficient	Std. Error	z-stat	p-value	CI Lower	CI Upper
Pre_avg	-0.4329	0.0655	-6.61	0.000	-0.5612	-0.3046
Post_avg	-0.1492	0.2926	-0.51	0.610	-0.7227	0.4242
Tm5	-0.3993	0.2613	-1.53	0.126	-0.9115	0.1128
Tm4	-2.2567	0.4817	-4.68	0.000	-3.2009	-1.3126
Tm3	0.3266	0.2683	1.22	0.223	-0.1992	0.8524
Tm2	0.1435	0.2252	0.64	0.524	-0.2979	0.5849
Tm1	0.0214	0.1374	0.16	0.876	-0.2478	0.2907
Tp0	0.0605	0.1139	0.53	0.595	-0.1626	0.2837
Tp1	0.2407	0.1878	1.28	0.200	-0.1274	0.6087
Tp2	-0.2774	0.3696	-0.75	0.453	-1.0019	0.4471
Tp3	-0.5664	0.1202	-4.71	0.000	-0.8019	-0.3308
Tp4	-0.2036	1.0956	-0.19	0.853	-2.3510	1.9438

years relative to treatment, while the vertical axis shows the estimated change in $\ln(\text{wage})$.

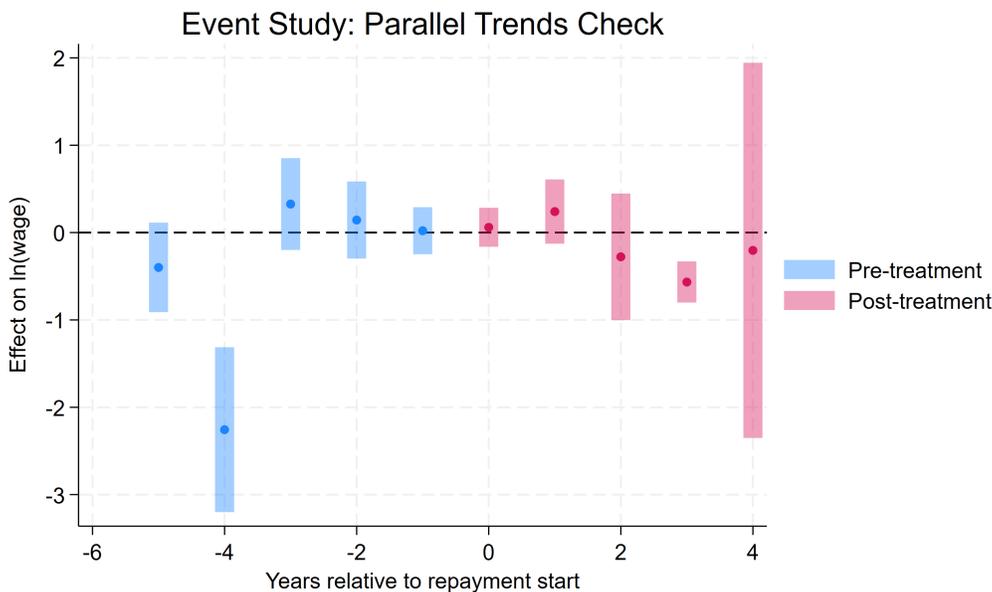


Figure 1: Event-study plot of ATT estimates before and after repayment starts.

Most of the pre-treatment points are close to zero, but the joint test of all pre-treatment coefficients strongly rejects the null that they are all zero ($\chi^2(10) = 261.23$, $p < 0.001$), confirming that early and late repayers differ systematically even before repayment begins.

Table 2 reports the overall Average Treatment Effect on the Treated (ATT) combining post-treatment periods. The point estimate is small in magnitude and very imprecise.

Table 2: Average Treatment Effect on the Treated (ATT)

	Coefficient	Std. Error	z-stat	p-value
ATT	-0.0573	0.3531	-0.16	0.871

With a high p -value (0.871), there is no statistically significant effect of loan repayment on wages detectable in these results. Combined with the strong pre-treatment differences documented in Table 1, this means we cannot interpret the ATT as the causal effect of starting repayment: there is no clear, robust relationship between the onset of student loan repayment and annual wages in this sample.

Conclusion

Within this sample of NLSY97 borrowers we do not see a significant, consistent shift in yearly salaries at the start of student loan repayment. The parallel trends assumption that supports a causal difference-in-differences interpretation is violated by the substantial pre-treatment differences between borrowers who begin repayment earlier and those who begin later, who already have distinct salary trajectories. These pre-existing differences suggest that borrowers differ well before repayment begins, which makes it difficult to isolate any causal effect of repayment itself. We cannot say that repayment itself consistently increases or decreases wages since, after payback begins, the estimated post-treatment effects are insignificant and imprecise, and the aggregate ATT is statistically indistinguishable from zero.

It is important to read these results with several caveats. Our data window is relatively short and includes only the early labor-market years; few respondents fully repay their debts, so we cannot analyze the impacts of being debt-free. We cannot isolate salary impacts from labor-supply or sorting mechanisms and cannot rule out omitted-variable bias since we do not have access to important controls like hours worked, occupation and sector, institutional selectivity, major, and comprehensive socioeconomic background. Furthermore, complexity like postponement, forbearance, or consolidation may be overlooked by our treatment measure, which is the first recorded drop in outstanding balance. The paper's contribution is therefore descriptive: we show that early and late repayers are on distinct salary paths well before payback begins, and we find no discernible wage split at repayment initiation. To determine the causal effect of student debt payment on wages and to better link wage dynamics to search behavior, sector choice, and labor supply, future research utilizing richer administrative loan data or policy-driven variation in repayment terms is required.

References

- Beyer, H., Hastings, J., Neilson, C., & Zimmerman, S. (2015). Connecting student loans to labor market outcomes: Policy lessons from Chile. *American Economic Review*, *105*(5), 508–513.
- Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200–230.
- Daniels Jr., G. E., & Smythe, A. (2019). Student debt and labor market outcomes. *AEA Papers and Proceedings*, *109*, 171–175.
- Ji, Y. (2021). Job search under debt: Aggregate implications of student loans. *Journal of Monetary Economics*, *117*, 741–759.
- Rothstein, J., & Rouse, C. E. (2011). Constrained after college: Student loans and early-career occupational choices. *Journal of Public Economics*, *95*(1–2), 149–163.
- U.S. Bureau of Labor Statistics. (2023). *National longitudinal survey of youth 1997 (nlsy97)*. National Longitudinal Surveys, U.S. Department of Labor. (<https://www.nlsinfo.org/content/cohorts/nlsy97>)

Who Steps Forward? Demographic Patterns in Living vs. Deceased Kidney Donation

Siya Motwani, Lidia Rodriguez

Introduction

CHOOSING to donate an organ offers insight into both medical and social phenomena, revealing how biological eligibility, social expectations, and structural factors determine who becomes a donor. Deceased donors offer a population-level baseline: individuals whose medical condition and circumstance of death make donation possible. Living donation, on the other hand, adds an additional behavioral layer, shaped by social norms, relationships, and opportunity. By comparing living and deceased donation patterns, we can better understand how medical constraints and human behavior interact to shape donor supply. Research has examined these groups separately, consistently finding that women are more likely to become living donors and that racial minorities are underrepresented across both donor types. Few studies, however, directly compare the demographic characteristics of living and deceased donors to understand why these patterns diverge. Directly comparing the demographic characteristics of living and deceased donors allows us to distinguish factors driven by biological eligibility and circumstances of death from those driven by social pressures, structural inequities, or cultural expectations surrounding living donation. This paper addresses that gap by analyzing and comparing the gender and racial/ethnic composition of living and deceased kidney donors using the OPTN STAR database. In doing so, it seeks to uncover whether these differences arise from social pressures, systemic inequities, or broader cultural expectations surrounding donation. By contrasting the demographic "baseline" of deceased donors with the selective patterns of living donation, this study contributes to a deeper understanding of how social and structural dynamics influence participation in organ donation and what that means for promoting diversity and equity within the transplant system.

Literature Review

Current literature on organ donation consistently shows that demographic patterns differ between living and deceased donors, yet most studies examine these populations in isolation. Early work by Dobson found that women are more likely than men to become living organ donors, suggesting that gender differences may stem not only from altruistic motives but also from relational expectations and social pressure (Dobson, 2002). Building on this, Ross and Thistlethwaite (2021) demonstrated that women remain overrepresented and Black individuals underrepresented among living kidney donors, highlighting the influence of structural and cultural dynamics, such as caregiving norms, economic barriers, and healthcare access, on donation behavior. Racial disparities have also been linked to trust and perceptions of fairness within the transplantation system. In a seminal study, Siminoff et al. (2006) found that African American patients and families expressed greater mistrust toward medical institutions and perceived organ allocation as less equitable, contributing to lower consent and registration rates compared to White respondents. These attitudinal barriers reveal that differences in donation rates are not only about medical suitability but also reflect historical contexts that shape willingness to donate. In contrast, deceased donation operates on a different set of determinants. Rather than voluntary action, deceased donation reflects population-level patterns driven by factors like mortality, cause of death, medical eligibility, and consent processes. Across multiple national datasets, the most prominent deceased donors are White and male. A Centers for Disease Control and Prevention (CDC) analysis of 70,414 U.S. deceased donors from 2010–2017 found that 59.6% were male and 66.2% were White (Abara et al., 2019). Similarly, a JAMA Surgery report confirmed that White individuals consistently represent the highest proportion of deceased donors in the United States (Kernodle et al., 2021). The gender reversal between donor types, which is women predominating in living donation but men in deceased donation, has drawn growing academic attention. Yee et al. (2021) found that women account for about 60% of living donors but only around 40% of deceased donors, emphasizing that these differences arise from distinct mechanisms, such as women's greater socialization into caregiving and relational obligations versus men's higher likelihood of dying in trauma-related or donation-compatible circumstances. Despite previous work on gender and racial disparities in both living and deceased donation, few studies analyze these donor types side-by-side to determine whether their demographic differences arise from biological, structural, or social factors. The existing literature establishes patterns within each group, but rarely interrogates how these patterns compare or why they diverge. This study contributes to this gap by directly comparing the demographic populations of living and deceased kidney donors. By comparing these two donor pathways, this analysis helps to understand which disparities reflect medical eligibility and circumstances of death (relevant to deceased donation) versus social pressures, cultural norms, and structural barriers (more for living

donation). In doing so, it provides a more integrated understanding of how demographic inequities emerge across the entire organ donation system, offering a foundation for targeted interventions and more equitable policy design.

Research Method & Data

This study uses a quantitative approach to analyze geographic and demographic trends in kidney donation across the United States. Three datasets were collected from the Organ Procurement & Transplantation Network (OPTN), which contains national statistics on organ donors. Each dataset reports annual kidney donor counts but categorizes donors differently: by donor-recipient relationship, donor state, donor age, donor sex, and donor race/ethnicity. A key variable shared across all datasets was donor type, such as living versus deceased donation. Because donor type and year were consistent throughout all datasets, these two variables were used to merge the datasets into a single file. After merging, donor counts were aggregated by state and a choropleth map was used to illustrate the distribution of kidney donation across the United States. There are several methodological considerations to keep in mind when interpreting these findings. The datasets only align on donor type and year, meaning that not every demographic variable appears for every observation after merging, causing some combined records to be incomplete. Additionally, our analysis is descriptive and reports total donor counts without adjusting for state population size, healthcare access, or transplant demand, which may exaggerate the contribution of states with a larger population. Additionally, OPTN data collection may vary slightly from year to year. Despite these limitations, this methodology allows for a replicable and informative exploration of national donation patterns.

Data Analysis

Our analysis examines geographical variation in kidney donors, the distribution of living versus deceased donors by relationship, and major trends in kidney donation in the United States over time. Figure 1 displays the total number of kidney donors by state. The choropleth map above depicts a range of colors for each state based on the number of donors. As highlighted, the states of California, Texas, New York and Florida, colored by the darkest shades, are the states with the highest number of donors and will, therefore, be those states of focus in this paper. The map also shows variation across regions, with northeastern, southern, and western states generally exhibiting darker shading compared to several states in the upper Midwest region. This figure highlights that kidney donor totals are not evenly distributed across the country. Additionally, in Figure 2, donor type is compared by donor-recipient

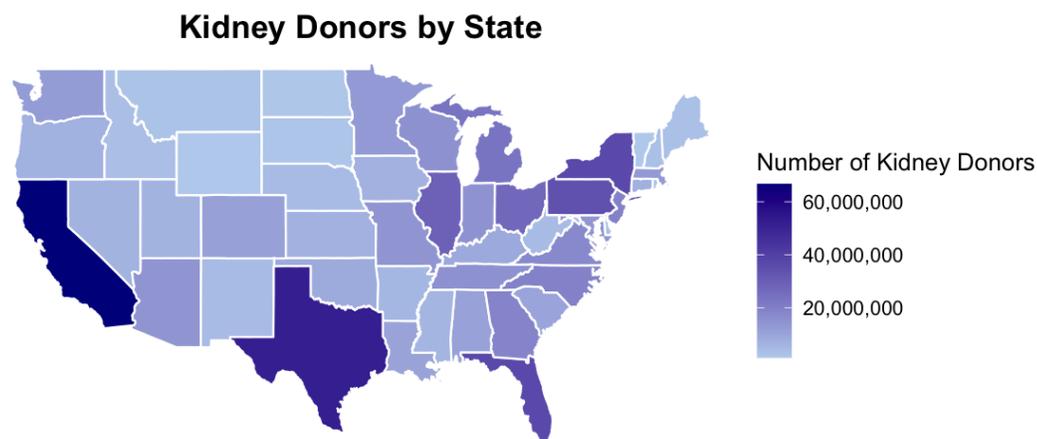


Figure 1: Kidney Donors by State

relationship. Deceased donors, expectedly, represent the largest contribution to kidney donation in the dataset. Living donors appear in multiple relationship categories, with full siblings, children, spouses, and parents showing higher percentages among the living donor population. Other living donor categories, such as unrelated or paired

donation, appear considerably smaller in comparison. This figure illustrates the distribution of donor categories within the sample rather than their outcomes or motivations. Figure 3 presents the counts of living versus deceased

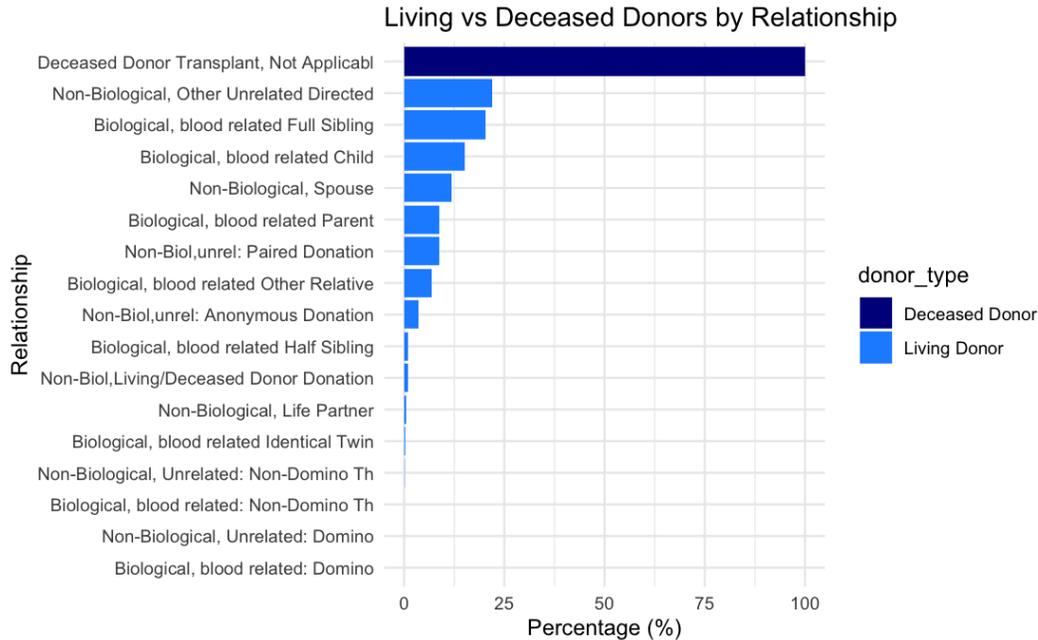


Figure 2: Living vs Deceased Donors by Relationship

kidney donors from 2000 to 2024. The number of deceased donors increased during this period, particularly from around 2015 onward. In contrast, the number of living donors does not show the same degree of increase and remains more stable throughout the timeline. There is also a visible decrease in living donor numbers around 2020, followed by a slight increase in subsequent years. This visualization depicts differences in the trajectories of the two donor types without making conclusions about the causes of these patterns.

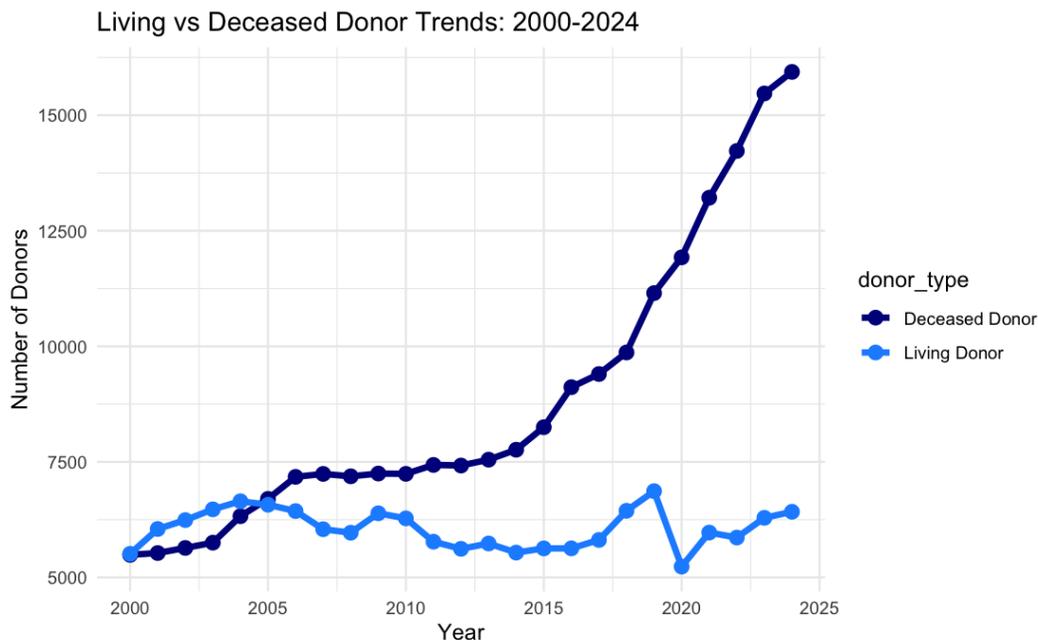


Figure 3: Living vs Deceased Donors Trends from 2000 to 2024

Together, these figures summarize the key characteristics of available donor data: (1) donor totals differ across U.S. states, (2) deceased donors make up a large portion of kidney donations, and (3) deceased donor counts have increased over time, while living donor counts have remained more consistent. These trends suggest that living and

deceased donation operate through different mechanisms, one shaped more by social and structural factors, the other by population-level patterns such as mortality and medical suitability. Although limited to the available dataset, these observations highlight why examining the two donor types together is important for understanding where disparities emerge and what may be driving them.

Shortcomings

Despite providing a descriptive overview of kidney donation patterns in the United States, this study has several important limitations that constrain our ability to fully answer our research question. First, as noted by Ross and Thistlethwaite (2021), the underrepresentation of Black living donors may partially stem from higher rates of health co-morbidities in Black communities, which make potential donors medically ineligible. However, this factor alone cannot fully account for the disparities observed. Our OPTN dataset only captures certain demographic variables, such as age, sex, race/ethnicity, and donor-recipient relationship, but does not include other contextual factors like socioeconomic status, education, or cultural influences that may affect an individual's willingness to donate. Additionally, medical eligibility constraints, such as comorbidities or other health conditions, are not captured across all datasets, particularly those for living donors, which may disproportionately affect certain racial groups. While we are observing clear patterns in donor type distribution across states over time, we cannot determine the mechanisms driving these differences beyond assumptions. Future work should address these gaps by integrating additional datasets that capture social, economic, and health-related variables, allowing for a more comprehensive understanding of our findings.

Conclusion

This study performs a descriptive analysis of kidney donation patterns across the United States, with a focus on overall donor counts by state, donor type, and donor-recipient relationships. While our current data does not allow for definitive conclusions about demographic overrepresentation or misrepresentation, it shows important patterns in donation volume and geographic distribution. These observations suggest further exploration into disparities in living versus deceased donation and provide a foundation for studies that incorporate additional demographic and contextual factors. Future research that includes variables such as socioeconomic status, and cultural attitudes could help define the mechanisms driving differences in donation behavior and guide strategies to promote equity and participation across diverse populations, so that it is clear what policies or political changes need to be made to incentivize different demographic groups to donate to reach a more diverse population and better matching practices.

References

- Abara, W. E., Collier, M. G., Moorman, A., et al. (2019). Characteristics of deceased solid organ donors and screening results for Hepatitis B, C, and Human Immunodeficiency Viruses — United States, 2010–2017. *MMWR Morbidity and Mortality Weekly Report*, *68*, 61–66.
- Dobson, R. (2002). More women than men become living organ donors. *BMJ*, *325*(7369), 851.
- Kernodle, A. B., Zhang, W., Motter, J. D., et al. (2021). Examination of racial and ethnic differences in deceased organ donation ratio over time in the US. *JAMA Surgery*, *156*(4), e207083.
- Ross, L. F., & Thistlethwaite, J. R. (2021). Gender and race/ethnicity differences in living kidney donor demographics: Preference or disparity? *Transplantation Reviews*, *35*(3), 100614.
- Siminoff, L. A., et al. (2006). Racial disparities in preferences and perceptions regarding organ donation. *Journal of General Internal Medicine*, *21*(9), 995–1000.
- Yee, E., et al. (2021). Sex disparities in organ donation: Finding an equitable donor pool. *Journal of the American Heart Association*, *10*(19), e020820.

Housing Market Responses to New Transit: An Exploratory Analysis of Silver Line Phase II in Northern Virginia

Katelyn Kim, Elizabeth Odife

This paper seeks to see how the opening of the Washington, D.C. Metrorail Silver Line Phase II in 2022 affects residential housing prices in the developing suburbs of Northern Virginia. Existing literature focuses primarily on more mature rail systems, like the BART in the Bay Area. Unlike prior research, we are examining an area that was more auto-dependent before the opening of the stations. This allows us to explore early housing market response to new transit connectivity. We use a descriptive, exploratory analysis approach to compare trends over time of how the opening of the stations impacted housing prices. We analyzed Zillow Home Value Index data from 2018 to 2025 across 12 zip codes in Reston, Herndon, Sterling, and Ashburn, Virginia, categorizing each by distance from the newly opened stations. We find that on average, zip codes closer to the stations tend to exhibit lower average home values relative to those farther away.

Introduction

INFRASTRUCTURE investments, such as major transit extensions, often reshape urban and suburban housing markets. In the Washington, D.C. region, the Metrorail Silver Line Phase II extension opened in 2022, linking these suburban and emerging neighborhoods to hubs like Dulles Airport and regional job centers. This extension was part of one of the nation's largest rail projects, helping to improve mobility, accessibility, and reduce car dependency. Although this makes surrounding areas more desirable, there are concerns about how this can affect housing affordability due to transit premiums. After Phase I in 2014, many areas around the stops saw sharp increases in property values, amounting to a 138% increase in value since 2013 (Ingram, 2025). With current debates over affordability and inclusionary zoning, examining how transit development affects housing prices in these areas highlights who benefits, who faces greater burdens, and how policy might respond.

Existing research has documented that closer proximity to transit often yields premiums on property value, typically ranging from a 3% to 12% increase (Rennert, 2022). This is especially the case near newly opened stations, but effects vary by factors such as distance, neighborhood characteristics, and local housing market conditions. Additionally, prices may rise before official openings in some areas due to expectations (LaRoche, 2019). Prior research has primarily focused on regions such as Berkeley, California, and Minneapolis, Minnesota, which are considered more mature urban areas and are already transit-linked. Our research focuses on three areas in the Northern Virginia region: Reston, Herndon, and Loudoun, which are rapidly developing communities where many residents commute toward Washington, D.C. for work. The value of examining the Silver Line Phase II in these areas is that they previously lacked rail access and were largely auto-oriented. Using city-level measures of station distance and population density from 2018 to 2025, this paper explores how the opening of new transit stations influences housing prices and how these effects may differ from those observed in transit-mature urban environments. Understanding these dynamics is important for informing regional planning efforts, especially as concerns about affordable housing continue to increase.

Given the current literature on transit premiums and housing prices, we hypothesize that the expansion will cause an increase in housing prices, particularly for homes in closer proximity to the newly opened stations.

Literature Review

Research across urban economics consistently shows that rail transit investments can influence residential property values, though the effects vary substantially across regions, transit types, and neighborhood characteristics. Through meta-analyses, Debrezion, Pels, and Rietveld (2007) and Mohammad, Graham, Melo, and Anderson (2013) highlight a general pattern that homes closer to rail stations tend to result in higher prices, with the effect declining sharply as distance increases. This distance gradient appears across many metropolitan areas and is strongest within roughly a half mile to one mile of stations, emphasizing the importance of modeling proximity. Additionally, based on a quasi-experimental approach using transportation innovations in London and applying hedonic price regressions with fixed effects, there is emphasis on the fact that residential property value effects depend highly on station context, including land-use patterns, zoning, and neighborhood density (Gibbons & Machin, 2005). Across the United States and other parts of the world, areas with established transit-oriented development, higher density, or strong local amenities often show stronger capitalization effects. Conversely, low-density, auto-oriented suburbs may experience more varied or delayed responses to new transit access.

However, more recent empirical studies reveal that these premiums are neither universal nor stable over time. For instance, it was found that single-family homes near Minneapolis's light rail line initially show price increases, but these gains are highly sensitive to the choice of the control group and the definition of the post-treatment window, sometimes disappearing altogether (Pilgram & West, 2018). Pilgram and West used a difference-in-differences framework that compared price appreciation of properties near and farther away from stations. These findings suggest that appreciation due to station openings may be modest, temporary, or overshadowed by broader housing market trends. Similarly, LaRoche (2019) studied newly opened BART stations in suburban East Bay communities using a difference-in-differences approach by comparing treatment ZIP codes with neighboring control ZIP codes before and after new station openings. She reported mixed results, with some areas experiencing higher per-square-foot home values, while others showed no significant change. It is argued that in regions like the Bay Area, where many suburban zones already have multiple transportation options and mature urbanization, the marginal value of a new station may be limited. Existing literature also underscores the role of anticipation effects on housing prices. McDonald and Osuji (1995) use a regression model to isolate the effects of individual housing and neighborhood attributes on sale prices to show that residential land values in Chicago rose before the new transit service began because households capitalized expected future accessibility improvements into current prices.

While prior research provides important insights into how rail transit affects residential property values, most studies focus on dense, transit-mature metropolitan areas or suburban areas that are already integrated into older rail networks. For example, work on BART extensions in the Bay Area (LaRoche, 2019) and on the light rail in Minneapolis (Pilgram & West, 2018) evaluate systems where transit access is not new and where housing markets may have already capitalized transit benefits long before the studied extensions. These mixed and often modest effects imply that in mature systems, the marginal value of an additional station can be limited or difficult to isolate, helping to explain the variation in findings across these studies. In contrast, the Northern Virginia region, where Silver Line Phase II occurred, differs greatly, as it was newly introduced into previously auto-dependent suburbs that are still developing. Given that Phase II opened recently in late 2022, the market is still adjusting and there is limited post-treatment data for causal claims. For this reason, we utilize an exploratory, descriptive approach that examines residential property prices across distance bands, population density levels, and city-specific trends from 2018 to 2025. This helps establish early empirical patterns and identify whether these patterns resemble or diverge from those found in mature systems that have already been studied. Additionally, by focusing on Herndon, Reston, and Ashburn, areas that lacked rail service before 2022 and are undergoing significant demographic and land-use change, this research extends the literature to a new and understudied suburban context and provides information on how major transit expansions affect residential property markets outside of mature urban rail systems.

Methodology and Data

This research was an exploratory analysis on the relationship between zip code distances from new Silver Line stations and the changes in housing value; thus, no dense quantitative research was necessary for our analysis. From the data gathered, we constructed graphs based on changes in housing values, per zip code, to provide a visual understanding of the relationship between our variables. This enables us to assess discrepancies within our data. Data was collected from various publicly available sources, including Zillow (Zillow, 2025), US Zip Codes (Zipcodes, 2025), Data Commons (Data Commons, 2025), and World Population Review (Review, 2025), to gather zip code distances, housing values, and population changes.

The first step in our data collection process was identifying the station addresses of the new Silver Line expansion, including Herndon Station, Reston Town Center, Loudoun Gateway, and Ashburn. We elected to exclude the Dulles Airport and Innovation Center stops because they were located in two out of the three areas of interest, Sterling and Herndon, Virginia; thus, inclusion would be redundant. The next step was the identification and elimination of zip codes within our cities of interest: Herndon, Sterling, Ashburn, and Reston. It is imperative to note that Herndon and Reston are both part of Fairfax County, which inevitably results in some zip code overlap. The same applies to Ashburn and Reston, which are part of Loudoun County. However, such overlap does not complicate the data, as the zip codes are distinct and directly related to one of the four stations previously identified.

The zip codes were obtained from the United States Zip Codes website (Zipcodes, 2025), and the monthly household values were obtained from Zillow (Zillow, 2025). We initially began with fifteen different zip codes, but when inputting the information into Zillow, our zip code numbers decreased to twelve, given the fact that some of the listed zip codes were merely PO boxes with no retrievable data. From there, we documented the changes in the monthly Zillow Home Value Index for all homes with seasonal adjustment, from January 2018 until September 2025. After obtaining that data, we utilized the Google Maps measurement tool to measure the distance from each zip code to its respective station and categorized them: Less than or equal to one mile, less than or equal to two miles, less than or equal to three miles, and less than or equal to four miles. The colors themselves included light blue, dodger blue, gray blue, and dark blue. Dark blue represents zip codes that were less than or equal to four miles, and

light blue represents zip codes less than or equal to one mile.

In addition, there were difficulties in acquiring population data. Due to the government shutdown, typical table estimations obtained from the Census Bureau were not available to the public, inhibiting our ability to gather accurate estimates of population changes. To combat this, we compared annual population changes, by county, from two websites. One was Data Commons (Data Commons, 2025), which went up to 2023, and the other was from World Population Review (Review, 2025), which provided an estimation for the projected 2024 and 2025 data. In the event there were any discrepancies in the reported data from the years 2018 to 2023 in either source, we would average the results to obtain a rough estimation. From this point on, we used the information provided to graph the changes in housing prices per zip code relative to their distance from the stations of interest.

Analysis

The general plot follows this color pattern: Zip codes in dark blue are less than or equal to 4 miles, zip codes in gray blue are less than or equal to three miles, zip codes in dodger blue are less than or equal to 2 miles, and zip codes in light blue are less than or equal to one mile. In addition, we incorporated a gray line that indicates the month the silver line expansion was completed, which was November 2022. Contrary to our initial hypothesis, the data indicate that the majority of housing values for zip codes in closer proximity to the station are lower on average than those further away. This strongly contradicts previous literature that has shown the opposite effect occurring. To gain further understanding of the reasoning behind such differences, we inspected the data by city following a similar color pattern, except for precise mileage. In the event that two zip code distances were roughly the same value with a minor difference between the two, the one “furthest” away was given a dark turquoise color for visual aid. Furthermore, we included an orange line to indicate the mean housing value within the city per year.

Upon separation of our data isolated by city, we find that the housing prices amongst the closest zip codes are substantially lower than those further away. The two cities that demonstrate this best include Reston and Herndon, with their closest zip codes containing lower housing prices relative to the other zip codes within that area. In contrast, although Ashburn and Sterling data indicate that the housing prices in the closest zip code are lower relative to the furthest zip code, both serve as a median point amongst the other zip codes in their city. This result coincides with the average housing values per city represented by the orange line. Consequently, the relatively small differences in their prices relative to the average might be interpretable as the station line having a relatively insignificant effect on housing values within those areas.

These unexpected results, however, do contradict our initial hypothesis that the closer proximity would result in inflated housing values. Potential reasoning as to why this might be the case includes the housing types available within the respective cities. In Herndon, utilizing Zillow (Zillow, 2025) data, we find that there are approximately 618 rental spaces available in comparison to only 119 spaces available for purchase, which means over 80% of our housing data consists of apartments or other similar rental types. When narrowing by zip code, we find that the majority of homes for sale are located in 20170, our farthest zip code, whilst the apartments are closer to 20192. We follow this logic for other cities and find similar results. In Reston, there are approximately 882 rentals available, and only 112 are available for purchase. In Ashburn, there are approximately 646 available for rent, and only 261 are available for purchase. Lastly, in Sterling, there are 418 spaces available for rent and only 98 available for purchase. To summarize, over 80% (except Ashburn with 70%) of our housing data consists of apartments, and the apartments are concentrated within closer zip codes. This means that although the closer zip codes have lower housing values, they are not less populated. It can be theorized that to accommodate population needs within those areas, apartment complexes have become more popular, rather than typical housing units. To better assess our population data, we constructed a graph of the differences in population by city. We find a general upward trend, with Ashburn being the exception, but no visible spikes in population density. Therefore, we cannot further assess the cause behind the prevalence of apartment complexes relative to the population.

Shortcomings and Next Steps

There are several limitations to our research. Due to the government shutdown, we were unable to collect data from sources such as IPUMS or the Census Bureau, which limits the precision of our population estimates and overall findings. Additionally, our analysis does not take into account additional external factors that may influence housing prices independently of the Silver Line expansion, possibly resulting in omitted variable bias. Such factors include the presence of universities or major shopping centers, which could affect internal migration patterns within the city. To fully isolate those causal effects would require more in-depth analysis and data collection, which is currently unavailable.

Ultimately, while our study provides useful preliminary insights, the results should be interpreted with caution

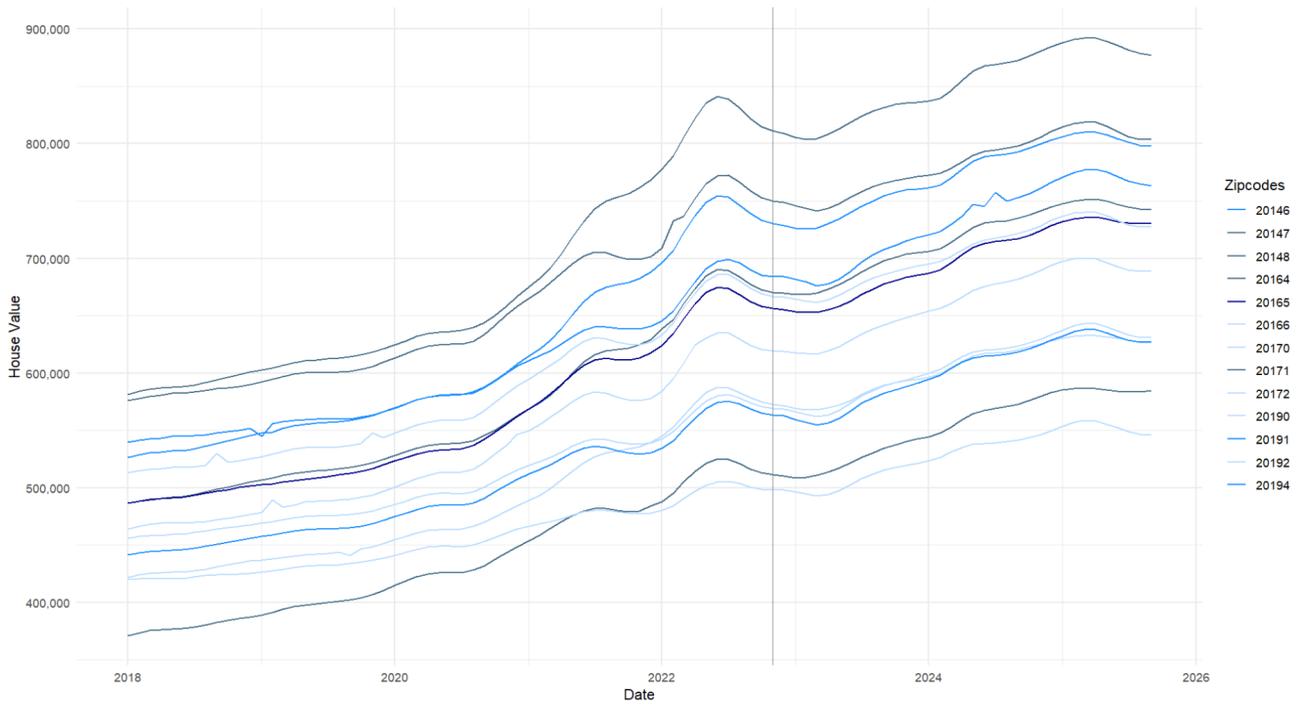


Figure 1: Changes in housing data by zip code.

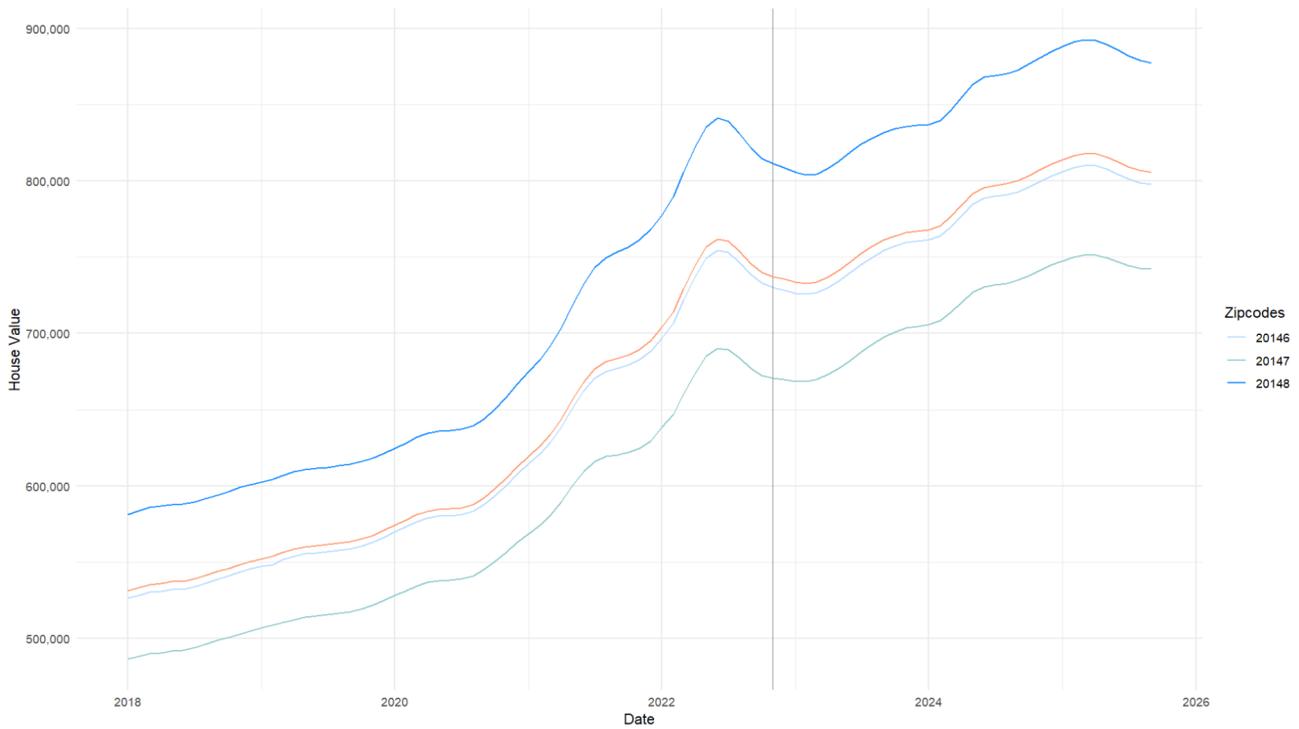


Figure 2: Ashburn Housing Values

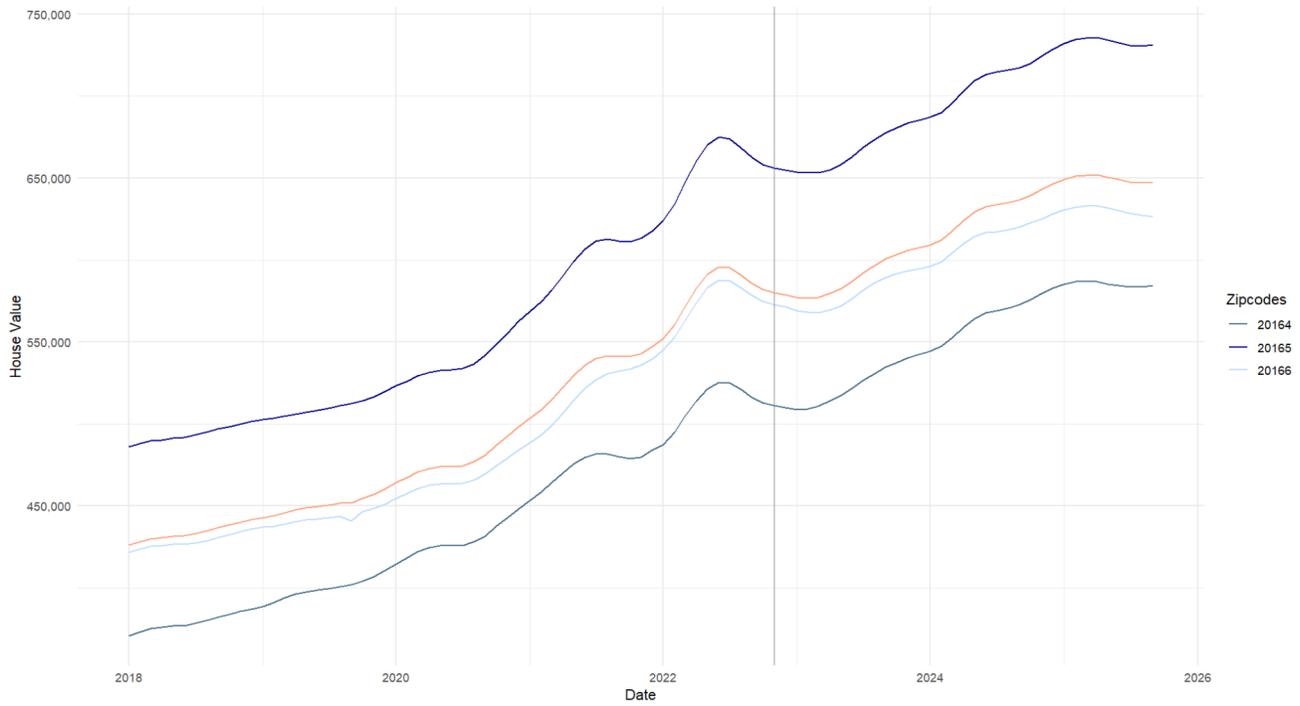


Figure 3: Sterling Housing Values

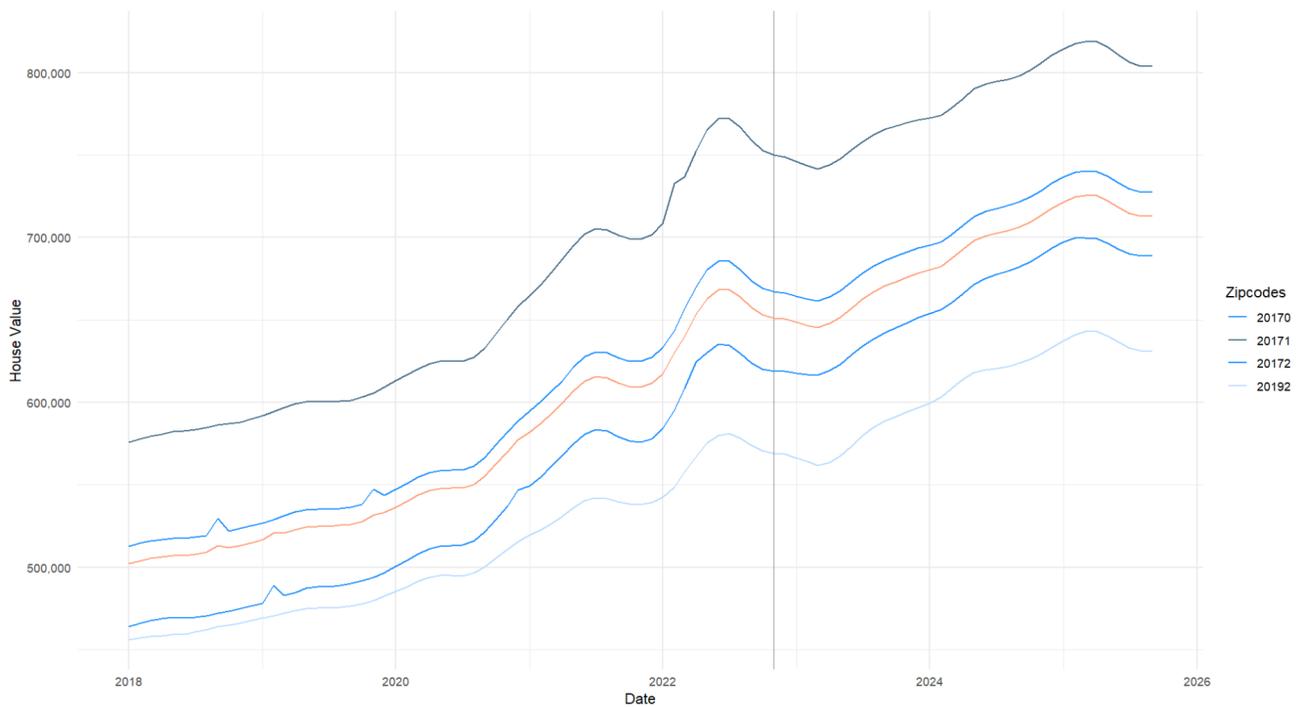


Figure 4: Herndon Housing Values

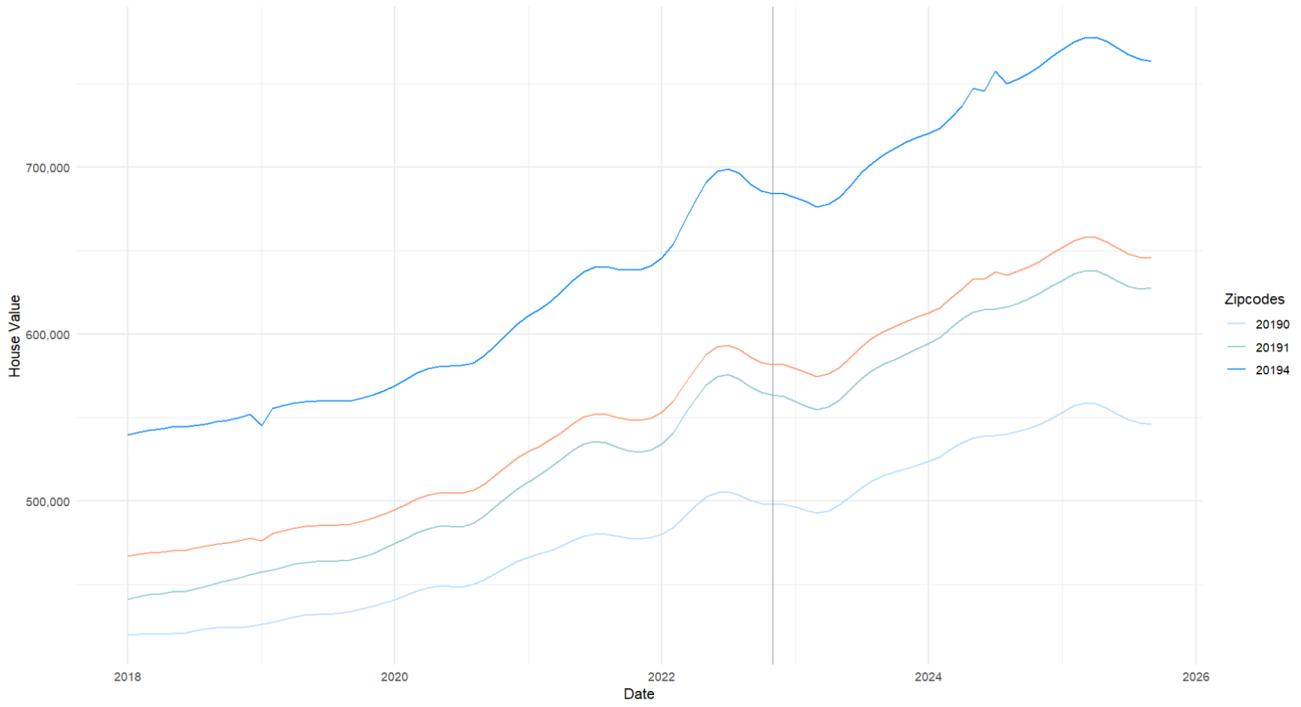


Figure 5: Reston Housing Values

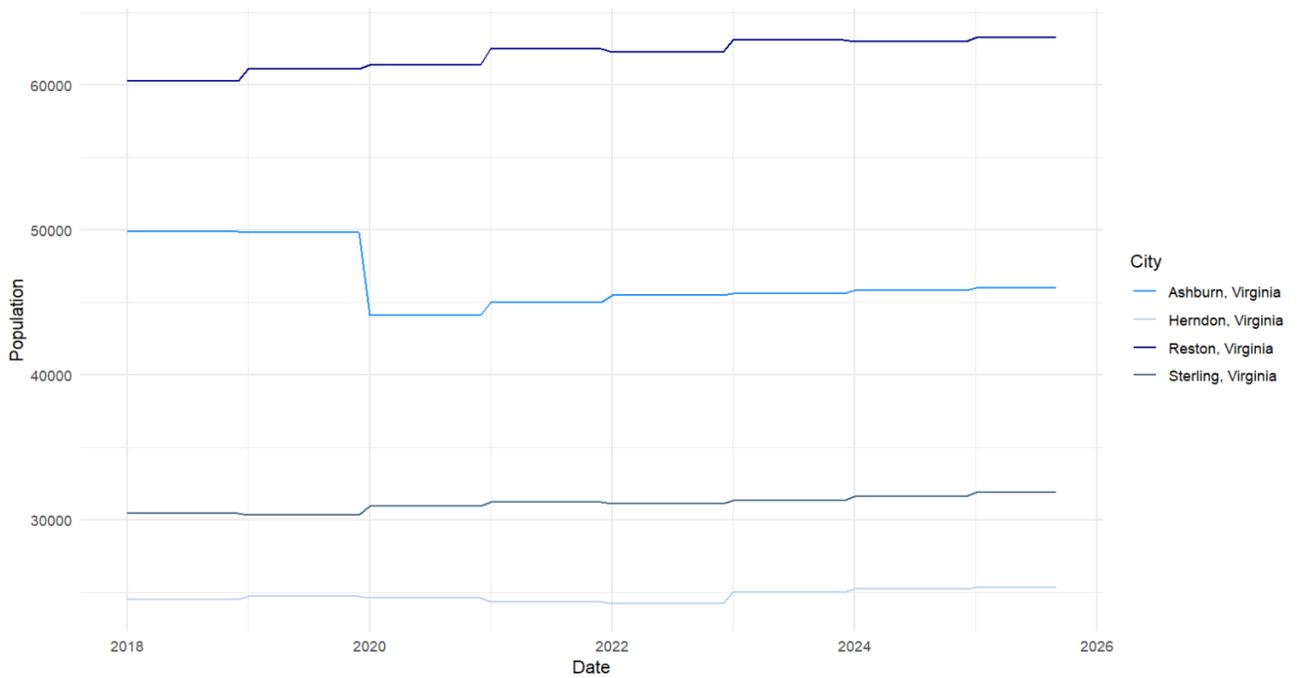


Figure 6: Population changes by City

until more comprehensive data becomes available. That being said, if we had additional resources, greater lengths would be taken to identify the reasoning behind such price changes. The first would be concrete data on population density by zip code, not by city. With that data, we would be able to analyze internal migratory patterns within the city that could serve as an explanation for our housing prices. The second would be the careful collection of housing and rental units sold and listed by zip code, per year. From there, we might be able to find a pattern; there could be a correlation between the increase in population density and the available apartment complexes, which would help our analysis. In fact, we hypothesize that the prevalence of apartment complexes is indicative of population changes, but our limited data does not enable us to accept or reject that theory.

There is much to discuss and dissect from our results, and we encourage other researchers to go further in depth to answer the questions we were unable to answer. Although there are constraints at the moment due to governmental changes, we sincerely believe that this project is worth looking into.

Conclusion

With the growth in concern regarding affordable housing amidst inflationary pressures and increased costs of living, it is imperative to assess how seemingly beneficial advancements, such as train line expansions, could impact the residents within that area due to changes in migratory patterns. Our initial hypothesis assumed that closer proximity to the stations of interest would inevitably lead to rising housing prices within that area, most likely due to population changes. However, our findings reveal the opposite: properties closer to the stations tend to have lower housing values.

This unexpected outcome suggests that additional factors may be influencing these patterns. When assessing the ratios between rental properties and spaces available for purchase, we find large discrepancies, with a higher concentration of apartment complexes and rental spaces located in zip code areas closer to the station. Consequently, this could indicate higher population density and provide evidence of migratory impacts from the silver line expansion. Therefore, it is encouraged that for further research, housing data is isolated by housing type, which would require more data collection, but is an arguably worthwhile endeavor.

Further research must be conducted to assess why that is the case. Potential explanations could include neighborhood sizes, housing regulations, government pricing controls, among many others that this paper does not delve into. From the information gathered, what can be acknowledged is that the silver line expansion might not have much of an impact on housing prices, but could affect other variables not analyzed, such as the creation of major shopping centers or universities.

While our results contradict initial expectations, they open the door for further research into the complex relationship between transit accessibility and housing markets. Although we are unable to make concrete conclusions, our findings suggest that affordability might not be the concern, but rather availability. There are significantly fewer housing units available compared to rental spaces, so although the cost of living near station areas might appear lower, it is not to be taken as evidence of affordability. Rather, it is encouraged that future research tracks changes in apartment prices over time rather than general data. Researchers should explore these relationships by using more granular data and controlling for additional variables to better understand the broader effects of transit development in emerging areas.

References

- Data Commons. (2025). *Data commons populations*. Retrieved 2025-11-23, from <https://datacommons.org/explore?q=population/>
- Debrezion, G., Pels, E., & Rietveld, P. (2007, August). The Impact of Railway Stations on Residential and Commercial Property Value: A Meta-analysis. *The Journal of Real Estate Finance and Economics*, 35(2), 161–180. Retrieved 2025-11-20, from <https://doi.org/10.1007/s11146-007-9032-z> doi: 10.1007/s11146-007-9032-z
- Gibbons, S., & Machin, S. (2005, January). Valuing rail access using transport innovations. *Journal of Urban Economics*, 57(1), 148–169. Retrieved 2025-11-23, from <https://www.sciencedirect.com/science/article/pii/S0094119004001020> doi: 10.1016/j.jue.2004.10.002
- Ingram, A. (2025, April). *Silver Line Extension to Ashburn: Real Estate Expectations vs. Reality*. Retrieved 2025-11-23, from <https://www.myhaven.homes/blog/silver-line-extension-to-ashburn-real-estate-expectations-vs-reality/>
- LaRoche, K. C. (2019). The Effect of New BART Stations on Housing Prices.
- McDonald, J. F., & Osuji, C. I. (1995, August). The effect of anticipated transportation improvements on residential land values. *Regional Science and Urban Economics*, 25(3). Retrieved 2025-11-23, from <https://trid.trb.org/View/462087>

- Mohammad, S. I., Graham, D. J., Melo, P. C., & Anderson, R. J. (2013, April). A meta-analysis of the impact of rail projects on land and property values. *Transportation Research Part A: Policy and Practice*, 50, 158–170. Retrieved 2025-11-23, from <https://www.sciencedirect.com/science/article/pii/S0965856413000207> doi: 10.1016/j.tra.2013.01.013
- Pilgram, C. A., & West, S. E. (2018, March). Fading premiums: The effect of light rail on residential property values in Minneapolis, Minnesota. *Regional Science and Urban Economics*, 69, 1–10. Retrieved 2025-11-23, from <https://www.sciencedirect.com/science/article/pii/S0166046217301886> doi: 10.1016/j.regsciurbeco.2017.12.008
- Rennert, L. (2022, September). A meta-analysis of the impact of rail stations on property values: Applying a transit planning lens. *Transportation Research Part A: Policy and Practice*, 163, 165–180. Retrieved 2025-11-20, from <https://www.sciencedirect.com/science/article/pii/S0965856422001677> doi: 10.1016/j.tra.2022.06.013
- Review, W. P. (2025). *World population review*. Retrieved 2025-11-23, from <https://worldpopulationreview.com/>
- Zillow. (2025). *Zillow housing index*. Retrieved 2025-11-23, from <https://www.zillow.com/home-values/102001/united-states/>
- Zipcodes, U. (2025). *United states zipcodes*. Retrieved 2025-11-23, from <https://www.unitedstateszipcodes.org/>

A Descriptive Analysis of the Post-COVID-19 Labor Market for Recent Graduates

Alexander Klautky

We examine the post-COVID-19 labor market for young college graduates in the United States using a CBSA-by-quarter panel from 2020–2024 combining CPS microdata, Indeed postings, OEWS-based Bartik shocks, and IPEDS BA completion flows. First, we document descriptive trends in unemployment, underemployment, young graduate relevant postings, vacancy rates, and BA supply, and plot a Beveridge curve for young graduates which closely mirrors the aggregate US labor market. Second, we estimate a log matching function for flows from unemployment to employment, finding a matching elasticity of 0.36 with respect to unemployment. Third, we use a Bartik instrumented postings as a measure of local labor demand in our 2SLS models for unemployment and underemployment. Despite strong first stages, estimates were small, negative, and statistically insignificant suggesting that either our specifications suffered from omitted variable bias or measurement error and failed to pick up the true effect, or the true effect is small, and rising unemployment is better explained by wage adjustment, job quality, or career trajectory.

Introduction

SINCE the post COVID-19 normalization of unemployment in 2022, the labor market for young graduates has been a fiercely debated topic among students, recruiters and political commentators (U.S. Bureau of Labor Statistics n.d.; Smith 2023). Students and young graduates' express frustration over high underemployment, low wages, degree mismatch, and the rising number of applications needed to secure one return offer (Smith 2023; Abel, Deitz, and Su 2014; Bell and Blanchflower 2011). Young graduates express concern over the role of AI in job displacement and in the hiring process (Smith, Johnson, et al. 2024; Lee and Williams 2023). On the other hand, 80 percent of employers feel that young graduates lack problem solving skills and 65 percent believe that universities fail to prepare students to meet industry needs (Sutherland and Hughes 2024; Rodriguez and Gray 2023). Employers and students alike report using AI to aid in applying for and filtering applications respectively (Lee and Williams 2023; Smith, Johnson, et al. 2024). Some political commentators blame universities and question the value of a college degree while others note low wages and high costs of living for young graduates (Smith 2023; Abel, Deitz, and Su 2014; Bell and Blanchflower 2011). Beyond social discourse, there is growing concern among economists that individuals who graduated during or right after the COVID-19 pandemic may suffer from life-long cyclical scarring like cohorts during the Great Recession of 2008 (Kahn 2010; Oreopoulos, von Wachter, and Heisz 2012; Hughes and Carter 2024).

Considering the growing social and academic concern over the topic, our study aims to examine the condition of the post-COVID-19 labor market for young graduates by utilizing multiple linear regression and Bartik instruments to estimate job matching efficiency and the effect of labor demand on unemployment (Blanchard and Diamond 1989; Barnichon and Figura 2015; Bartik 1991; Goldsmith-Pinkham, Sorkin, and Swift 2020; Hershbein and Kahn 2018). In consonance with research on the labor market during the 2008 recession, we hypothesize that demand has a positive effect on both unemployment and underemployment and expect matching efficiency to decrease (Kahn 2010; Bell and Blanchflower 2011; Abel, Deitz, and Su 2014; Barnichon and Figura 2015; Oreopoulos, von Wachter, and Heisz 2012; Hughes and Carter 2024).

Literature Review

We base the framework of our analysis on the search and matching model of unemployment (job findings vs separations) and the Beveridge curve relating vacancies to unemployment (Blanchard and Diamond 1989; Barnichon and Figura 2015). Movement along the Beveridge curve is classically interpreted as a change in aggregate labor demand while shifting of the curve is attributed to changes in matching efficiency (Barnichon et al. 2012; Şahin et al. 2014). Recent estimates of aggregate matching functions show that the residual matching efficiency is heavily cyclical and falls during periods of recession implying that shifts in matching efficiency can account for a significant share of unemployment variation (Barnichon and Figura 2015). Related work shows that entering the labor force during a recession is associated with persistent scarring effects on earnings and job quality (Kahn 2010; Oreopoulos, von Wachter, and Heisz 2012). Even outside of recessions, young graduates experience a higher rate of mismatch and underemployment relative to the aggregate labor force.

A separate literature studies the measurement of labor demand and the effect of exogenous demand shocks on unemployment. Vacancy data and online job postings data are now widely accepted and frequently used proxies for local labor demand to measure the impact of shocks on the skill composition of available jobs (Hershbein and Kahn 2018). In parallel, shift-share (Bartik) instruments use predetermined local industry shares interacted with national growth rates to construct exogenous local demand shifters. We base our construction of Bartik instruments on recent work which clarifies the identification and assumptions implicit to the method (Bartik 1991; Goldsmith-Pinkham, Sorkin, and Swift 2020). The labor demand studies and the matching model studies primarily focus on the aggregate labor market, rather than the labor market for young graduates, and rarely include matching efficiency or demand-unemployment models for young graduates (Barnichon and Figura 2015; Abel, Deitz, and Su 2014). In our relatively short time frame (2020-2024), we follow a common assumption of graduation-in-recession literature of treating local labor supply flow as slow moving, predetermined, and weakly associated with cyclical shocks (Kahn 2010; Oreopoulos, von Wachter, and Heisz 2012). Therefore, we exclude the improbable association between local labor supply and unemployment from our broader framework. Instead, we focus on the primary shifters of the Beveridge curve, exogenous demand shocks and matching efficiency.

Research Methods

Panel Structure

Our analysis is conducted using a core panel of CBSA \times Year \times Quarter which we construct in section 4. The core panel contains variables for unemployment, underemployment, jobs finding, postings, vacancy rate, Bartik index, a JOLTS control, and BA flows on a common set of CBSAs over time. Let i denote CBSA codes and q calendar quarters; the core panel spans 261 CBSA codes across 20 quarters (2020-2024). All variables were constructed independently then left joined into the core panel dropping all rows with missing CBSA codes and rows with insufficient data for quarters yielding a balanced panel that has strong coverage across all variables.

Estimating Demand on Unemployment and Underemployment

To calculate the association between labor demand and unemployment, we estimate a two-stage least squares model (2SLS) where the outcome is the unemployment rate for young graduates and the endogenous regressor is our cumulative postings-based demand measure. Let $u_{iqt}^{\text{BA},22-27}$ denote the unemployment rate for young graduates in CBSA i and quarter (q, t) and $u_{iqt}^{\text{nonBA},22-27}$ the unemployment rate for young individuals without a BA in CBSA. Let $Dsum_{iqt}$ denote the cumulative BA postings demand measure equal to the sum of the four quarter lags. Our main 2SLS specification is:

$$u_{iqt}^{\text{BA},22-27} = \alpha_i + \lambda_{qt} + \beta Dsum_{iqt} + \theta u_{iqt}^{\text{nonBA},22-27} + \varepsilon_{iqt}$$

with $Dsum_{iqt}$ instrumented using four lags of the Bartik demand shifter:

$$Dsum_{iqt} = \alpha_i + \lambda_{qt} + \delta u_{iqt}^{\text{nonBA},22-27} + \pi_1 E_{1,iqt} + \pi_2 E_{2,iqt} + \pi_3 E_{3,iqt} + \pi_4 E_{4,iqt} + \nu_{iqt}$$

where $E_{\ell,iqt} = \text{lag}_{\ell}(\text{Bartik}_{it})$ for $\ell = 1 \dots 4$, and α_i and λ_{qt} denote CBSA and quarter by year fixed effects in both stages, standard errors are clustered at the CBSA level. We expect the Bartik index to be a relevant for postings demand because it interacts predetermined occupational shares with national occupation growth shocks, so CBSAs that are more specialized in nationally expanding occupations should experience associated changes in vacancies and postings. The exclusion restriction is plausible because CBSA fixed effects, time fixed effects, and non-BA controls are used, however, variables outside the scope of our analysis like wages or job quality may associated with the Bartik instrument and unemployment. The main specification is redefined to estimate two placebos, one for individuals with a BA aged 27–35 and another for individuals aged 22–27 without a BA where $u_{iqt}^{\text{nonBA},22-27}$ is omitted as a control. To obtain the main estimate and placebo estimates for underemployment, we replace the dependent variable $u_{iqt}^{\text{BA},22-27}$ with the variable for underemployment UnderEmp_{iqt} while keeping $Dsum_{iqt}$ as the endogenous regressor instrumented with the four lags of the Bartik index. The underemployment estimates are computed using the same specification for unemployment with a control for aggregate underemployment minus the placebos.

Estimating the Matching Function and Matching Efficiency

To estimate the matching function, we use a log matching function on CBSA \times quarter flows and unemployment stock for young graduates. For each CBSA i and quarter q , let M_{iq} be the CPS weighted amount of unemployment

to employment status transitions (matches) for young graduates. Let U_{iq} be the CPS weighted stock of unemployed young graduates at the start of a given quarter. We estimate the fixed effects regression:

$$\log M_{iq} = \lambda_i + \tau_q + \alpha \log U_{iq} + \varepsilon_{iq}$$

where λ_i is CBSA fixed effects, τ_q is quarter by year fixed effects, and α is the elasticity of matches in terms of unemployment. Standard errors are clustered at the CBSA level. We extract the estimated time fixed effects τ_q and interpret them as $\log A_q$ to construct quarter level relative matching efficiency:

$$A_q^{\text{rel}} = \exp(\tau_q - \bar{\tau})$$

Data

Unemployment Rate

The primary outcome is the quarterly unemployment rates for recent graduates. We started by downloading CPS monthly microdata files from the US census website, between 2020 Q1 and 2024 Q4. We restricted the CPS sample to individuals who were aged 22-27, with a BA, and years 2020-2024 and drop cells with non-positive person weights. When CBSA codes were available, cells were grouped into unique CBSA codes, when unavailable a variable was generated from state and county codes and merged into CBSA groups via an external county-CBSA crosswalk. CBSA level totals for each CBSA code, year, and month are created by summing person weights over individuals in the labor force and over individuals unemployed to compute monthly unemployment rates. CBSA monthly totals are then aggregated to calendar quarters by summing the weighted CBSA cells. This yields a CBSA x quarter x year panel with unemployment rates of young graduates and the associated quarterly labor force weights which are merged into the core panel. A control for aggregate unemployment was constructed in the same manner and joined to the core panel.

Underemployment Rate

Variables for the occupational BA-share and underemployment were created to measure how many young college graduates are working in jobs that do not require a BA for each CBSA i and month m . Drawing on the CPS monthly microdata, we restrict the sample to individuals who have a BA, positive person weights, valid CBSA codes, and are aged 22 to 27. The level of the construction thus far is the reference month m where each observation is on worker j , in CBSA i , with CPS weights w_{jm} , and detailed occupation code occ .

We define how BA-heavy each occupation is in the overall labor market using CPS observations for all individuals in the labor force. We computed for each occupation o the CPS weighted share of workers with a BA or higher and define the BA share of occupation o as:

$$\text{BAshare}_o = \frac{\sum_{j \in o} w_j \mathbf{1}\{\text{BA}+\}}{\sum_{j \in o} w_j}$$

The index is bounded between 0 and 1, with a small score representing a lower BA-share and a larger score representing a higher BA-share for a given occupation. The BA share index is then mapped back onto a new young graduates dataset, giving each young graduate a score for their occupation. Occupations such as retail and food services scored low on the scale while occupations such as banking and software engineering had high scores.

We define underemployment as the share of employed young graduates working in occupations where most workers do not have a BA at the CBSA month level. Let τ be the cutoff for a BA job. We set τ to 0.6 so that occupations where $\text{BAshare}_o < 0.6$ are classified as non BA occupations and occupations where $\text{BAshare}_o > 0.6$ are classified as BA occupations. For each CBSA i and month m , we compute the CPS weighted underemployment rate for young graduates:

$$\text{UnderEmp}_{im} = \frac{\sum_{j \in (i,m)} w_{jm} \mathbf{1}\{\text{BAshare}_{o(j)} < \tau\}}{\sum_{j \in (i,m)} w_{jm}}$$

The numerator counts, in CPS weighted units, the number of employed young graduates in CBSA i and month m working in non-BA occupations while the denominator counts the total CPS weighted number of young workers within a CBSA. The resulting variable UnderEmp_{im} is a rate between zero and one that measures the rate of underemployment among recent graduates in each labor market for each month. The underemployment rate is then joined into the core panel by CBSA x year x quarter. An aggregate underemployment control was constructed in the same way and merged to the core panel.

Job-Finding Rate

Variables for unemployment stock U_{iq} and flow from unemployment to employment M_{iq} were created to calculate the job finding rate (f_{iq}) for young graduates. Drawing on the CPS monthly microdata, we restrict the sample to individuals who are aged 22-27, with at least a bachelor's degree, and have positive person weights. Individuals are identified across months using the CPS household and person identifiers (`hrhhid`, `pulinen0`) while their labor force status is recorded as employed, unemployed, or not in the labor force. For each individual observed in consecutive months we identify an unemployment to employment transition where an individual is unemployed in month $m - 1$ and employed in month m and remains in the same CBSA. The monthly flow of matches, measured in levels, represents the CPS-weighted number of young graduates in CBSA i who move from unemployment to employment between month $m - 1$ and m .

$$M_{im} = \sum_j w_{j,m-1} \mathbf{1}\{U \rightarrow E \text{ in } i\}.$$

The corresponding unemployment stock is the CPS-weighted number of unemployed young graduates for any given CBSA i at the start of each month:

$$U_{im} = \sum_j w_{j,m-1} \mathbf{1}\{U \text{ in } i\}.$$

The monthly measures for M_{im} and U_{im} are aggregated to the CBSA \times quarter level by summing both variables over months within each calendar quarter yielding M_{iq} and U_{iq} respectively. M_{iq} represents the number of matches for young graduates in CBSA i during quarter q while U_{iq} represents the size of the unemployed risk set at the start of those months for any given CBSA. The job finding rate for young graduates in CBSA i and quarter q is defined as:

$$f_{iq} = \frac{M_{iq}}{U_{iq}}.$$

f_{iq} measures the fraction of unemployed young graduates whose status changed to employed within the quarter. The job finding rate f_{iq} for each CBSA \times year \times quarter is merged into the core panel.

Job Postings Index

We construct a job postings index by using the daily metro data from Indeed's Hiring Lab job postings tracker. The raw data reports a measure of job postings expressed as a percent deviation from a pre-COVID-19 level J_{id} for each CBSA i on day d . We transform this percent difference measure into a relative level by computing L_{id} :

$$L_{id} = 1 + \frac{J_{id}}{100}.$$

L_{id} is averaged across quarters and CBSA to create a quarterly average level P_{iq} for postings:

$$P_{iq} = \frac{1}{D_{iq}} \sum_{d \in q} L_{id},$$

where D_{iq} is the number of daily observations available for any given CBSA i in quarter q . Values constructed with less than 30 days or missing CBSA ids are dropped. P_{iq} is then log transformed and the resulting CBSA \times quarter \times year postings file is joined with the core panel.

Bartik Index

We calculate the CBSA \times year \times quarter Bartik index from the raw BLS Occupational Employment and Wages Statistics (OEWS) data. For each discrete occupation o , CBSA i and calendar year t , let E_{iot} denote OEWS employment for any given occupation, CBSA, and year. First, we summed all employment counts across all OEWS rows that share the same CBSA, occupation, and year. Using a fixed baseline at $t_0 = 2020$, we computed the CBSA \times occupation employment shares:

$$s_{io}^{(0)} = \frac{E_{iot_0}}{\sum_{o'} E_{io't_0}}$$

So that the sum over all occupations o of $s_{io}^{(0)} = 1$ for each CBSA. Then we aggregated the new table to the national level. E_{ot}^{US} denotes the total U.S. employment in occupation o in year t and was obtained by summing state-level employment. For each occupation \times year cell we compute the annual log change in national employment:

$$g_{ot} = \log E_{ot}^{US} - \log E_{o,t-1}^{US}$$

All missing first differences were set to 0. The Bartik variable for CBSA i in year t is a share-weighted average of the national occupation-specific shocks:

$$B_{it} = \sum_o s_{io}^{(0)} \cdot g_{ot}$$

B_{it} can be interpreted as an index of predicted growth for local labor demand driven purely by national shifts in occupational demand and predetermined local occupational shares. Finally, we assign the value of each year to each quarter within a year and merge the resulting `Bartik_entry` variable to the core panel.

Vacancy Rate

We constructed a proxy for the CBSA \times quarter vacancy facing young graduates by combining our job postings panel with our CPS measure of BA employment. Starting with postings, we extract the total number of active postings P_{iq} from each CBSA i and calendar quarter q . We use CPS microdata aggregated to the CBSA \times quarter level for young graduates to obtain our measure for the CPS-weighted size of the labor force LA_{iq}^{BA} and the CPS-weighted number of unemployed workers U_{iq}^{BA} . We merge the two along common sets of CBSA code, year, and quarter. We then compute the implied employment stock for young graduates in CBSA i during quarter q :

$$E_{iq}^{BA} = LA_{iq}^{BA} - U_{iq}^{BA}.$$

E_{iq}^{BA} counts CPS-weighted employed young graduates. We define the vacancy rate as the share of jobs that are unfilled, where postings are treated as open vacancies and employed workers as filled positions. For each CBSA i and quarter q we compute:

$$v_{iq}^{BA} = \frac{P_{iq}}{P_{iq} + E_{iq}^{BA}}.$$

We drop all cells where $P_{iq} + E_{iq}^{BA} = 0$. The resulting variable v_{iq}^{BA} is our CBSA \times quarter measure of the vacancy rate for young graduates. v_{iq}^{BA} was joined into the core panel.

Check if any other variables need to be described in data. Vacancy rate.

BA Flows IPEDS

We constructed a flow variable, `flow_ba_it`, to measure the new supply of recent graduates entering each CBSA. Drawing on IPEDS, we restrict the sample to individuals aged 22–27 between the years of 2020 and 2024. The sample is further restricted to only include cells that have institutional identifiers and a completions count. We only keep cells that have a completions count which corresponds to BA-associated indicators (`credential_level`, `creddesc`, or `awardlevel_desc`). Because indicators for BA change across the sample, searching for multiple indicators is necessary. We then only keep four columns: institutional identifier (`unitid`), calendar year (`_year`), zip code (`_zip`), and the chosen completions count (`_n`). Rows missing `_year` are dropped, and rows with missing or negative values of `_n` are dropped. For each institution-by-year, we collapse differences across program, sex, and other variables by summing `_n`, yielding the total count of bachelor's completions per institution per year.

A fourth variable, `cbasa`, is created to map institutions onto their corresponding CBSA. A new file lists institutions and their corresponding zip codes. Using the crosswalk bridge created in Section 2.1, five-digit CBSA codes are generated for each institution's five-digit zip code. The resulting zip-code–CBSA crosswalk is left joined to the institution-by-year file. The zip-code variable is then dropped, and any rows missing `cbasa` are dropped.

Finally, we aggregate completions to the CBSA-by-year level by summing `_n` for each institution j across each CBSA i within each year t . For each CBSA i and year t , let n_{jt} denote the total number of bachelor's degrees completed at institution j in year t , and let $\mathcal{J}(i)$ be the set of institutions located in CBSA i . The BA supply measure `flow_ba_it` is then

$$\text{flow_ba_it} = \sum_{j \in \mathcal{J}(i)} n_{jt},$$

the total number of bachelor's degrees awarded by institutions in CBSA i in calendar year t . The final data set is left joined to the core panel preserving CBSAs and their associated BA flows for each year.

Analysis

Descriptive Statistics and the Beveridge Curve

1, is constructed by pulling cells from the young graduate unemployment rate, underemployment rate, vacancy rate, BA postings index, and BA flows aggregated across each CBSA and quarter. Variation is measured as the percentage

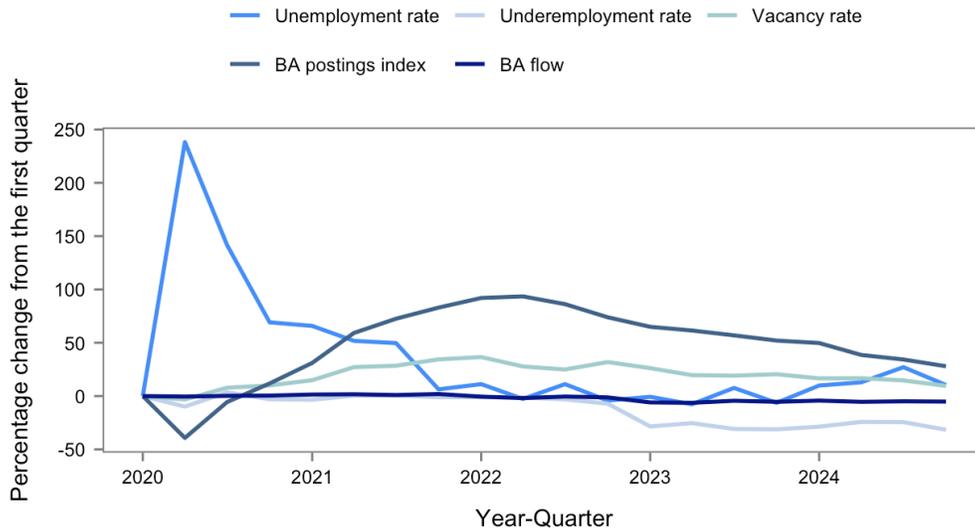


Figure 1: Evolution of unemployment, underemployment, vacancy rate, postings index, and BA supply for recent graduates

change from the first quarter of 2020. The unemployment rate spikes during 2020 indicating the COVID-19 labor market shock and subsequently normalizes during 2022. After 2022 the unemployment rate has very slowly but steadily increased to five percent at the end of 2024. The BA postings index dips during 2020 recovering to a peak in Q2 2022 and declines steadily afterwards. The vacancy rate dips slightly during 2020 and increases between 2020 and 2022 and then slowly decreases. The BA supply flow remains nearly flat while underemployment is initially flat but begins to decrease during Q2 2022. Note the inverse relationship between the vacancy rate and unemployment rate; to better understand this relationship we plot a Beveridge curve for young graduates.

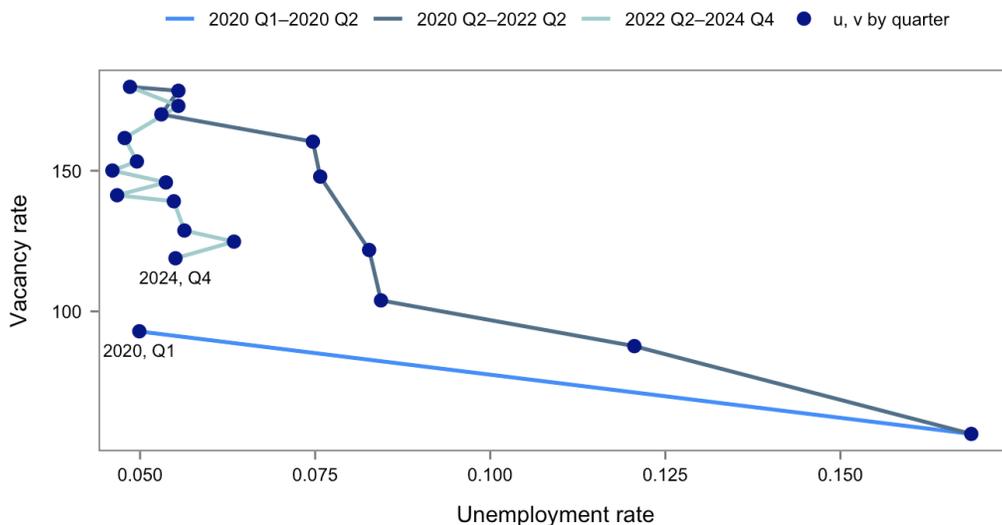


Figure 2: The Beveridge Curve for Young Graduates

2, We constructed the Beveridge curve for young graduates by mapping vacancy and unemployment rates across the 5-year period per quarter. The curve exhibits three phases 1. Q1-2020- Q2-2020, COVID shock, there is movement along the curve outwards as the recession decreases vacancies and increases unemployment. 2. Q2-2020 - Q2-2022, post COVID shock, the curve steepens as unemployment decreases and vacancies increase. 3. Q2 - 2022 - Q4-2024, post normalization the curve becomes unusually steep as vacancies rapidly decrease while unemployment increases slowly. The three selected phases for the curve are the only possible periods that exhibit a monotonic increase or decrease in vacancies across time. The first phase aligns with the spike in unemployment and the drop in vacancies between Q1-2020 and Q2-2020. The second phase begins when unemployment and vacancies begin to increase and ends when unemployment is normalized and vacancies peak. Across the third phase vacancies fall and

unemployment slowly rises. The curve is downward sloping across all three phases and progressively becomes steeper. This indicates that a decrease in the postings index for recent graduates is associated with a smaller increase in unemployment rate over time. This suggests but does not prove that the stress of decreasing vacancies is taken up by other factors like wages or job downgrading. Our Beveridge curve for young graduates is consistent with the BLS's Beveridge curve for broader labor force. This suggests that the shifting relationship between vacancies and unemployment for recent graduates is not dissimilar to the shifting relationship for the entire labor force.

Matching Efficiency

We estimate the relationship between log unemployment stock and log matching stock using the regression model for matching described in section 3.3. Table 1 presents the results of the regression used to estimate the elasticity of matches with respect to unemployment within the matching function.

Table 1: Matching Function Regression Results

	$\log M_{iq}$
$\log U_{iq}$	0.359*** (0.034)
Observations	714
Adjusted R-squared	0.681

Notes: Standard errors are in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Our estimate suggests that within any given CBSA and quarter, a one percent increase in unemployment stock U_{iq} is associated with a 0.36 increase in the number of transitions from employed to unemployed M_{iq} . Because the elasticity is significantly below one, cells with higher unemployment tend to have lower job finding rates M_{iq}/U_{iq} regardless of matches in levels. Together, this estimate and the Beveridge curve indicate that markets with more unemployed graduates tend to be associated with more matches and lower job findings rates which is predicted by the matching function and the search and matching model.

Across 2020–2024, the quarter level matching efficiency estimate A_q^{rel} hovered around 1 with swings of roughly $\pm 20\%$ indicating cyclical variation rather than a large macro collapse or boom in matching efficiency. Compared with the Beveridge curve, our matching efficiency estimates suggests that the dramatic movement along and slope changes of the Beveridge curve are more reflective of changes in demand and unemployment than a deterioration of matching efficiency for young graduates. The small decreases in efficiency line-up with the steep decline in vacancy rates observed from 2023–2024 but this small decrease is not sufficient to explain a significant part of the weakening association between vacancy rates and unemployment rates.

The Effect of Demand on Unemployment and Underemployment

We now turn to the 2SLS estimates of the effect of posting-based labor demand on unemployment and underemployment for recent graduates. The estimates were computed from the framework described in section 3.2: the endogenous regressor is the cumulative demand postings index $Dsum_{igt}$ instrumented with four lags of the Bartik index E1 through E4. CBSA and quarter-by-year fixed effects are used and, for unemployment models, the unemployment rate of young individuals without a BA is a control. Table 2 presents the 2SLS estimates for unemployment among young graduates:

In table 2 for the main specification (1), the coefficient on cumulative demand is small, negative, and statistically insignificant: -0.0023 with a standard error of 0.0015 . The first stage results provide good evidence for instrument relevance: the coefficients on the four Bartik lags range from 19.7 to 23.1 and are highly significant while the joint F statistic on E1 through E4 is 123.5 which is far above the conventional weak instrument threshold of 10. The Hansen J test statistic of 0.245 (insignificant) indicates that for this specification and dataset we failed to reject the null hypothesis of instrument validity, meaning that there is no evidence of over-identifying restrictions. This is consistent with the exclusion assumption that occupational Bartik shocks affect young graduate unemployment only through posting-based demand. These results taken literally suggest that changes in our postings-based index are associated with small and meaningless changes in unemployment for young graduates. However, these estimates could also be the result of either omitted variable bias from failing to control for a local labor market factor that impact both the job postings index and unemployment, or measurement error in the postings index itself.

Table 2: IV coefficients on cumulative demand main specification and placebos

	(1) BA 22–27 U	(2) Placebo BA 28–35 U	(3) Non-BA 22–27 U
Cumulative demand	−0.0023 (0.0015)	−0.0010 (0.0026)	−0.0094* (0.0049)
Hansen J test	0.245	9.551**	8.492**
First-stage coefficients on Bartik instrument			
Lag 1 Bartik index (E1)	23.088*** (3.122)	23.088*** (3.122)	22.892*** (3.172)
Lag 2 Bartik index (E2)	28.353*** (1.674)	28.353*** (1.674)	28.171*** (1.686)
Lag 3 Bartik index (E3)	25.592*** (1.379)	25.592*** (1.379)	25.614*** (1.379)
Lag 4 Bartik index (E4)	19.735*** (1.351)	19.735*** (1.351)	19.697*** (1.362)
Joint F-stat: E1–E4	123.598***	123.598***	120.932***

Notes: Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Despite a strong first stage, the coefficients on cumulative demand for the placebo specification (2) is small, negative, and insignificant while the coefficient on cumulative demand for the placebo specification (3) is moderately small, negative, and marginally significant. The results for model 3 suggest that the effect is actually larger for young individuals without a BA those with a BA, however, we should not interpret this literally because the Hansen J test for both placebo specifications are significant indicating that over-identifying restrictions were violated. Together, the placebo results imply that our identification assumptions fare better for the actual sample, young graduates.

Table 3: IV coefficients on cumulative demand main specification and placebos

	BA 22–27 Underemp
Cumulative demand	−0.010 (0.012)
Hansen J test	2.810
First-stage coefficients on Bartik instrument	
Lag 1 Bartik index (E1)	23.286 (3.140)
Lag 2 Bartik index (E2)	28.233 (1.748)
Lag 3 Bartik index (E3)	25.508 (1.333)
Lag 4 Bartik index (E4)	20.019 (1.320)
Joint F-stat: E1–E4	121.842

Notes: Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 presents the estimate for underemployment among young graduates. Here, the coefficient on the cumulative demand -0.010 is small, negative and statistically insignificant. The first stage was strong with coefficients for the four Bartik indexes ranging from 20 to 23.3 while the joint F statistic was 121.8. In addition, the coefficient on the Hansen J test was statistically insignificant 2.810 indicating that we failed to reject the null hypothesis of instrument validity. The combination of strong first-stage relevance and weak second stage results suggest that measurement error in the underemployment index, limited power, and the short window may have contributed to the lack of a significant effect.

Shortcomings

While we believe that our study will help shape a broader understanding of the labor market for young undergraduates, we faced significant limitations in a few key areas. Observing effects across CBSAs is challenging for this demographic because young graduates are a very mobile group. Young graduates' mobility directly hindered our ability to accurately observe the relationship between labor demand and unemployment across CBSAs. If young graduates leave weak markets for stronger markets, mobility will bias our estimate towards zero causing us to underestimate

the true relationship between labor demand and unemployment. Another major limitation is possible underlying variation in our categorization of a BA dominated occupation. The true category of an occupations could have shifted causing us to treat some BA occupations as non-BA occupations and vice versa. This likely caused measurement error distorting our estimated relationship between labor demand and unemployment.

Conclusion

Overall, this paper provides an integrated descriptive and quasi-experimental view of the post-COVID-19 labor market for recent college graduates. Using a CBSA x year x quarter panel that combines CPS microdata, Indeed postings, OEWS-based Bartik shocks, and IPEDS completion flows, we document labor market indicators. Using these indicators, we plot a postings rate to unemployment rate Beveridge curve for young graduates which closely tracks the aggregate US curve: sharp outward movement during the covid shock, rapid normalization through 2022, and steepening from 2022 through 2024. Over the same period, underemployment trends downwards while supply flows remain roughly flat. Our matching function estimates imply an elasticity of matches with respect to unemployment of 0.36 and a quarter level matching index which fluctuates ± 20 percent around its mean and suggests no evidence of any large, continued collapse in the efficiency that unemployed young graduates are matched to available jobs.

Relative to prior literature (Kahn 2010; Oreopoulos, von Wachter, and Heisz 2012), we find no observable effect of labor demand on unemployment or underemployment. In our main 2SLS specification, postings-based cumulative demand had a small, negative, and statistically insignificant effect on unemployment and underemployment despite strong first stage results and over-identification tests that supported our instruments validity. Given young graduates high geographic mobility, and the possibility of shifting categories of BA and non-BA occupations, we suspect that there is potential for measurement error in both the postings variable and underemployment measure. A cautious reading of our estimates suggest that within our time frame and identification strategy, COVID-19 associated demand shocks do not appear to generate large or precise estimated change in unemployment or underemployment for recent graduates. This implies that shifts in wages, job quality, or career trajectory could better explain rising unemployment in recent graduates. We conclude that a more refined estimate for the effect of labor demand on unemployment is needed, and suggest that future research incorporates wages, job quality, and career trajectories into the model.

References

- Abel, J. R., Deitz, R., & Su, Y. (2014). Are recent college graduates finding good jobs? *Current Issues in Economics and Finance*, 20(1), 1–8.
- Barnichon, R., & Figura, A. (2015). Labor market heterogeneity and the aggregate matching function. *American Economic Journal: Macroeconomics*, 7(4), 222–249.
- Bartik, T. J. (1991). *Who benefits from state and local economic development policies?* W. E. Upjohn Institute for Employment Research.
- Bell, D. N. F., & Blanchflower, D. G. (2011). Underemployment in the UK in the Great Recession. *National Institute Economic Review*, 215(1), R23–R33.
- Blanchard, O. J., & Diamond, P. (1989). The beveridge curve. *Brookings Papers on Economic Activity*, 20(1), 1–76.
- Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8), 2586–2624.
- Hershbein, B. J., & Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7), 1737–1772.
- Hughes, I. M., & Carter, S. R. (2024). The long-term effects of COVID-19 on graduate employment: Cyclical scarring and job market disruptions. *Journal of Labor Economics*, 42(1), 1–25. doi: 10.1086/jle.2024.59842
- Kahn, L. B. (2010). The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics*, 17(2), 303–316.
- Lee, R., & Williams, T. B. (2023). The growing role of AI in recruitment: How job seekers and employers use technology in hiring. *Journal of Human Resource Management*, 55(3), 180–200. doi: 10.1086/jhrm.2023.459833
- Oreopoulos, P., von Wachter, T., & Heisz, A. (2012). The short- and long-term career effects of graduating in a recession. *American Economic Journal: Applied Economics*, 4(1), 1–29.
- Rodriguez, C., & Gray, S. H. (2023). Employers' complaints about recent graduates: A survey of hiring managers in 2023. *International Journal of Workforce and Education*, 38(2), 112–130. doi: 10.1086/iwse.2023.245701

- Smith, J. (2023, May 15). The reality of being a graduate in 2023: Unemployment, underemployment, and the blame game. *The Guardian*. Retrieved from <https://www.theguardian.com/education/2023/may/15/the-reality-of-being-a-graduate-in-2023-unemployment-underemployment-and-the-blame-game>
- Smith, S., Johnson, D., et al. (2024). Public perceptions of artificial intelligence and job displacement: A survey study. *Journal of Technology and Society*.
- Sutherland, J. T., & Hughes, E. J. (2024). Employer opinions on graduate employability and the role of higher education. *Journal of Higher Education Policy and Management*, 46(1), 45–67. doi: 10.1086/00015678
- The labor market for recent college graduates: Evidence of growing underemployment. (2023). The labor market for recent college graduates: Evidence of growing underemployment. *Journal of Labor Economics*, 41(2), 123–145. doi: 10.1086/712561
- U.S. Bureau of Labor Statistics. (n.d.). *The beveridge curve (job openings rate vs. unemployment rate), seasonally adjusted*. Job Openings and Labor Turnover Survey. Retrieved 2025-11-23, from <https://www.bls.gov/charts/job-openings-and-labor-turnover/job-openings-unemployment-beveridge-curve.htm> (Job Openings and Labor Turnover Survey)



Regional Analysis



The Cost of Saying No: Ohio SB 52 and Renewable Energy Development

Zoe Epperson, Elizabeth Krotine

Introduction

ON October 11th, 2021, Ohio Senate Bill 52 (SB 52) took effect, subjecting large solar and wind farm projects to review by county commissioners and allowing said commissioners to create exclusion zones within their respective counties. The bill was sponsored by Senators Rob McColley (R) and Bill Reineke (R) and was signed into law by Governor Mike DeWine (R). This new step is in addition to the previous rigorous application process handled by the Ohio Power Citing Board, which requires multiple application letters, public notices, and a public hearing. Because counties can now restrict renewable energy development within their borders, we can examine the specific economic impacts on GDP and the value of housing of a county's decision to restrict. Figure 1 shows the counties that have created these exclusion zones where large solar and wind projects are prohibited, and Figure 2 shows when each county enacted their restrictions.

Ohio Counties with Exclusion Zones for Solar and Wind

Counties in blue enacted restrictions under Ohio SB 52

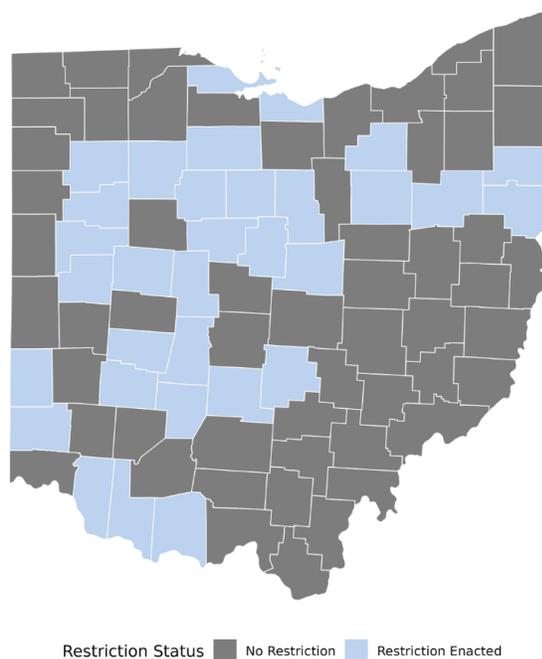


Figure 1: Ohio counties with solar/wind restrictions (as of 9/30/25), self-made restrictions dataset.

Large-scale wind and solar projects are viewed in many rural Ohio counties as disruptive intrusions into long-established agricultural and residential landscapes, as well as a threat to household wealth and property values. Supporters emphasize the aesthetic impacts of solar and wind on rural landscapes, fears about soil and water contamination from materials such as crystalline cadmium telluride in solar panels, and the belief that converting cropland to energy production represents a near-permanent loss of farmland, as reflected in claims that “those fields will never be farmed again” and that “it takes 100 years to restore an inch of topsoil” (Barnet, 2023). What’s more, opponents of large-scale renewable projects often argue that these perceived land-use and environmental risks may manifest as lower property values. There is some evidence for this: “A study from the Lawrence Berkeley National Laboratory found that homes near large solar farms do have lower resale values, though only by 1.5%” (Gottsacker, 2025). The decision to restrict is framed as a way to protect rural character, farmland, and property values in Ohio’s rural counties.

Activism groups, such as Ohio Citizen Action, have opposed the law, claiming that “the added restriction places an unfair burden on renewable energy development, while fossil fuel projects remain largely unaffected.” Opponents of the bill argue that these restrictions infringe on the rights of farmers living on or near affected land to decide how best to use their property. Further, these restrictions prevent developers and landowners from freely contracting over renewable energy projects, artificially constraining consumer and supplier choice and hiding the true demand and economic value of wind and solar. There are also real economic benefits to allowing wind and solar farms to be built in these counties. Counties that choose to welcome renewable energy projects generate revenue that helps fund libraries, police forces, parks, and their education system. In Paulding, they “now have a library system ... in every town” and “a very active parks board that [they] didn’t have before” (Gottsacker, 2025).

While there have been intensive discussions and research about the economic value of renewable and non-renewable energy in Ohio in general, there has been little analysis yet on the impacts of SB 52. We aim to assess how renewable restrictions may be related to home values and GDP.

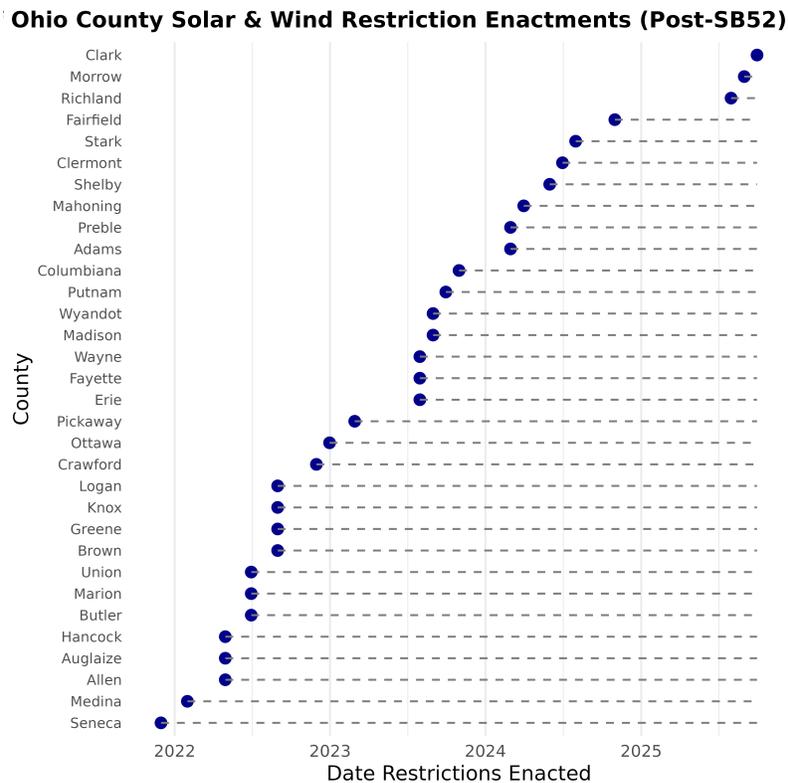


Figure 2: Timeline of solar & wind restriction enactments in Ohio by county (Post-SB52), self-made data-set.

Literature Review

The current view of renewable energy projects by the Ohio Chamber of Commerce informs our analysis of this bill. In recent years, the Ohio Chamber of Commerce’s views on renewable energy has changed. Although opposed to renewable energy projects initially, they now recognize the economic benefits of a statewide initiative on energy and disagree with SB 52. They explain that numerous highly-ranked companies that could provide jobs and tax revenue for Ohioans may choose renewable energy generated outside of the state, rather than in-state non-renewable energy, if given a choice. A brief published by Kingwood Solar notes that “denial of renewable project applications may jeopardize Ohio’s future economic growth and energy stability.” SB 52 could stall economic growth in the energy sector, and therefore growth for the whole economy.

Economist Colton Smith explains, “while renewables such as solar and wind do create problems such as land and habitat losses, high water use, and the use of potentially hazardous materials in the manufacturing process, the benefits of sustainable electricity generation, acid rain declines, emissions reductions, and climate change mitigation still outweigh these costs, especially when compared to the alternatives.” Smith suggests that Ohio should not put as many restrictions on wind and property lines and should protect the state’s renewable energy portfolio (Smith, 2018).

In his dissertation on the impact of renewable energy adoption on state-level housing prices, Checchi finds that solar and geothermal energy have opposite associations with housing prices. Using fixed-effects and instrumental-variable estimation, he reports a positive correlation and causal effect of solar energy consumption on housing prices, and a negative correlation and causal effect of geothermal energy consumption (Checchi, 2024).

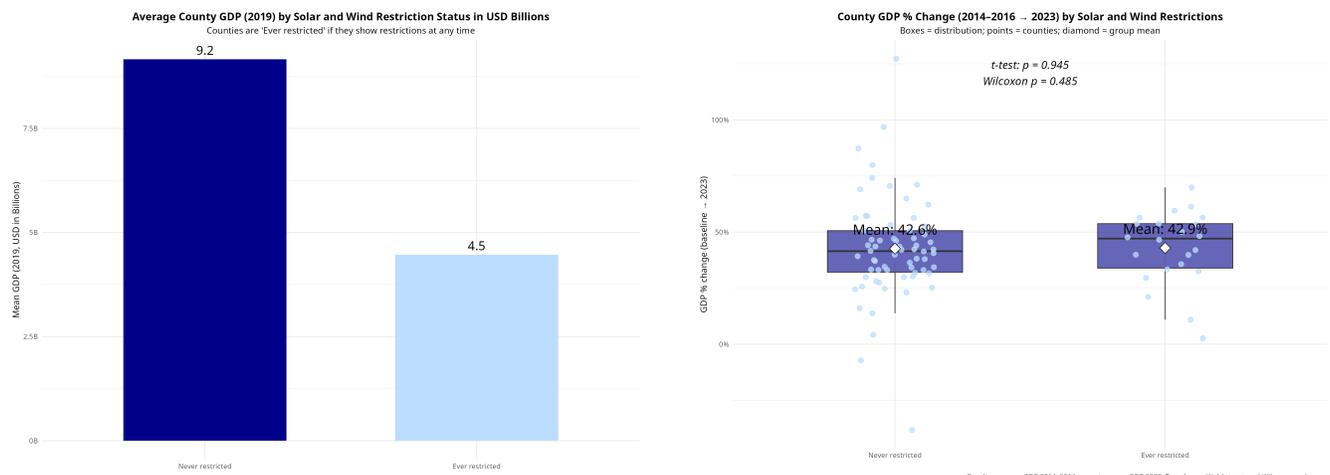
Data and Methods

To obtain an up-to-date and thorough dataset on individual county restrictions, we completed deep research into each county’s current policy on wind and solar farms. This research included scouring local news sources and county websites, as well as cross-checking with a map created by the Ohio Citizen Association, documenting counties with restrictions that impacted five townships or more as of 2023. We, however, decided to consider those who enacted any restrictions to be a “restricted county.” We then found the dates when each of the policies was enacted, giving us a clear timeline of the bill’s outcome. County enactment dates are shown in Figure 2.

For our analysis, we chose to look at variables that are likely to respond relatively quickly to changes in renewable energy restrictions and that are available from datasets that are updated regularly. We obtained annual county-level GDP from the Bureau of Economic Analysis, and we use Zillow’s Home Value Index (ZHVI), updated monthly and at the county-level, as a measure of property values. Using county GDP data from before the passage of Senate Bill 52, we first examine whether counties that eventually adopted restrictions differed systematically from those that did not. We then use changes in GDP and median housing values after the adoption of county restrictions to test whether restricting is associated with different economic trends. Additionally, we run both a t-test and a Wilcoxon signed-rank test to determine if the differences between counties that did or did not restrict are statistically significant. We selected a Wilcoxon signed-rank test because it was less impacted by outliers than a paired t-test, and counties with cities and urban populations were major outliers in our GDP analysis.

Analysis

We use county GDP to compare the overall scale of economic activity in restriction and non-restriction counties. Because we use total GDP rather than GDP per capita, our measure reflects changes in aggregate output and may partly capture differences in population growth or industrial structure, rather than changes in economic well-being per resident. In the left panel of Figure 3, the mean GDP in 2019 is shown for the counties based on their current restriction status. Counties that have never restricted themselves have a much higher average GDP. This could be due to outlier counties such as Cuyahoga County or Franklin County, which have a much higher GDP and did not enact any restrictions, as most wind and solar developers do not wish to build large projects in metro areas. The counties with lower GDP are often in rural, more conservative areas, which accompanies the decision to restrict renewable energy for two major reasons: these areas are the most likely to be impacted by large solar and wind farms, and this bill was championed by conservative politicians who earn higher support in rural areas.



(a) Average County GDP (2023) by restriction status.

(b) County GDP % change (2019–2023) by restriction status.

Figure 3: County GDP levels (2023) and GDP growth (2019–2023) by solar and wind restriction status.

Table 1: County GDP% Change 2014-2016 → 2023

Group 1	Group 2	Mean Group 1	Mean Group 2	T Statistic	DF	P Value
Never Restricted	Ever Restricted	0.426	0.429	-0.07	50.18	0.945

In the right panel of Figure 3, we compare the change in GDP from a 3-year period well before SB 52 to 2023, two years after SB 52. We chose these years as we wanted our base year to predate COVID-19, as some counties were hit harder than others, and our ending year to be the closest possible to 2025. It should be noted that for this graph specifically, we only counted counties that enacted restrictions before 2024 as “restricted.” Counties that restricted had, on average, a 0.3 percentage point higher change in GDP over the selected years. The large p-values from our significance tests determine that this difference is statistically insignificant at the 5% level. One major limitation of this visualization is the fact that we only have two years of GDP data post-SB52. This means that it is possible that in future years, the decision to restrict could have a larger impact than we can not yet determine.

In Figure 4, we use home values as an indicator of the effects of restricting. Before SB 52 was passed in 2021, the average median housing value yearly was higher for counties that had restrictions enacted compared to those that did not. The two lines seem to follow a fairly similar trend prior to SB 52’s passage. After SB 52 was passed into state law, the difference in average median home values between counties that have passed restrictions and those that have not has widened slightly. Counties that have used SB 52 to restrict the development of renewable energy farms have average median home values that are increasing at a higher rate than those with no restrictions. This would support the fears many Ohioans had about solar and wind farms: the creation of these renewable energy farms will hurt housing values. However, there are limitations to this analysis. First, the restrictions in each county were enacted at different times after the bill was passed, so a staggered difference-in-differences approach would be necessary to estimate any possible effect. Second, there could be other factors that have influenced home values in these counties. Finally, by giving each county equal weight, the comparison does not account for differences in population size between counties or the number of households affected.

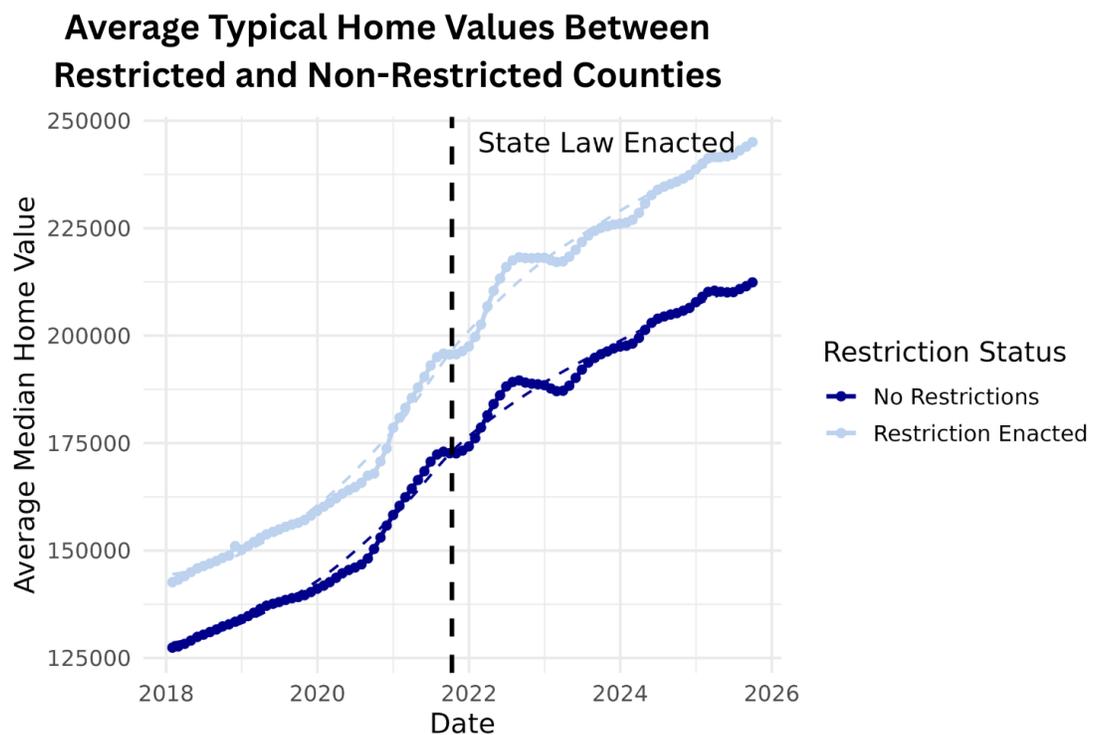


Figure 4: Average typical home values between counties with and without restrictions.

Conclusion

Our analysis of economic data on a county-by-county basis has given us insignificant results for the correlation between GDP, median housing value, and the restriction status of counties. We have, however, gained insights into which counties restricted and why they did. Counties that restrict tend to be rural and have lower total GDP. Ohio Senate Bill 52 has only been in effect since 2021. Many solar and wind projects in counties without restrictions that are currently being built or will be built in the future are not yet causing economic impacts. As renewable energy is becoming increasingly popular, the differences in the economies that allow and those that restrict renewable energy will become clearer. Our analysis of the effects of the bill is exploratory, and we cannot draw any strong conclusions. In the future, we are interested in re-running our tests with more data after restrictions have been enacted for longer periods of time. Additionally, a staggered difference-in-difference model could be used to estimate the causal effect of restrictions. Extensions from this project could include an analysis of the effect of only restricting renewables or, alternatively, only restricting fossil fuels in different areas, or an analysis of the statewide impacts of these restrictions in the long run.

References

- Barnet, K. (2023). *Brown county soil and water conservation board member warns knox county of solar project pitfalls*. Retrieved from <https://mountvernonnews.com/stories/652798093-brown-county-soil-and-water-conservation-board-member-warns-knox-county-of-solar-project-pitfalls>
- Checchi, H. (2024). *The impact of renewable energy adoption on state-level housing market prices* (Unpublished doctoral dissertation). Sacred Heart University.
- Gdp by county, metro, and other areas*. (2025). <https://www.bea.gov/data/gdp/gdp-county-metro-and-other-areas>. (Accessed: 2025-11-19)
- Gottsacker, E. (2025). *Wind energy is bringing major dollars to some ohio counties. others aren't blown away*. Retrieved from <https://www.statenews.org/section/the-ohio-newsroom/2025-08-26/wind-energy-is-bringing-major-dollars-to-some-ohio-counties-others-arent-blown-away>
- Housing data*. (2025). HousingData-ZillowResearch. (Accessed: 2025-11-19)
- Kowalski, K. (2025). The ohio chamber of commerce emerges as an unexpected ally for solar projects. *Ohio Capital Journal*.
- Smith, C. (2018). The solar and wind economy in ohio: Industry analysis and policy implications. *The Journal of Sustainable Development*(20), 43–61.
- Utility-scale solar*. (n.d.). Retrieved from https://www.ohiocitizen.org/utility_scale_solar

Housing Pressure Points: Trends in Homelessness Across Ohio's Urban Counties

Yeva Butovska, David Smith

Introduction

HOMELESSNESS remains a persistent challenge in urban America, shaped by overlapping economic, social, and housing factors. In Ohio, the metropolitan counties of Cuyahoga (Cleveland), Franklin (Columbus), Hamilton (Cincinnati), Montgomery (Dayton), and Summit (Akron) account for most of the state's unhoused population. Each faces distinct affordability pressures and demographic changes, making them useful for exploring how economic and housing conditions relate to homelessness at the regional level. This paper presents an exploratory analysis using Point-in-Time (PIT) counts, eviction filings, and population-adjusted homelessness rates from 2007 to 2024. The findings highlight general correlations between economic stress, housing availability, and homelessness, without inferring direct causation.

The academic literature provides a unified framework for interpreting these regional trends by highlighting how affordability pressures, economic shocks, and institutional capacity jointly shape homelessness. Glynn and Fox show that even small rent increases can raise homelessness rates in metropolitan areas with low vacancy, a pattern consistent with rising rents and lagging incomes in Ohio's large counties (Glynn & Fox, 2019). Evans, Phillips, and Ruffini (2019) add that homelessness tends to surge after recessions as job losses and rental arrears accumulate, helping explain Ohio's peaks around both the 2008 financial crisis and the COVID-19 downturn (Evans, Phillips, & Ruffini, 2019). Complementing these economic explanations, Lim et al. (2018) demonstrate that supportive housing and coordinated health and social services substantially reduce homelessness and related public costs, indicating that local institutional resilience can moderate economic pressures (Lim et al., 2018). Taken together, these studies suggest that rising rents and income stagnation increase vulnerability, while differences in regional capacity to absorb economic shocks shape the depth and duration of homelessness across Ohio's counties.

Methodology

This study draws from several complementary data sources to explore these questions. Annual Point-in-Time (PIT) counts from 2007-2024 (U.S. Department of Housing and Urban Development, 2007) provide standardized estimates of sheltered and unsheltered homelessness across each Continuum of Care region. The Eviction Lab (The Eviction Lab at Princeton University, 2024) contributes county-level measures of eviction filings, median rent, and poverty rates, offering a lens into housing market pressure and socioeconomic context. To account for population differences, annual county population data were taken from the U.S. Census Bureau's 2000–2010 Intercensal Estimates (U.S. Census Bureau, 2000) and 2010–2020 County Population Totals (Vintage 2020) series (U.S. Census Bureau, 2010). By dividing PIT counts by annual population estimates, homelessness rates per 10,000 residents were derived, producing a per-capita metric that allows meaningful comparison across regions of different size and growth rates.

Results

Across Ohio's major counties, homelessness rates reflect broad economic cycles and housing market conditions rather than any single cause. Population-adjusted trends show peaks after 2008 and 2020, periods marked by widespread job loss and housing instability, but the data cannot confirm direct causal relationships. Instead, these patterns indicate strong correlations between economic downturns, affordability stress, and the prevalence of homelessness. Cuyahoga and Hamilton counties consistently exhibit the highest rates, often between 15 and 20 per 10,000 residents, coinciding with higher eviction activity and older rental housing stock. Franklin County shows a steady increase through the 2010s as Columbus experienced rapid population growth and rising rents. Montgomery and Summit counties display lower and more stable rates, reflecting differences in local housing markets and economic composition. These variations suggest that homelessness in Ohio's urban centers is closely associated with, though not necessarily caused by, affordability pressures and rental market dynamics.

Eviction data further support this pattern. In Cuyahoga County, filings spiked between 2007 and 2010, declined during the mid-2010s, and rose again after pandemic-era moratoria expired. These temporal parallels between eviction and homelessness point to a potential link, but the available data are descriptive and cannot confirm a

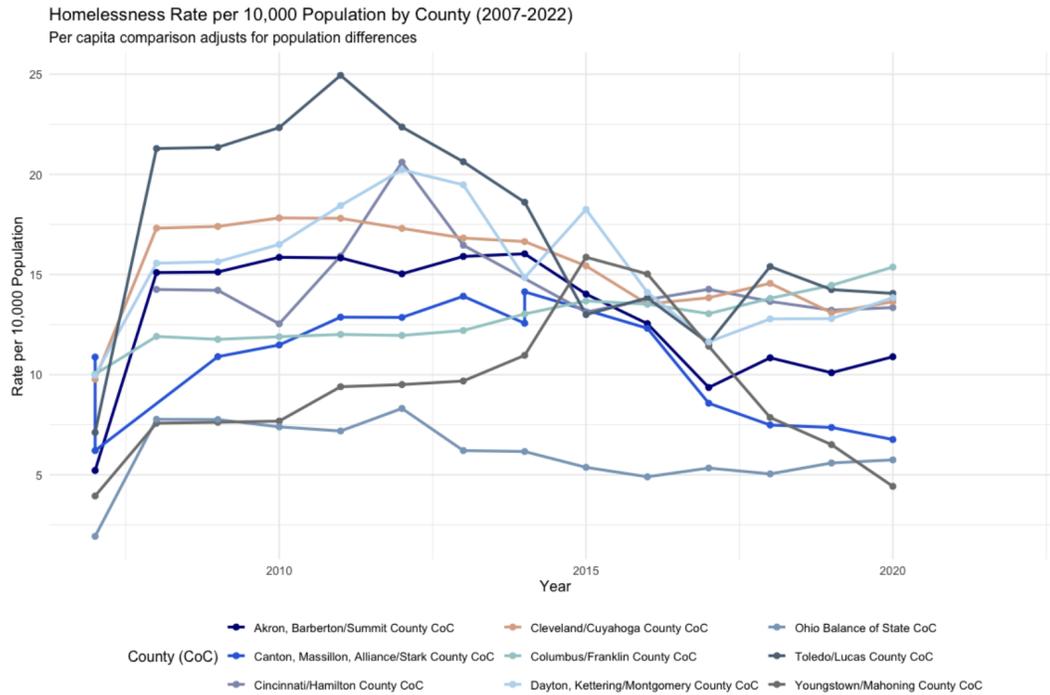


Figure 1: Homelessness Rate per 10,000 Population by County (2007-2022)

direct sequence. Similar cycles appear in other counties, especially in rural areas, suggesting that eviction rates can serve as useful indicators of broader housing stress, even if the causal mechanisms remain untested.

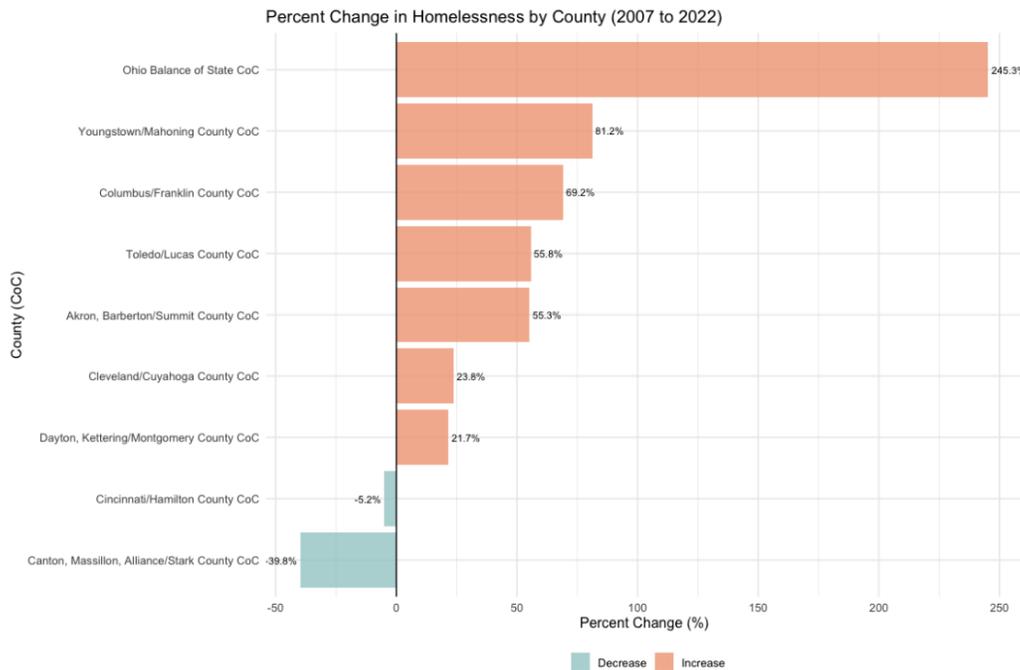


Figure 2: Percent Change in Homelessness by County (2007-2022)

Ohio’s homelessness landscape is shaped by both long-term structural forces and region-specific conditions. Urban centers with strong population growth face affordability pressures, while many rural and small-town counties appear increasingly exposed to economic vulnerability and limited support infrastructure.

One positive note is that most homeless people in Ohio are considered sheltered.

Taken together, these observations underscore that homelessness in Ohio’s largest counties co-moves with key

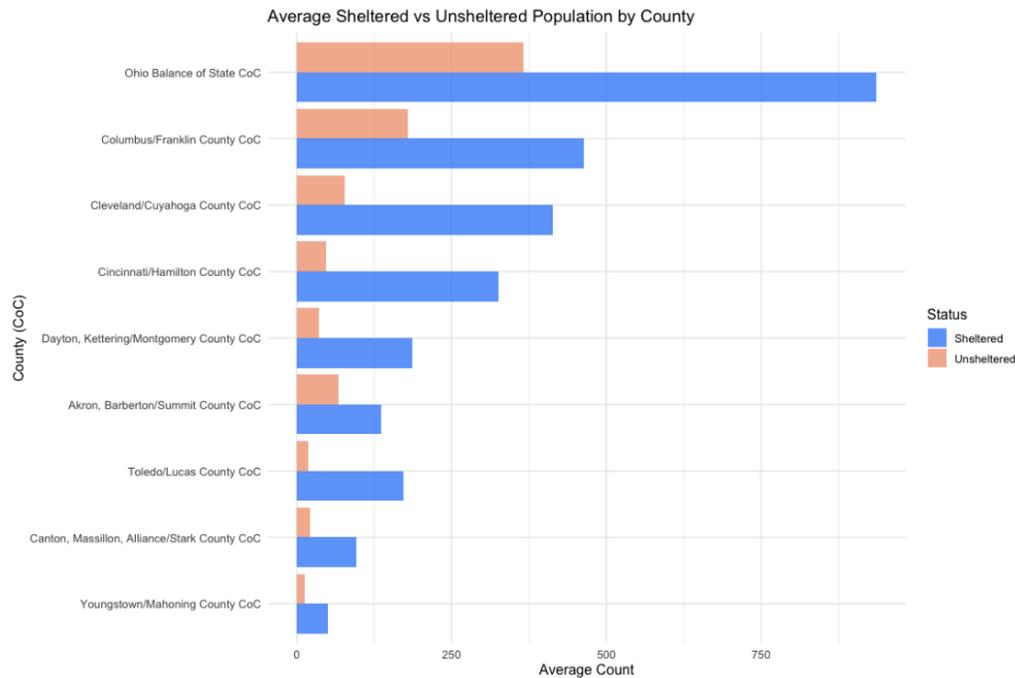


Figure 3: Average sheltered vs unsheltered homeless counts by county (2007-2022)

economic and housing indicators. Further research incorporating programmatic data and longitudinal tracking would be needed to determine causation or the impact of specific policy interventions.

These findings support a framework in which homelessness arises from the interaction of three factors: baseline affordability, economic vulnerability, and housing market instability.

Conclusion

Homelessness in Ohio's five largest counties reflects the broader structural tension between housing costs and income security. The evidence suggests that increases in eviction activity coincide with rising homelessness, particularly in regions with constrained affordable housing stock. Counties like Cuyahoga and Hamilton face continued affordability challenges, while Franklin's trajectory illustrates how economic growth can coincide with housing stress when supply lags behind demand. Montgomery and Summit counties, by contrast, demonstrate more gradual trends, suggesting less volatility in their local rental markets. The results imply that effective monitoring of eviction and rent indicators could provide an early signal for when and where homelessness is likely to rise. While causality cannot be established from descriptive data alone, these patterns highlight the importance of sustained attention to affordability and housing stability as central components of regional economic health.

References

- Evans, W., Phillips, D., & Ruffini, K. (2019). Policies to address homelessness. *NBER Working Paper*(26550).
- Glynn, C., & Fox, E. (2019). *Dynamics of homelessness: Estimating population changes in large metropolitan areas*.
- Lim, S., et al. (2018). Supportive housing and health: Longitudinal outcomes among a cohort of chronically homeless individuals with serious mental illness. *American Journal of Public Health*, 108(8), 1050–1057.
- The Eviction Lab at Princeton University. (2024). *National database of eviction filings and housing indicators*.
- U.S. Census Bureau. (2000). *County intercensal population estimates*. (Data series spanning 2000–2010)
- U.S. Census Bureau. (2010). *County population totals (vintage 2020)*. (Data series spanning 2010–2020)
- U.S. Department of Housing and Urban Development. (2007). *Point-in-time counts of homelessness*. (Data series spanning 2007–2024)

What are the Effects of Game-Day Rainfall on Major League Baseball Home Game Attendance?

Erin Butler, Matteo Aron, Paulo Aguiar

Introduction

WEATHER has long been an overlooked factor influencing sports events, despite evidence that temperature is a significant factor in a team's performance (Koch, 2013). This influence extends past the scope of athletic performance and into the economic success of sports teams, where sports fans and ticket sales are often at the mercy of outdoor weather conditions. The economic health of sports teams plays a vital role in sustaining both the local and broader metropolitan economy. Game days can account for an outsized share of revenue for nearby businesses and drive the foot traffic that keeps them afloat (Giri, Sagan, & Podgursky, 2024).

The literature on the effects of temperature is broader than precipitation, so the goal of this paper is to isolate the effect of weather on Major League Baseball (MLB) ticket sales by analyzing the relationship between precipitation and home game attendance. In doing so, we aim to deepen understanding of fan behavior. This relationship has important implications for both team revenue forecasting, stadium dome rationale, and local business preparation. To do this, we employ a linear regression model with lagged rainfall and date of game to control for prior rainfall and seasonality, regressing attendance on rainfall. Our hypothesis is that precipitation discourages home fans from attending MLB games.

Literature Review

A substantial body of research has examined the determinants of sports attendance, emphasizing economic, performance-based, and environmental factors. Pang Yu and Wang Fengchen study NFL stadium attendance rate by examining factors ranging from economic indicators, such as personal income and unemployment rate, to team performance, such as the home win percentage. Although temperature is included in their model, it ranks only tenth among the seventeen variables, suggesting that weather is not among the deciding factors of attendance outcomes.

Other research highlights that weather may exert indirect effects on attendance by shaping on-field performance. Koch and Panorska's study of weather's impact on Major League Baseball games demonstrates that temperature significantly alters teams performances on the field. In warm weather, runs scored, batting average, slugging percentage, on-base percentage, and home runs significantly increase, while walks significantly decrease. This correlation in temperature and performance could, in turn, impact fans' decisions to attend a game. Together, these studies suggest that while environmental conditions have not been the main focus of attendance models, through its direct and indirect impacts on attendance, further investigation is warranted.

The economic significance of this investigation is supported by the findings in Giri et al, who examine the economic spillover effect of Major League Baseball in Saint Louis games onto neighboring businesses on gameday by measuring the size of the spread of the spillover, the business types affected, and spending data. They found significant spending patterns on game day at Busch Stadium, with notable revenue increases in businesses such as restaurants, bars, grocery stores, and liquor stores.

Data and Methodology

To examine the relationship between daily precipitation and Major League Baseball (MLB) game attendance, multiple data sets were used to give accurate results. For the precipitation data, two data sets were combined: the National Oceanic and Atmospheric Administration's (NOAA) and Meteostat weather data. NOAA's weather dataset is a comprehensive record of daily weather observations across United States airport weather stations, including precipitation. Since MLB stadiums are typically located within a few miles of major airports, NOAA's dataset provided accurate and reliable observations that closely reflected weather conditions on game-day. To ensure further accuracy, these datapoints were supplemented by an additional dataset from Meteostat with precipitation data from additional weather stations.

The second primary data source was Baseball Reference, which provided complete records of Major League Baseball game statistics. This dataset was limited to home games played between the 2018–2025 seasons and provided information such as the opponent, rank of home team, time and date of game, and attendance of the

game. The dataset excluded 2020 due to the COVID-19 related restrictions, resulting in zero or near-zero stadium attendances.

The data sources were then merged by matching each MLB game to the nearest weather station dataset for the same date. After the merge, the dataset contained 16,653 Major League Baseball games from 30 MLB teams.

To estimate the relationship between rainfall and game attendance, the following ordinary least squares (OLS) model was used:

$$\text{Attendance}_{it} = \beta_0 + \beta_1 \text{rain_mm}_{it} + \beta_2 \text{rain_prev_5days}_{it} + \beta_3 \text{DayOfYear}_{it} + \beta_4 \text{DayOfYear}_{it}^2 + \varepsilon_{it}$$

This model serves to isolate the direct effect that rainfall has on game attendance, controlling for two variables: rainfall in the prior five days and seasonality. `rain_mmit` represents the rain in millimeters on the day of the game. `rain_prev_5daysit` is the sum of the total rainfall of the previous five days. This variable was included to control for the effects of rainfall during the week leading up to the game, isolating solely the effect of game-day rainfall. Past literature on the effects of lagged rainfall on MLB game attendance found that lagged cumulative day rainfall leads to an increase in attendance of approximately 130 additional fans (Ge, Humphreys, & Zhou, 2017). `DayOfYear` and `DayOfYear2` both represent the calendar date, controlling for the fluctuations in demand related to the time in the season the game was played. `DayOfYear2` was specifically chosen as a control as it allows the model to fit a curved line, instead of relying only on the linear trend `DayOfYear` provides, improving model fit by capturing curvatures in attendance trends throughout the season.

Results

After using the OLS regression model outlined, the main finding is that there is a negative correlation between the amount of precipitation (`rain_mm`) and the game-day attendance of Major League Baseball (MLB) games. The OLS regression results are as follows:

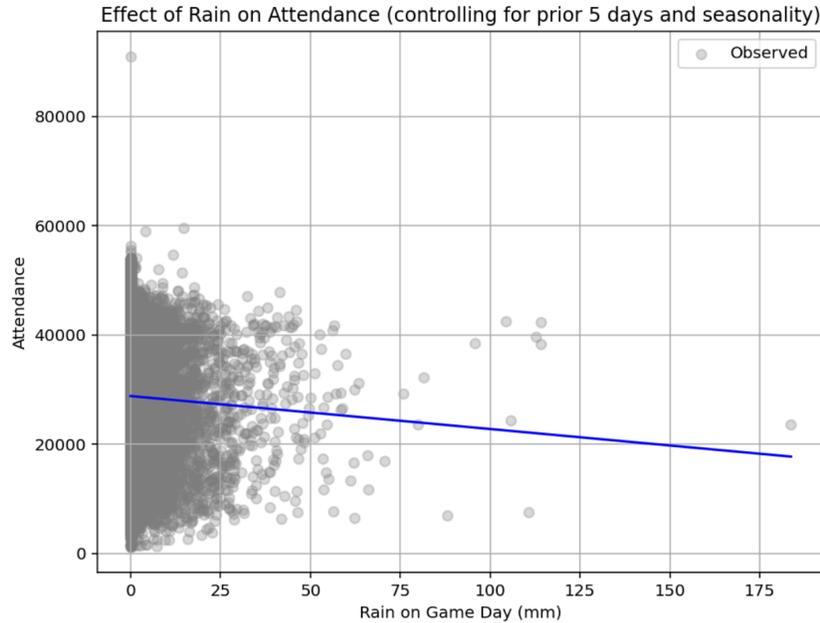
The R^2 value for the OLS is 0.022, meaning that our model explains 2.2 percent of the variation in MLB attendance. While this is a statistically low-value, our F-statistic is below 0.001, meaning the overall model is statistically significant. These values confirm our hypothesis that precipitation does have a negative impact on game-day attendance, however the relatively low R and R^2 values means that attendance is driven by many other factors that our model does not include.

OLS Regression Results	
Dep. Variable:	Attendance
Model:	OLS
Method:	Least Squares
Date:	Wed, 19 Nov 2025
Time:	19:11:04
No. Observations:	16653
Df Residuals:	16648
Df Model:	4
R-squared:	0.022
Adj. R-squared:	0.022
F-statistic:	94.91
Prob (F-statistic):	5.70e-80
Log-Likelihood:	-1.7932e+05
AIC:	3.586e+05
BIC:	3.587e+05
Covariance Type:	nonrobust

Table 1: OLS model summary

Variable	coef	std err	t	$P > t $	[0.025, 0.975]
Intercept	1.292e+04	1077.579	11.986	0.000	[1.08e+04, 1.5e+04]
rain_mm	-60.3715	13.179	-4.581	0.000	[-86.204, -34.539]
rain_prev_5days	-22.2195	4.509	-4.928	0.000	[-31.057, -13.382]
DayOfYear	149.6138	12.417	12.050	0.000	[125.276, 173.951]
I(DayOfYear ²)	-0.3395	0.033	-10.203	0.000	[-0.405, -0.274]

Table 2: Coefficient estimates



Furthermore, the results provided a quantitative effect on how much attendance dropped with each millimeter (mm) of rain. With a coefficient of -60.37 , rain_mm explains that with each millimeter of rain, attendance would drop on average by 60 fans. This result is statistically significant, but relatively small in magnitude. The control variables coefficients, rain_prev_5days , DayOfYear , and DayOfYear^2 , also provided valuable insights into lagged precipitation effects on gameday attendance and season trends. rain_prev_5days had a coefficient of -22.22 , meaning that with every millimeter of rain in the previous five days, attendance dropped by an average of 22 fans. This was statistically significant, but still of relatively small magnitude. DayOfYear , with a coefficient of 149.61, reflects a steady increase in attendance by an average of 150 fans for each passing day in the season. This increase steadily decreases, as is reflected by DayOfYear^2 's coefficient of -0.3395 . These results indicate that attendance is relatively low at the beginning of the season, slowly increasing with each successive match day as teams progress towards playoffs, with attendance declining late in the season.

Although gameday rainfall and recent rainfall have a statistically significant impact on gameday attendance, its effects are small in magnitude. Larger variations in gameday attendance can be explained by other variables.

Limitations

There are several limitations in the data we were able to collect. Primarily, due to the large range of dates, we were limited to one source for the data, introducing the risk that there could be other variables with a large effect on attendance that we were unable to control for. These variables could range from ticket prices, macroeconomic conditions (e.g. unemployment and inflation rates), marketing strategies, or other factors that influence attendance. Our dataset was unable to account for games that were cancelled due to rain, as those games were simply removed in favor of the make-up game. Additionally, the rain data we collected does not specify whether the precipitation was rain or snow, nor does it indicate the strength of the storms on a given day as precipitation is a volume metric that does not indicate the severity of the rain. We also did not control for teams with dome roofs, which could cause the effect of rain on attendance to appear smaller. A major limitation to our study is that we did not account for regional climate. This could potentially result in a disproportionate effect of precipitation on game-day attendance throughout the United States as average rainfall is expected to vary regionally.

Conclusion

In this paper, we aim to analyze how rainfall's impact on in-stadium attendance for Major League Baseball (MLB) home games. Using an Ordinary Least Squares (OLS) regression model, we find statistically significant evidence that precipitation is negatively correlated with fan attendance.

While the analysis confirms the negative correlation between precipitation and game-day attendance, the OLS model reveals that rainfall influences attendance alongside other determinants, such as lagging rainfall, date of the match, and other variables not investigated in this study.

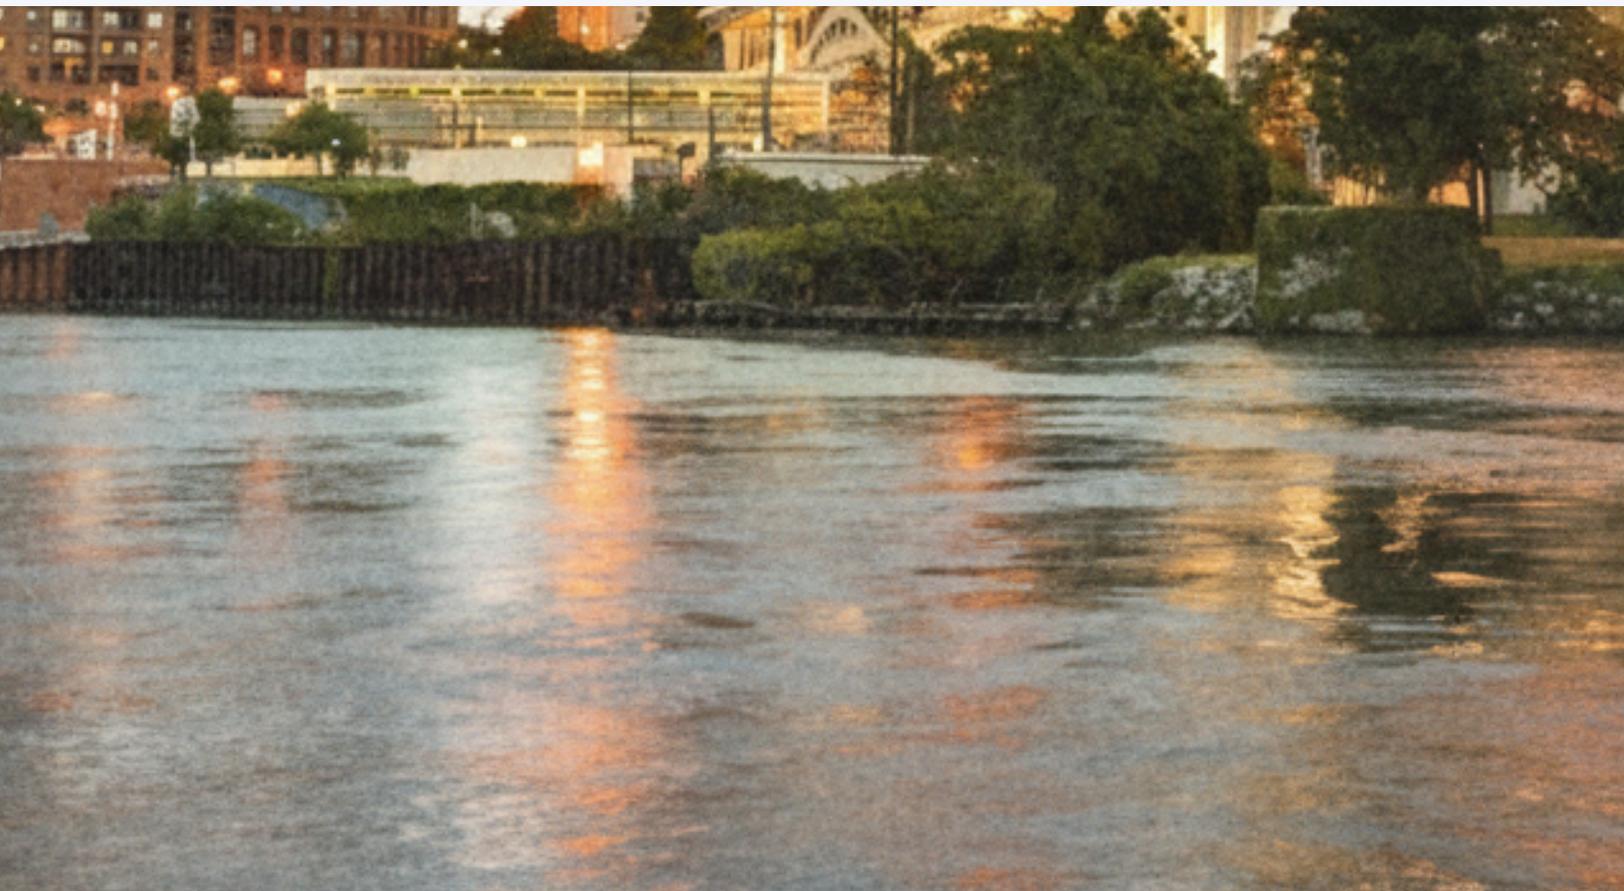
This paper shows that understanding the magnitude of the influence of weather on attendance has economic implications for both MLB teams and surrounding local businesses. These findings on how demand fluctuates with rainfall could inform teams on pricing strategies and stadium architecture. Ultimately, this study broadens the literature available on how environmental factors shape sport fans' choices in live sporting events.

References

- Baseball-Reference. (n.d.). *Major league baseball game statistics (2018–2025)*. Sports Reference LLC. Retrieved from <https://www.baseball-reference.com> (Retrieved 2025-11-24)
- Ge, Q., Humphreys, B. R., & Zhou, K. (2017). Are fair weather fans affected by weather? rainfall, habit formation and live game attendance [Working Paper]. (17-24). Retrieved from https://researchrepository.wvu.edu/econ_working-papers/18 (WVU Research Repository Working Paper)
- Giri, A., Sagan, V., & Podgursky, M. (2024). The ballpark effect: Spatial-data-driven insights into baseball's local economic impact. *Applied Sciences*, *14*(18), 8134. Retrieved from <https://www.mdpi.com/2076-3417/14/18/8134> (Retrieved 2025-11-24) doi: 10.3390/app14188134
- Koch, B. L. D., & Panorska, A. K. (2013). The impact of temperature on major league baseball. *Weather, Climate, and Society*, *5*, 359–366. Retrieved from https://journals.ametsoc.org/view/journals/wcas/5/4/wcas-d-13-00002_1.xml (Retrieved 2025-11-24) doi: 10.1175/WCAS-D-13-00002.1
- Meteostat. (n.d.). *Historical weather data (2018–2025)*. Meteostat Weather Service. Retrieved from <https://meteostat.net> (Retrieved 2025-11-24)
- National Oceanic and Atmospheric Administration. (n.d.). *Daily weather observations (2018–2025)*. NOAA National Centers for Environmental Information. Retrieved from <https://www.ncei.noaa.gov> (Retrieved 2025-11-24)
- Pang, Y., & Wang, F. (2024). Forecasting stadium attendance using machine learning models: A case of the national football league. *Studia Sportiva*, *18*. Retrieved from <https://journals.muni.cz/studiasportiva/article/view/38549> (Retrieved 2025-11-24) doi: 10.5817/StS2024-2-12



Index



Trade Centrality Index

Fiona Arado, Bo Cao, Daniel Chase, Isma'il Seddon

Understanding the key players in trade requires an analysis of relationships between countries. Creating an index that measures centrality of countries in trade suggests the use of graph theoretic frameworks, as these are strong models for representing relationships between entities. We therefore make use of graph theory foundations and add nuance by incorporating value added in trade. We create an index for several countries of interest over time and offer analysis to explain our findings.

Introduction

TRADE policy has become an increasingly highly-charged issue in the United States. Many politicians, including President Donald Trump, have argued that the U.S. has become too dependent on imports. President Trump has used this philosophy to defend his recent tariff policy, which has had profound impacts on both the U.S.'s geopolitical relationships and overall macroeconomic trends. Given the breadth of the impact of trade policy, constructing metrics can help to increase understanding of trade dynamics and to more effectively quantify a country's position in the global trade network. In this paper, we tackle exactly that problem and construct an index to measure trade centrality. Trade centrality is a useful measure in that it can predict a country's susceptibility of being impacted by changes in foreign trade policy. Centrality can be measured on both the sector and product level, which allows for more fine-tuned inquiry into a country's trade weaknesses.

On the surface, it might seem trivial to determine which countries stand at the center of global trade – one can simply rank nations by their exports and GDP. Such an approach is straightforward and easily digestible, but it loses the nuance of trade we attempt to capture in our index. Moreso, it is the relationships and dependencies between countries that we are interested in. Since we're interested in the trade dynamics between countries, we will employ a network analysis of trade rooted in graph theory. Such an approach enables us to consider each individual relationship between pairs of countries, which we will piece together to get a more comprehensive view. See Figure 1 for an example of a graph.

We will further our analysis by not just considering raw imports and exports between countries, but the true value added by each country as they pass products between each other. This notion of "true value added" is captured by Trade in Value Added (TiVA), and we will use such information to get an understanding of which countries are truly producing value in the supply chain. To our knowledge, using both TiVA and network analysis is novel, and in this paper, we discuss our methodology for integrating these.

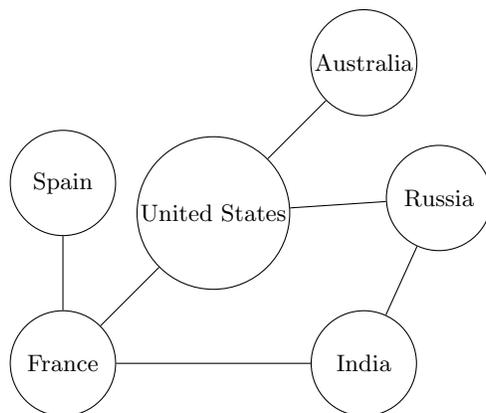


Figure 1: Example of a graph with six countries.

Literature Review

Although trade has long been a popular topic for economic indices, many choose to focus on narrow aspects. For instance, Caldera and Iacoviello's Geopolitical Index uses topic words in news articles to understand geopolitical relationships to estimate the risk of economic disasters (Caldera & Iacoviello, 2022). This index offers a unique understanding of mapping relationships between countries, but the authors limit their analysis to economic disasters and do not consider trade relationships.

There is substantial literature on using network analysis for trade. For example, the Centre for Prospective Studies and International Information (CEPII) has constructed an index that visualizes the World Trade Network, describing the topology of the network of trade through network analysis (Benedictis, Nenci, Santoni, Tajoli, & Vicarelli, 2013). However, one drawback to this index is that it does not choose to take the product-level networks into account in its overall measure of centrality. Instead, it prioritizes spatial relationships and physical accessibility of countries to one another.

In this paper, we base the first part of our methodology on Korniyenko, Pinat, and Dew's 2017 framework used by the International Monetary Fund (IMF). Korniyenko et al outline a method that calculates, at the product level, which countries are "central players" by using outdegree centrality (Korniyenko, Pinat, & Dew, 2017). Calculating centrality at the product level is particularly advantageous as it gives us the opportunity to calculate weightings on particular products based on their determined "criticality" (e.g. healthcare imports would be considered more critical than furniture and can be weighted higher). While we do not calculate weight by industry based on criticality in this work, we believe this is a promising avenue to explore and plan to do so in future work. Therefore, Korniyenko et al provides us with an opportunity to incorporate this further, product-level analysis.

As previously mentioned, we add complexity to our examination of trade by incorporating Trade in Value Added. TiVA adds depth to our analysis by measuring what the value a country adds to a product at the different stages of production, rather than just at the finished product. It involves creating calculations based on the input and output of products by country, and then weighting them based on the additional value that they bring to a product. An example of incorporating TiVA into trade analysis is Johnson and Noguera's 2014 study, where they used TiVA to analyze trade wedges and their effects on forty-two countries using a multi-sector structural gravity model (Johnson & Noguera, 2014). Their use of TiVA allowed them to take into account complicated patterns of bilateral trade, but their analysis largely focused on spatial factors and regional trade agreements as opposed to the centrality of a country in the larger global trade network. By incorporating TiVA into our analysis, we, too, are able to find more nuanced patterns of trade dependencies and identify weak spots of the United States and its trade partners.

Methodology

We model international trade as a directed and weighted network in which countries are nodes and exports are edges. Let $w_{ij,t}$ denote the gross export value (in USD) from country i to country j in year t . For each importer i , we compute its average import per trading partner as

$$\langle w_{ij,t} \rangle = \frac{1}{k_{j,t}} \sum_{i=1}^N w_{ij,t}, \quad (1)$$

where N is the number of countries in the sample and $k_{j,t}$ is the number of countries j imports from in year t . this term captures how much country i typically imports from each partner and is used as a normalization factor. Incorporating the methodology from Korniyenko, Pinat, and Dew, the outdegree-based centrality of exporter i in year t is defined as

$$c_{i,t} = \sum_{j=1}^N \frac{w_{ij,t}}{\langle w_{j,t} \rangle}. \quad (2)$$

Therefore, a country's centrality is high when it has a high volume of exports. The degree of centrality increases further when a country's trade partners import a high volume of diverse goods.

Gross trade flows overstate a country's true contribution when production is fragmented across borders. To account for this, we incorporate Trade in Value Added (TiVA) data from the OECD. Let $\theta_{i,t} \in (0, 1]$ denote the share of domestic value added in country i 's exports in year t . We construct TiVA-adjusted export flows

$$w_{ij,t} = \theta_{i,t} w_{ij,t} \quad (3)$$

so that countries engaging mainly in assembly (low $\theta_{i,t}$) contribute less to measured flows than countries generating high domestic value added. Replacing $w_{ij,t}$ with $v_{ij,t}$ in the previous expression yields our added centrality value:

$$C_{i,t}^{\text{TiVA}} = \sum_{j=1}^N \frac{v_{ij,t}}{\langle v_{j,t} \rangle}, \quad \langle v_{j,t} \rangle = \frac{1}{k_{j,t}} \sum_{i=1}^N v_{ij,t}. \quad (4)$$

This index can be interpreted as the structural importance of country i in generating value added for global trade network in year t . To compare centrality scores over time, we adjust all TiVA values for inflation.

Data Sources

We sourced our data from the Organization for Education Co-operation and Development (OECD). We selected 30 countries for analysis that we believe will offer interesting analysis and comparisons. These countries represent major players in world trade. Our selected countries are: Australia, Canada, Chile, Colombia, Czechia, France, Germany, Ireland, Israel, Italy, Japan, Korea, Mexico, New Zealand, Poland, Türkiye, UK, US, Bangladesh, Brazil, China, Egypt, India, Indonesia, Pakistan, Philippines, Russia, Saudi Arabia, Singapore, Vietnam.

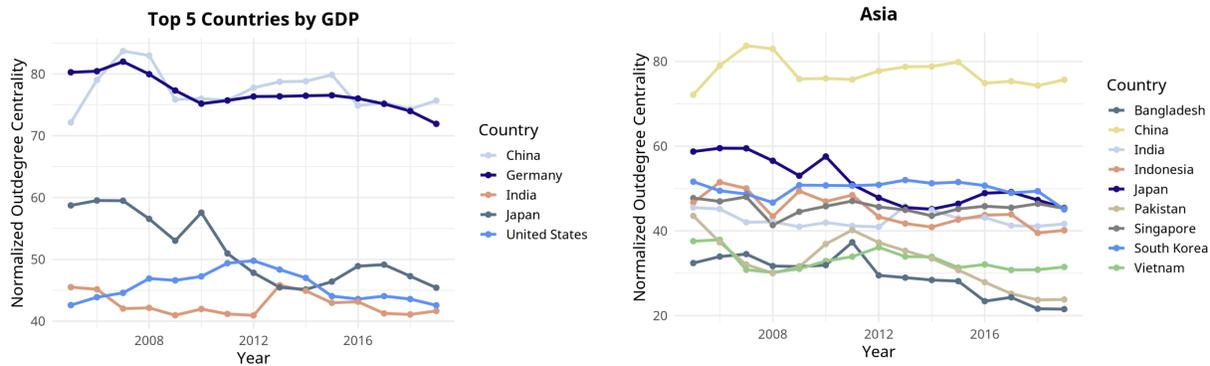


Figure 2: Left: Trade centrality index for the top five countries by GDP. Right: Trade centrality index for selected countries in Asia.

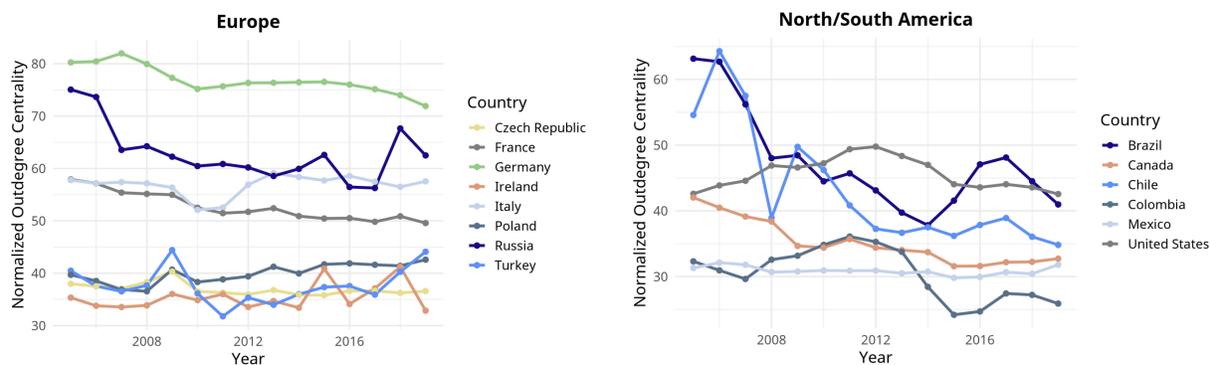


Figure 3: Left: Trade centrality index for selected countries in Europe. Right: Trade centrality index for selected countries in North/South America.

Results and Discussion

In Figures 2 and 3, we report our results. These figures show the extent to which each country functions as a supplier of value added within the global value chain over time. This allows us to diagnose underlying trade shifts and analyze relationships between countries.

The Americas

The left panel of Figure 3 below displays the normalized outdegree centrality over time for North and South American countries. This figure shows that the U.S. has consistently remained a key export hub with a steadily high index of around 45 -50. Before 2008, Brazil and Chile were very central exporters with an index of over 60. After 2008, we observed a sharp decline of 20 index points over the next 10 years. This is likely due to a decline in commodity prices resulting from decreased demand from China for raw materials. Chile’s export structure is heavily concentrated in the copper industry and is therefore vulnerable to price shocks. This is something that can be analyzed in further depth by constructing an industry-level centrality index.

Brazil’s time-series profile illustrates that export centrality does not necessarily move in tandem with domestic economic performance. Before 2014, we saw strong economic growth, yet a weakening centrality index. This is

likely due to increasing dependence on a small number of large commodity trading partners. Throughout their 2014 recession, we saw a sudden spike in export centrality. This is indicative of the slowing domestic consumption as businesses seek to sell to alternative foreign consumers. The case for Brazil exemplifies how having a higher or lower export centrality is not always an indicator of the economic health of a country, but rather can show whether the value created in a country is demanded domestically or by foreign countries.

Top Five by GDP

The left panel of Figure 2 shows export centrality for the top five GDP countries. While all five countries have a relatively high level of export centrality—indicating their importance in the global value chain—Germany and China’s indices are nearly twice that of the United States, Japan, and India. Germany is clearly the most stable hub in the global value chain, likely due to its large, diversified market of global consumers. The Chinese export position post-2011 began to decline due to a national shift to seek to boost domestic consumption. The Chinese case demonstrates the capability of this index measure as a tool for the analysis of macroeconomic policies. Japan’s structural decline as a global exporter of value over the last decade is shown clearly in the figure. This coincides with the long-term decline in key export industries such as automobiles and electronics.

Europe

The Russian export centrality appears to be in a stable decline. A large form of export is the export of oil and gas into Western Europe. It appears that our measure of export centrality aligns with the price of oil and gas. When the commodity is rare and in high demand, Russia fortifies its position as the central player in energy exports.

Asia

A striking observation from Figure 2 is the sharp decline in export centrality of both Pakistan and Bangladesh. Both countries have an extremely narrow export structure, dominated by the textiles industry. Alternative sources of production with looser labor laws and lower minimum wages, such as Vietnam, have increased their share as an exporter of textiles. Pakistan has faced a problem of political instability which—coupled with chronic energy shortages—has bruised their global competitive edge. These shocks have increased the volatility in export volumes and discouraged diversification.

Future Work

For future work, we hope to segment our analysis by industry, as our methodological framework makes such analysis straightforward. We also hope use our framework to analyze “shock” events, such as COVID or the financial crisis.

References

- Benedictis, L. D., Nenci, S., Santoni, G., Tajoli, L., & Vicarelli, C. (2013). *Network analysis of world trade using the baci-cepii dataset* (Working Papers No. 2013-24). CEPII. Retrieved from <https://www.cepii.fr/CEPII/fr/publications/wp/abstract.asp?NoDoc=6112>
- Caldara, D., & Iacoviello, M. (2022, April). Measuring geopolitical risk. *American Economic Review*, 112(4), 1194–1225. Retrieved from <https://www.aeaweb.org/articles?id=10.1257/aer.20191823> doi: 10.1257/aer.20191823
- Johnson, R. C., & Noguera, G. (2014, Jan). A portrait of trade in value-added over four decades. *The Review of Economics and Statistics*, 99(5), 896–911. Retrieved from <https://cep.lse.ac.uk/seminarpapers/21-05-14-NG.pdf> doi: https://doi.org/10.1162/rest_a_00665
- Korniyenko, M. Y., Pinat, M., & Dew, B. (2017, Feb). *Assessing the fragility of global trade: The impact of localized supply shocks using network analysis* (IMF Working Papers No. 2017/030). International Monetary Fund. Retrieved from <https://ideas.repec.org/p/imf/imfwpa/2017-030.html> doi: None
- Trade in value added (tiva) 2025 edition: Origin of value added in gross imports/exports.* (2025). Retrieved from <https://data-explorer.oecd.org>



External Submissions



Tax Portfolio Composition and Revenue Volatility: A Comprehensive Analysis of State Tax Systems, 2005-2024

Ryan Hydrick, Peter Constant

State governments face the critical challenge of providing consistent services, while managing revenue streams that fluctuate with economic conditions. Previous studies establish that sales taxes generally stabilize revenue collections while income taxes increase volatility. However, the conditions under which these patterns hold, the magnitude of their effects, and the role of specific tax subcategories remain inadequately understood. We use data from the US Census Bureau to assess state tax portfolio composition and revenue volatility across all 50 states from 2005-2024. With that data, we employ linear regression models with four distinct specifications to assess the volatility effects of six broad tax categories: property, general sales, selective sales, license, income, and other taxes. Contrary to conventional wisdom, tax volatility may be overstated for states with relatively standard tax portfolios and stable economic conditions. In states with special circumstances, it is important to know how economic conditions can affect revenue sources.

Introduction

WHEN state tax revenues drop unexpectedly, governments face an immediate choice: cut services, drain reserves, increase fees, or institute new taxes. Unlike the federal government, states cannot run deficits or print money—thus making steady tax revenue essential for consistent public services. With this challenge, states must exercise conservatism in how they forecast revenue and account for potential fluctuations in revenues. Due to the complex nature of tax composition and revenue, there is a gap between conceptual economic theory and government practice. Previous research examines individual tax types in isolation, leaving gaps in selective sales, general sales, license, and severance taxes.

States pursuing fiscal stability need evidence-based guidance, not assumptions based on older research. The consequences of ignoring revenue volatility concerns can put constituency services at risk. Our research puts together entire state tax portfolios and explores their relationship with revenue volatility in the last 20 years, from 2005 to 2024.

Literature Review

The scale and depth of state government services, regulation, and administrative roles are immense. However, state governments do not have the same financing ability as the federal government, and forty-nine out of fifty states have balanced budget requirements mandating annual budget reconciliation (Seegert, 2016, p. 1). When states face budget shortfalls, limited financing options force state governments to cut services (Seegert, 2016, p. 3). When facing such financial constraints, the ability for states to collect steady tax revenue streams is critical for the consistent and predictable provision of state government services.

Over the past three decades, state tax revenue volatility rose from 2.9% to 10.8% by 2016 (Seegert, 2016, p. 1). While numerous factors influence the consistency of revenue collections, changes in tax policy account for 59-67% of all state tax revenue volatility, while GDP and all other differences account for 13-22% and 19% respectively (Seegert, 2016, p. 3). This finding suggests that state governments have considerable control over their revenue stability through thoughtful tax policy design.

Cornia and Nelson (2010) note that while state governments cannot directly alter the volatility and growth of their economies, they can adjust their tax portfolios to minimize the effects of the business cycle and enable the steady provision of services (p. 55). Similarly, research suggests that states should design tax portfolios that ensure revenue stability in times of economic uncertainty. Sobel and Wagner (2003) argues that highly volatile taxes with high income elasticities may appear useful when revenues are high but create substantial challenges for states during economic downturns. Of course, many decisions are highly dependent on the structure of a state's economy. States with low levels of growth might be willing to accept some revenue volatility for higher growth potential, while other states may opt to prioritize stability at the expense of potentially higher revenues (Cornia & Nelson, 2010, p. 23).

Growing research addresses the specific aspects of tax policy that drive revenue volatility. Some researchers argue that greater diversification of tax revenue sources could lead to more stable revenue collections (Kwak, 2013; Yan, 2012). Others focus specifically on the imposition of taxes on capital gains income as a driver of revenue volatility (Dauchy & Balding, 2013; Williams & Vasché, 2005). However, a vast portion of the literature remains focused on the volatility effects of specific tax types, both broad and narrow.

Income taxes are generally understood to be more volatile than sales taxes (Kwak, 2013; Mattoon & McGranahan, 2012; Seegert, 2016, p. 5). The majority of studies identify corporate income taxes as the most volatile general tax type due to their substantial reliance on business cycle fluctuations (Cornia & Nelson, 2010; Kwak, 2013; Mattoon & McGranahan, 2012). Seegert argues that revenue volatility increased substantially during the dot-com bubble, while Boyd notes that income and sales tax collections were notably volatile following the 2001 and 2007 recessions (Boyd, 2022; Seegert, 2016). The effects of the COVID-19 pandemic lockdowns on revenue volatility are yet to be fully understood.

By comparison, sales taxes are understood to be less volatile (Boyd, 2022; Cornia & Nelson, 2010; Kwak, 2013; Mattoon & McGranahan, 2012). However, the volatility of selective sales taxes in relation to general sales taxes is unclear, in part because not all selective sales taxes are equally volatile. Kwak (2013) finds sales taxes on motor vehicles and non-durable goods to be strong revenue stabilizers. Similarly, Cornia and Nelson (2010) finds sales taxes on motor fuels and motor vehicle licenses are revenue stabilizers, while also highlighting the stabilizing effect of taxes on alcoholic beverages (p. 33). However, they find tobacco sales taxes to be destabilizing, partially due to legislated changes in their imposition (Cornia & Nelson, 2010, p. 32).

Boyd (2022) elaborates on the influence of legislated changes, noting that sales taxes are often imposed on a quantity of goods rather than their value, so tax revenue collections often do not keep pace with inflation (p. 10). Legislators often forgo raising sales taxes to keep pace with prices, meaning the extent to which sales taxes act as revenue stabilizers could be explained by the extent to which legislative changes are made (Boyd, 2022; Cornia & Nelson, 2010). Such findings suggest that apparent revenue stability may result from a lack of efforts to ensure revenue growth keeps pace with inflation rather than true economic stability. Walczak (2019) also notes that many states have adopted individual income tax codes with inflation adjustments to ensure collections rise with inflation.

Conversely, Yan (2012) finds that increased reliance on sales taxes raises revenue instability, arguing that portfolio diversification, along with the imposition of license taxes, leads to higher revenue stability. Severance taxes on the extraction of natural resources are understood to be incredibly volatile; however, only a few states rely heavily on such revenues (Boyd, 2022, p. 5). Schaufele (2016) finds that resource volatility substantially increases GDP volatility in Canadian provinces Newfoundland, Labrador, Alberta, and Saskatchewan, suggesting that tax revenue volatility in these jurisdictions may not be due solely to the imposition of taxes on severance revenues, but rather the volatile nature of these areas' economies (p. 469).

Research indicates state tax revenue volatility is a growing issue that directly impedes states' ability to provide critical services, but that states can design tax portfolios to withstand uncertain economic conditions. The literature establishes that income taxes, specifically corporate income taxes, increase revenue volatility, while sales taxes generally have a stabilizing effect on revenue collections. Specific studies highlight the stabilizing nature of particular types of selective sales taxes, while others note underlying factors that may influence the volatility of destabilizing selective sales taxes. However, significant gaps remain in understanding the full spectrum of tax source volatility. The effects of license and severance taxes remain less clear, both due to a lesser research focus on these taxes and the reality that these tax sources comprise a smaller percentage of state tax portfolios. A comprehensive analysis examining all tax sources simultaneously is needed to provide policymakers with clear guidance on tax portfolio design.

Methodology

To test the relationship between state tax portfolios and revenue volatility, this study employs multiple linear regression models. We analyze the six broad categories of tax revenue using four distinct model specifications to thoroughly address their influence and ensure the robustness of our findings. All statistical analysis, including regression modeling and influence diagnostics, was conducted in RStudio using data prepared in Microsoft Excel.

Tax portfolio data from the US Census Bureau's Annual Survey of State Government Tax Collections Revenues are categorized into six general types: property, general sales, selective sales, income, license, and other (U.S. Census Bureau, 2024). Table 1 presents the average composition of state tax portfolios across all states. For this analysis, "Selective Sales", "Income", "License", and "Other" are treated as aggregates of their constituent specific taxes (e.g. tobacco, motor fuels sales taxes within "Selective Sales"). Selective sales taxes apply to the purchase of goods, while license taxes charge businesses and consumers for permission to operate or use those goods.

The volatility of state tax revenue is measured by calculating the mean and standard deviation of the last 20 years of annual total collections, then dividing the standard deviation by the mean to produce a coefficient of variation. Calculating the volatility as a coefficient of variation (CV) allows us to compare states with different tax bases. For example, a CV of 15% means revenue typically varies by 15% from the mean, regardless of whether that mean is \$2B or \$200B. The distribution of our state tax revenue volatility readings are demonstrated below in Figure 1.

The reliance on each tax type is measured as a percentage of the total state tax revenue it generates in a given year. To address short-term fluctuations and capture structural reliance, we calculated a 20-year average of these

Table 1: Average State Revenue Dependence by Tax Type (2005-2024)

General Tax Types	Specific Tax Types	Average Composition
Property	Property Taxes	2.95%
General Sales	General Sales	30.49%
Selective Sales	Alcoholic Beverages, Amusements, Insurance Premiums, Motor Fuels, Pari-Mutuels, Public Utilities, Tobacco Products, Other Sales	16.71%
License taxes	Alcoholic Beverages, Amusements, Corporations, Hunting/Fishing, Motor Vehicles, Motor Vehicle Operators, Public Utilities, Occupation/Business, Other License	6.77%
Income Taxes	Personal Income, Corporate Income	37.34%
Other Taxes	Death/Gift, Documentary/Stock Transfer, Severance, Other	5.75%

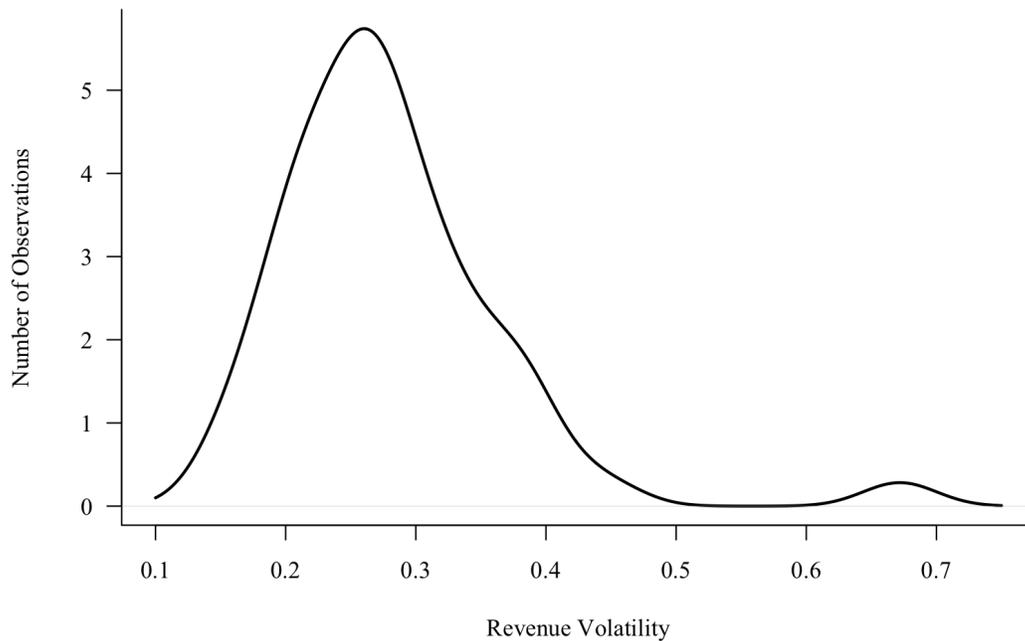


Figure 1: Distribution of State Tax Revenue Volatility

percentages for all 50 states. This approach mitigates distortionary effects from rare, drastic policy changes (e.g., the elimination or introduction of a major tax) and provides a stable measure of each state's tax portfolio.

The relationship between tax portfolio composition and revenue volatility is estimated using the following regression equation, where y represents revenue volatility (coefficient of variation), the β coefficients capture the effect of each tax type's share on volatility, C_i represents the control variables, and ϵ_i is the error term:

$$\hat{y}_i = \beta_0 + \beta_1(\text{Property Tax})_i + \beta_2(\text{General Sales})_i + \beta_3(\text{Selective Sales})_i + \beta_4(\text{License Taxes})_i \\ + \beta_5(\text{Income Taxes})_i + \beta_6(\text{Other Taxes})_i + \gamma C_i + \epsilon_i$$

To control major economic and demographic factors impacting revenue collection across the 20-year period, three control variables were used: average state-level private product remaining, average population change, and average unemployment rate. Varying levels of economic growth often have a substantial impact on revenue collections (Seegert, 2016). We account for this using private product remaining, a calculation of economic growth that captures private domestic consumption and investment. We use this instead of state GDP data to avoid double counting taxes collected and tax monies spent. This was factored in as the average year-over-year percentage change in private

product remaining, sourced by the Bureau of Economic Analysis (Bureau of Economic Analysis, 2024). Factoring in the average population change protects the model from changes in the tax base that are independent of tax policy. The average year-over-year percentage change in state population was collected from the Federal Reserve Bank of St. Louis (Federal Reserve Economic Data, 2024). Finally, the annual state unemployment rate across the 20-year period was also sourced from the Federal Reserve Bank of St. Louis to account for the drop in the income tax eligible population (Federal Reserve Economic Data, 2025).

Averaging these controls aligns them with the 20-year averages of the independent and dependent variables, creating a cross-sectional dataset suitable for analyzing the structural determinants of revenue volatility while holding economic, demographic, and labor market conditions constant.

Due to the compositional nature of the data (the six shares of tax revenue sum to 100%), the independent variables exhibit perfect multicollinearity. To resolve this, we adopt the standard approach of omitting one category from each regression, which becomes the reference case. We ran several of these regressions rotating the omitted tax type. The results from the three most informative specializations are reported, and we assessed the severity of the remaining multicollinearity using the Variance Inflation Factor (VIF), which is a standard diagnostic test for multicollinearity with linear regressions. We used a composite influence plot to identify outliers (points with high residuals) and high-leverage points (points with unusual predictor values), both of which can unduly influence the model. For simplicity, we refer to both of these observations as “outliers” throughout the article.

For the models themselves, we used a model specifically for broad tax categories, and then a model specifically for specific tax types. While all four model specifications were run for the specific tax types comprising each broad tax type, only models with compelling results are included in this analysis. The model for broad tax categories analyzed the six broad categories using four distinct model specifications—a Baseline Model, a Controlled Model, Robustness Model excluding outliers, and a Full Robustness Model—to thoroughly address their influence and ensure the consistency of our findings.

Model 1, the Baseline Model, does a simple regression of revenue volatility on the six broad tax categories, showing the raw correlation between tax portfolios and revenue volatility. Model 2, the Controlled Model, adds control variables to the baseline specification to isolate the relationship between volatility and tax structure from underlying economic and demographic conditions. Models 3 and 4, the outlier-excluded models, are crucial for addressing the generalizability of the findings. While states with atypically high reliance on specific revenue sources may not be standard, their experience with such dependence on certain sources may provide insights for other states considering dependence on similar sources. However, the experience of a state with extreme dependence on a specific source would be less useful for a state considering a marginal change to a more standard tax portfolio.

Results

Modeling tax portfolios and revenue volatility gives practitioners important context when trying to answer different policy-relevant questions. At its core, the model estimates how structural reliance on different tax types (the composition of a state’s tax portfolio) predicts revenue stability over time, while accounting for the economic environment in which that state operates. For practitioners, the model will allow them to ask questions such as, “If State A generates 35% of revenue from general sales taxes and 7% from property taxes while State B generates 39% from income taxes and 3% from property taxes, how does the compositional difference affect revenue volatility assuming everything else remains constant?”

As mentioned previously, we use four different models, both for the general and specific tax types. This section will first detail the four general models used before specific tax models.

General Tax Type

Our results reveal varying volatility effects across tax types, with findings highly sensitive to model specifications and baseline choices. Other taxes consistently exhibit the highest volatility. General sales and income taxes show stabilizing effects relative to other taxes in Models 1-3, with property taxes demonstrating a stabilizing effect in Model 1, and selective sales taxes exhibiting a stabilizing effect in Models 1 and 3. However, other taxes show no statistically significant effect in the full robustness model.

Selective sales taxes demonstrate the most reliable stabilizing effect. They significantly stabilize revenues relative to other taxes in Models 1 and 3, and relative to general sales and income taxes in Model 3. They lose their significance in Model 4. License taxes show no statistically significant influence on revenue volatility. General sales taxes demonstrate strong stabilization when other taxes are used as a baseline in Models 1 and 2, with diminished but significant stabilizing effects in Model 3. Effects against the other taxes baseline are insignificant in Model 4. Property taxes stabilize revenues versus other taxes in Model 1, but lose significance in Models 2-4. Income taxes demonstrate a strong stabilizing effect when other taxes are used as a baseline in Models 1-3.

Table 2: Broad Tax Type Models Tax Volatility

	Model 1 ^a			Model 2 ^b		
	General Sales	Income	Other	General Sales	Income	Other
Property Taxes	-0.211 (0.186)	-0.274 (0.188)	-0.655** (0.221)	-0.087 (0.186)	-0.168 (0.184)	-0.451 (0.227)
General Sales Taxes	-0.063 (0.073)		-0.445** (0.110)	-0.081 (0.075)		-0.363** (0.112)
Selective Sales Taxes	-0.231 (0.226)	-0.294 (0.206)	-0.675** (0.210)	-0.224 (0.293)	-0.305 (0.262)	-0.588 (0.293)
License Taxes	0.314 (0.193)	0.251 (0.199)	-0.130 (0.220)	0.354 (0.188)	0.274 (0.199)	-0.009 (0.216)
Income Taxes		0.063 (0.073)	-0.381** (0.095)		0.081 (0.075)	-0.283* (0.107)
Other Taxes	0.445** (0.110)	0.381** (0.095)		0.364** (0.112)	0.283* (0.107)	
Observations	50	50	50	50	50	50
R ²	0.387	0.387	0.387	0.478	0.478	0.478
Adjusted R ²	0.317	0.317	0.317	0.377	0.377	0.377
F Statistic	5.555** (df = 5; 41)			4.701** (df = 5; 41)		

Note: *p<0.05, **p<0.01

^a Baseline Model (no economic controls, all states included)

^b Controlled Model (economic controls, all states included)

Specific Sales Tax

To identify the components that drive sales tax volatility, we disaggregate selective and general sales taxes and apply our model to each subcategory. Motor fuel taxes exhibit the most consistent stabilizing effect across specifications, significant in Models 5-6 and 8. They lose significance in Model 7. Pari-mutuels taxes (primarily on gambling) significantly stabilize revenues in Models 5-6, but become insignificant when outlier states are excluded. General sales taxes stabilize revenues in Models 5-6, but show no significant effect in Models 7-8. Other sales taxes stabilize revenues in Model 7, with no significant influence in other specifications. Taxes on alcoholic beverages, amusements, insurance premiums, public utilities, and tobacco products do not show any statistically significant influence on revenue volatility across any specification.

Other Tax

While excluded from this article for brevity, our model ran all four specifications for specific other tax types. Across all four specifications, severance taxes on natural resource extraction was the only tax type to demonstrate any effect on revenue collection volatility. Severance taxes demonstrated a strong destabilizing effect in the baseline, controlled, and robustness model, although the effect did diminish in the robustness model. In the full robustness model, severance taxes did not have a statistically significant influence on revenue volatility. Due to the fact that severance taxes are the only apparent driver of revenue volatility within the other taxes category, the volatility of severance taxes is well represented through the other taxes category in the general model.

Discussion

Our findings both confirm and challenge prevailing assumptions about tax revenue volatility, offering key insights for state policymakers seeking to balance fiscal stability with revenue adequacy.

Table 3: Broad Tax Type Models Tax Volatility

	Model 3 ^c			Model 4 ^d		
	General Sales	Income	Other	General Sales	Income	Other
Property Taxes	-0.034 (0.152)	-0.007 (0.157)	-0.350 (0.193)	0.015 (0.150)	0.022 (0.151)	-0.111 (0.202)
General Sales Taxes	0.027 (0.060)		-0.316* (0.133)	0.007 (0.060)		-0.125 (0.154)
Selective Sales Taxes	-0.586** (0.201)	-0.559** (0.186)	-0.901** (0.227)	-0.366 (0.258)	-0.360 (0.238)	-0.492 (0.301)
License Taxes	0.462 (0.324)	0.489 (0.319)	0.146 (0.344)	0.183 (0.345)	0.190 (0.345)	0.058 (0.346)
Income Taxes		-0.027 (0.060)	-0.343** (0.123)		-0.007 (0.060)	-0.132 (0.146)
Other Taxes	0.316* (0.133)	0.343* (0.123)		0.125 (0.154)	0.132 (0.146)	
Observations	47	47	47	47	47	47
R ²	0.302	0.302	0.302	0.426	0.426	0.426
Adjusted R ²	0.216	0.216	0.216	0.306	0.306	0.306
F Statistic	3.542** (df = 5; 41)			3.530** (df = 5; 41)		

Note: *p<0.05, **p<0.01

^c Robustness Model (no economic controls, outlier-states excluded)

^d Full Robustness Model (economic controls, outlier-states excluded)

Income Taxes

The literature consistently identifies income taxes, particularly corporate income taxes, as highly volatile revenue sources due to their sensitivity to business cycles (Cornia & Nelson, 2010; Kwak, 2013; Mattoon & McGranahan, 2012; Seegert, 2016). Our findings provide some support for this view. Income taxes demonstrate stabilizing effects relative to other taxes in Models 1-3, suggesting they are less volatile than the most unstable revenue sources. However, income taxes show no statistically significant effect on revenue volatility in our full robustness model (Model 4), which controls for economic conditions and excludes outlier states.

This pattern suggests that income tax volatility may be more context-dependent than the literature implies. States with stable economic conditions and diversified economies may experience less income tax volatility than studies focused on recession periods would suggest. Conversely, states with economies heavily dependent on specific industries or with unusual tax structures may experience the high volatility that prior research emphasizes. This means that concerns about income tax volatility should be evaluated within each state's specific economic context rather than assumed to be universal.

Sales Taxes

Our results support the literature's general consensus that sales taxes serve as revenue stabilizers. General sales taxes demonstrate significant stabilizing effects in our baseline and controlled models (Models 5-6), aligning with findings from Boyd (2022), Cornia and Nelson (2010), Kwak (2013), and (Mattoon & McGranahan, 2012). However, this stabilizing effect loses statistical significance in our robustness and full robustness models (Models 7-8), suggesting it may be driven partially by outlier states or underlying economic conditions rather than inherent characteristics of general sales taxation alone.

The most striking finding of this analysis is the stabilizing effect of selective sales taxes. Unlike other tax categories whose effects diminish or disappear in more conservative model specifications, selective sales taxes demonstrate consistent revenue stabilization. In our general tax type robustness model (Model 3), selective sales taxes emerge as a significant stabilizer relative to general sales, income, and other taxes. The stabilizing effect diminishes when economic controls are added in Model 4, demonstrating selective sales taxes potential to limit revenue volatility during periods of economic uncertainty.

Table 4: Specific Tax Type Models

Tax Type	Model 5 ¹	Model 6 ²	Model 7 ³	Model 8 ⁴
General Sales	-0.199* (0.082)	-0.194** (0.067)	-0.089 (0.080)	-0.084 (0.067)
Alcoholic Beverages	4.947 (2.630)	3.390 (2.258)	0.952 (2.471)	1.467 (2.116)
Amusements	0.117 (0.572)	0.370 (0.497)	-1.460 (0.776)	-0.245 (0.706)
Insurance Premiums	1.042 (1.238)	1.332 (1.012)	0.743 (1.050)	0.786 (0.882)
Motor Fuels	-1.632* (0.715)	-1.723* (0.647)	-1.054 (0.643)	-1.564* (0.577)
Pari-Mutuels	-64.935* (30.667)	-74.686** (26.586)	-20.792 (33.791)	-31.193 (28.691)
Public Utilities	-0.375 (0.788)	0.148 (0.662)	-0.123 (0.660)	0.159 (0.562)
Tobacco Products	-0.126 (0.287)	-0.124 (0.244)	-0.286 (0.244)	-0.218 (0.210)
Other	-0.680 (0.382)	0.130 (0.690)	-0.776* (0.348)	-0.127 (0.623)
Observations	50	50	46	46
R ²	0.415	0.644	0.308	0.558
Adjusted R ²	0.283	0.528	0.134	0.391
F Statistic	3.151** (df = 9; 40)	5.569** (df = 12; 37)	1.777 (df = 9; 36)	3.471** (df = 12; 33)

Note: *p<0.05, **p<0.01

¹ Baseline General and Selective Sales Taxes (no economic controls, no excluded states)

² Controlled Sales Taxes (economic controls, all states included)

³ Robustness Model (no economic controls, outlier-states excluded)

⁴ Full Robustness Model (economic controls, outlier-states excluded)

Yan (2012) assertion that increased reliance on sales taxes raises revenue instability, and Cornia and Nelson (2010) identification of alcoholic beverages taxes as a revenue stabilizer, and tobacco taxes as a revenue destabilizer were not supported by our analysis. We found no statistically significant relationship between the imposition of alcoholic beverages or tobacco taxes on revenue volatility and found sales taxes, particularly selective sales taxes, to be a strong revenue stabilizer.

Motor Fuels

Our disaggregated analysis reveals that motor fuels taxes are the primary driver of selective sales tax stability, demonstrating significant stabilizing effects in Models 5-6 and 8. This finding strongly corroborates Cornia and Nelson (2010) and Kwak (2013) identification of motor fuels taxes as revenue stabilizers. The consistency of motor fuel tax collections likely reflects the relative price inelasticity of gasoline demand. Consumers continue purchasing fuel for commuting and essential travel even during economic downturns, providing states with steady revenue streams regardless of economic conditions.

This stability, however, comes with an important limitation noted by Boyd (2022) and Cornia and Nelson (2010). Selective sales taxes, including motor fuels taxes, are typically imposed on quantity rather than value, meaning their revenue collections do not automatically keep pace with inflation. The apparent stability of motor fuels tax revenues may therefore reflect static revenue rather than dynamic economic resilience. Policymakers relying on motor fuels taxes for stability must therefore implement periodic legislative adjustments to maintain revenue adequacy, creating a trade-off between the benefits of predictable collections and political challenges of regular tax increases.

Pari-Mutuels

Our findings regarding pari-mutuel taxes (primarily gambling revenues) present a more nuanced picture. These taxes demonstrate substantial stabilizing effects in baseline and controlled models (Models 5-6) but lose significance when outlier states are excluded (Models 7-8). This pattern suggests that gambling revenues provide stability primarily in states with substantial gaming industries (Nevada being the most obvious example), but do not represent a generalizable revenue stabilizer for typical states.

This finding aligns with the broader literature's recognition that not all selective sales taxes exhibit uniform volatility characteristics. It also emphasizes the importance of employing multiple model specifications: a researcher examining only Models 5-6 might conclude that all states should expand gambling taxation for revenue stability, while Models 7-8 reveal this recommendation applies only to states with already-substantial gaming sectors. For most states considering moderate expansions of gambling, pari-mutuel taxes are unlikely to meaningfully reduce revenue volatility.

Other

Our findings regarding "Other" taxes (which include death and gift taxes, documentary and stock transfer taxes, and severance taxes) present the study's most complex results. In baseline and controlled models (Models 1-2), other taxes consistently emerge as highly volatile, with nearly all other tax types demonstrating stabilizing effects by comparison. This volatility diminishes but remains statistically significant in the robustness model (Model 3), which excludes outlier states, but not economic controls.

However, other taxes show no statistically significant effect on revenue volatility in our full robustness model (Model 4), which both excludes outlier states and includes economic controls. This pattern strongly suggests that the apparent volatility of other taxes is driven primarily by resource-dependent outlier states, with Alaska and Wyoming being the most prominent examples, rather than representing a generalizable characteristic of these revenue sources.

Our supplementary analysis of disaggregated other taxes confirms this interpretation. Severance taxes on natural resource extraction demonstrate strong destabilizing effects in the baseline, controlled and robustness model, but this effect disappears when economic controls are added and outlier states excluded in the full robustness model. This finding aligns with Schaufele (2016) research on Canadian provinces, where apparent tax revenue volatility in resource-dependent regions may reflect underlying economic volatility rather than inherent characteristics of the tax itself.

For policymakers, the implication is clear: typical states without substantial natural resource extraction need not avoid "Other" taxes purely on volatility grounds. Severance taxes are indeed highly volatile, but only for the small number of states who rely heavily on them.

Property Taxes

Property taxes demonstrate statistically significant effects in Models 1-2 when compared to other taxes but show no significant influence in Models 3-4. This pattern suggests property tax stability may be overstated in simpler models that do not account for outlier states and economic conditions. Interestingly, the literature provides limited attention to property tax volatility at the state level, focusing instead on local government property taxation. Our findings suggest state-level property taxes may not provide the stability advantages often attributed to property taxation generally.

License Taxes

License taxes do not demonstrate a significant effect in Models 1-4. A lack of significant findings, combined with license taxes' relatively small share of state revenues (6.77%), suggest they are neither a major source of stability or instability for most states. License taxes may be best understood as minor revenue sources whose volatility effects are masked by larger tax categories and economic conditions.

Limitations and Future Research

Several limitations of this study suggest directions for future research. Our use of 20-year averages, while providing stable measures of structural tax reliance, may mask important temporal dynamics in the relationship between tax portfolios and volatility. Future research could employ panel data methods to examine how this relationship evolves over time and in response to specific economic shocks. For example, panel data would reveal whether states that increased their reliance on selective sales taxes after the Wayfair decision experienced a subsequent decline in revenue volatility relative to non-adopting states (*South Dakota v. Wayfair, Inc.*, 2018).

Additionally, while our control variables account for broad economic and demographic conditions, they do not capture state-specific factors such as tax base breadth, exemptions or rate structures that mediate the relationship between tax composition and volatility. More granular analysis could provide deeper insights into the mechanisms driving revenue stability.

Our analysis also focuses exclusively on volatility as an outcome, but policymakers must balance stability against other objectives including revenue adequacy, economic growth, and equity. Future research incorporating multiple policy objectives could provide more comprehensive guidance for tax portfolio design, helping policymakers understand when pursuing stability requires accepting trade-offs and when it aligns with other goals.

Finally, our findings apply specifically to the 2005-2024 period and may not be generalizable to different economic contexts. The relative stability of motor fuels taxes, for instance, may diminish as electric vehicle adoption accelerates. Similarly, the growth of remote work and e-commerce may alter the volatility characteristics of sales and income taxes in ways historical analysis cannot capture.

Conclusion

This comprehensive analysis of state portfolio composition and revenue volatility demonstrates that while general patterns exist, the relationship between tax structure and fiscal stability is more nuanced and context-dependent than prior literature suggests. Selective sales taxes, particularly motor fuels taxes, emerge as the most persistent revenue stabilizer, maintaining their stabilizing effects even when outlier states and economic controls are accounted for. General sales taxes show stabilizing effects that are less stable to model specification changes, while income taxes exhibit context-dependent volatility that may be less universal than commonly assumed. Other taxes, including severance taxes, demonstrate high volatility driven primarily by resource-dependent outlier states rather than representing a generalizable characteristic.

These findings underscore three critical lessons for policymakers and researchers. First, analyzing tax portfolios comprehensively rather than examining individual tax types in isolation provides more accurate guidance for policy design. Secondly, state-specific contexts matter enormously, as recommendations for Alaska or Nevada may be irrelevant or counterproductive for Iowa or Virginia. Third, the robustness of findings to different model specifications reveals important nuances that simpler analyses miss.

For states prioritizing revenue stability, increasing reliance on selective sales taxes offers the most promising path. However, policymakers must recognize that this stability comes at a cost: selective sales taxes require periodic legislative adjustments to maintain real revenue levels, creating political challenges that may offset the benefits of predictable revenues. For states with resource-dependent tax portfolios, having substantial severance tax revenues puts them at risk during economic downturns. For states with ordinary revenue structures, they should recognize that much conventional wisdom about tax volatility may be overstated for states with relatively standard tax portfolios and economic conditions. For states considering modest, marginal changes to typical tax portfolios, they do not need to be overly concerned about volatility effects.

For state policymakers seeking to promote fiscal stability while maintaining adequate revenues, the path forward requires balancing multiple objectives. Selective sales taxes offer genuine stability advantages but require active management to maintain real revenue levels. By comparison, resource-dependent revenues generate enormous boom-time collections but create severe bust-time challenges. The optimal tax portfolio depends on each state's economic structure, policy priorities, and willingness to accept trade-offs between stability and other goals.

Ultimately, this research confirms that states possess considerable control over their revenue volatility through thoughtful tax portfolio design. While economic cycles remain largely beyond state control, the composition of state tax systems can either amplify or dampen their fiscal effects. By understanding which tax sources provide stability across diverse contexts and which are context-dependent, policymakers can design tax structures that promote consistent service provision even amid economic uncertainty.

References

- Boyd, D. (2022). *State tax revenue volatility and its impact on state governments* (Tech. Rep.). The Pew Charitable Trusts. Retrieved from <https://www.pew.org/-/media/assets/2022/07/state-tax-revenue-volatility-and-its-impact-on-state-governments.pdf>
- Bureau of Economic Analysis. (2024). *Gdp by state*. Retrieved 2025, November 12, from [url{https://www.bea.gov/data/gdp/gdp-state}](https://www.bea.gov/data/gdp/gdp-state) (U.S. Department of Commerce)
- Cornia, G. C., & Nelson, R. D. (2010). State tax revenue growth and volatility. *Regional Economic Development*, 6(1), 23–58. Retrieved 2025-09-02, from <https://research.stlouisfed.org/publications/red/2010/01/Cornia.pdf> (Federal Reserve Bank of St. Louis)

- Dauchy, E. P., & Balding, C. (2013). Federal income tax revenue volatility since 1966. *SSRN Electronic Journal*. Retrieved from <https://doi.org/10.2139/ssrn.2351376>
- Federal Reserve Economic Data. (2024). *Total population by state*. Federal Reserve Bank of St. Louis. Retrieved 2025, September 2, from [\url{https://fred.stlouisfed.org/tags/series?t=population%3Bstate}](https://fred.stlouisfed.org/tags/series?t=population%3Bstate)
- Federal Reserve Economic Data. (2025). *Unemployment rate by state*. Federal Reserve Bank of St. Louis. Retrieved 2025, September 2, from [\url{https://fred.stlouisfed.org/release/tables?eid=840687&rid=116}](https://fred.stlouisfed.org/release/tables?eid=840687&rid=116)
- Kwak, S. (2013). Tax base composition and revenue volatility: Evidence from the U.S. states. *Public Budgeting & Finance*, 33(2), 41–74.
- Mattoon, R., & McGranahan, L. (2012). *State tax revenues over the business cycle: Patterns and policy responses* (No. 299). Retrieved from <https://www.chicagofed.org/publications/chicago-fed-letter/2012/june-299>
- Schaufele, B. (2016). Taxes, volatility, and resources in Canadian provinces. *Canadian Public Policy*, 42(5), 469–481.
- Seegert, N. (2016). Tax revenue volatility: Evidence from US states. *SSRN Electronic Journal*, 3. doi: 10.2139/ssrn.2789889
- Sobel, R., & Wagner, G. (2003). The impact of revenue diversification and economic base on state revenue stability. *Journal of Public Budgeting, Accounting & Financial Management*. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=437500
- South Dakota v. Wayfair, Inc. (2018). *South Dakota v. Wayfair, Inc.* (Vol. 585). United States Supreme Court. (Retrieved from https://www.supremecourt.gov/opinions/17pdf/17-494_j4e1.pdf)
- U.S. Census Bureau. (2024). *Annual survey of state government tax collections (STC)*. Retrieved 2025, September 2, from [\url{https://www.census.gov/programs-surveys/stc.html}](https://www.census.gov/programs-surveys/stc.html)
- Walczak, J. (2019). *Inflation adjusting state tax codes: A primer*. Tax Foundation. Retrieved from <https://taxfoundation.org/research/all/state/inflation-adjusting-state-tax-codes/>
- Williams, B., & Vasché, J. D. (2005). *Revenue volatility in California*. Retrieved from https://lao.ca.gov/2005/rev_vol/rev_volatility_012005.htm
- Yan, W. (2012). The impact of revenue diversification and economic base on state revenue stability. *Journal of Public Budgeting, Accounting & Financial Management*. doi: <https://doi.org/10.1108/JPBAFM-24-01-2012-B003>

An Analysis of U.S. GDP and its Effect on the EUR/USD Exchange Rate

Ethan Juhasz

This paper examines the effect of various macroeconomic variables such as GDP, interest rates, inflation and unemployment on the spot exchange rate between the U.S. Dollar and the Euro. While many classical economic theories offer explanations of the various macroeconomic variables that affect the numerous variables in the real world continue to challenge attempts to model and predict currency exchanges. Research appears to be divided on this topic; while some authors have found statistically significant models, others claim that attempts to predict exchange rates fail to differ themselves from a stochastic "random walk model" in any meaningful way. This paper utilized multiple models including univariate, multivariate, distributed lag and differenced regression models and found that while some variables including the Consumer Price Index are significant in predicting the USD/EUR exchange rate, most other variables including the primary regressor USGDP have either a weak or insignificant relationship with the exchange rate. These findings provide initial steps and direction for future research with more advanced modeling.

Introduction

THE market for foreign exchange has grown rapidly in recent years, with traders and investors looking for new ways to predict future exchange rates to make more informed financial decisions. One of the most important exchange rates is that between the United States Dollar (USD) and the Euro (EUR). Dal Bianco, Camacho, and Perez-Quiros argue that "fluctuations in the euro-dollar exchange rate are crucial not only for the economic transactions between the two major economic blocks but also for the rest of the countries." (Dal Bianco, Camacho, & Quiros, 2012) Researching how macroeconomic variables can be used to track the power of the United States dollar compared to other global currencies will not only aid foreign exchange investors but also assist policymakers in measuring U.S. global economic power.

This paper attempts to discern the effect of U.S. real gross domestic product (GDP) on the exchange rate between the USD and EUR. The dependent variable is the spot USD-EUR exchange rate, which is the monthly average value of the US Dollar in terms of Euros. The primary independent variable in this study is U.S. real GDP, measured in billions of 2017 U.S. dollars. The author hypothesizes that an increase in real GDP will lead to an increase in the exchange rate; since a higher real GDP generally indicates a stronger economy, this will increase demand for the currency from international investors and drive up the exchange value of the U.S. dollar.

Literature Review

Previous research has been done studying similar topics. Gwarnicki and Bahmani (2021) performed a multivariate time series analysis on the closing price of the USD/EUR exchange rate for the New York Stock Exchange, with dependent variables including United States GDP, European Union GDP, United States Interest Rates set by the US Federal Reserve for US-based investments, European Union Interest Rates for European Central Bank EU-based investments, United States Inflation, European Union Inflation, United States Unemployment Rate, and European Union Rate, as well as a number of dummy variables were employed for both interest rates as well as the day of the week. (Gwarnicki & Bahmani, 2021) For the period spanning 2013-2019, they found that "announcements for the variables representing the U.S. GDP, EU interest rate, and U.S. interest rate adjustments were statistically significant." (Gwarnicki & Bahmani, 2021) Gwarnicki and Bahmani argue that announcements in macroeconomic variables affect investor sentiment which therefore influence exchange rates; this is also corroborated by a study by McCoy (2020) on the effect of U.S. policy rates on the EUR/USD exchange rate. McCoy specifically studied interest rates and their effect on the exchange rate, and argues that, "Central bank quantitative easing (QE) programs, by compressing the term premium component of domestic bond yields, encouraged investors to rebalance their portfolios towards foreign, higher yielding assets. The induced capital outflows... have likely exerted a downwards pressure on the EUR/USD." (McCoy, 2020) Gwarnicki, Bahmani, and McCoy ultimately argue that the influence of macroeconomic variables on investor sentiment is powerful enough to affect exchange rates significantly.

In contrast, Dal Bianco et al. conducted a similar study with regressions utilizing macroeconomic fundamentals and regressions deliberately excluding them. The authors found that when predicting the USD/EUR exchange rate, it is difficult to beat the 'random walk model' which predicts values as stochastic, independent steps back and forth. Dal Bianco et al. argued that, "although in any case the reduced model can improve the random walk, all the equal accuracy tests considered in the table indicate that the differences in forecasting accuracy are not statistically

significant.” (Dal Bianco et al., 2012) Though Dal Bianco et al. cite the utilization of the monetary model when studying the effects that macroeconomic variables have; one should note that while the monetary model does include many macroeconomic variables including money supply growth, industrial growth, interest rates, and trade balances it does not include GDP specifically (Dal Bianco et al., 2012). Therefore, while Dal Bianco et al. do generally argue that most macroeconomic variables are not effective at beating the random-walk model, their utilization of a specific model leaves room for further research. While most literature indicates macroeconomic variables have at least some level of significance in predicting investor sentiment, their predictive power is ultimately mixed. Dal Bianco et al. provides contrary evidence arguing that predictors are not statistically significant in predicting exchange rates, but their conservative variable choice leaves a gap in existing research this report aims to fill.

Methodology

Regression and Specifications

This study’s primary regression equation is as follows:

$$USDEUR_t = \beta_0 + \beta_1 USGDP_t + \beta_2 EUGDP_t + \beta_3 FEDFUNDS_t + \beta_4 USCPI_t + \beta_5 UNEMPLOYMENT_t + \beta_6 USGDP_{t-1} + \beta_7 EUGDP_{t-1} + \beta_8 FEDFUNDS_{t-1} + \beta_9 USCPI_{t-1} + \beta_{10} UNEMPLOYMENT_{t-1} + \epsilon$$

All data was gathered from the Saint Louis Federal Reserve Economic Data (FRED) and used 102 observations measured from quarter 1 of 1999 to quarter 2 of 2024. The variables include United States Real GDP (USGDP), European Union Real GDP (EUGDP), United States Federal Funds Rate (FEDFUNDS), United States Consumer Price Index (USCPI), United States Unemployment Rate (UNEMPLOYMENT); furthermore, this study utilizes a distributed lag model which includes a lagged version of each variable at time t-1. Table ?? in the appendix indicates descriptive statistics for each variable, and Figures ?? through ?? display scatter plots for each non-lagged independent variable.

One concern that became present was the threat of nonstationarity in the dependent variable, displayed by the line graph of USDEUR in Figure ?. A Dickey-Fuller test for unit root shown in Table ? further confirms this; with a p-value of 0.5451, the test fails to reject the null hypothesis that a unit root is present. Therefore, a differencing model is also added to the regression, using DIFFUSDEUR as the variable, which is a differenced version of the original USDEUR. Figure ?, which is a line graph depicting DIFFUSDEUR, indicates that differencing the dependent variable significantly reduces its non-stationarity and should lead to more accurate results.

Independent Variables

USGDP is Real United States GDP. It represents the aggregate fair market value of goods and services produced by the United States. It is measured in billions of 2017 USD. As mentioned earlier, the author predicts an increase in USGDP will lead to an increase in the exchange rate, as when the U.S. economy becomes stronger the demand for the currency will drive up the value of the U.S. dollar; this is also supported by the results from Gwarnci and Bahmani (Gwarnicki & Bahmani, 2021).

EUGDP represents European Union GDP, or the aggregate GDP of each European Union country measured in Millions of Euros. Theoretically, EUGDP will have an inverse relationship with USGDP as a stronger E.U. economy will lead to a decrease in demand for the U.S. Dollar and therefore a decrease in the exchange rate, ceteris paribus.

FEDFUNDS represents the U.S. Federal Funds Effective Rate, which influences important consumer interest rates like mortgages and loans. A higher Federal Funds rate aims to slow down the economy; theoretically an increase in interest rates will lead to capital outflows that exert downward pressure on USD/EUR rate as discussed by McCoy (McCoy, 2020).

USCPI represents an aggregate price index of consumer goods and services in the U.S. and is used to measure inflation. Theoretically, a higher Consumer Price Index indicates higher inflation and therefore a weaker U.S. Dollar, so an increase in the Consumer Price Index should lead to a decrease in the USD/EUR rate.

Inverse to USCPI, UNEMPLOYMENT represents the current U.S. unemployment rate percentage. Since an increase in unemployment is generally indicative of a weaker economy, an increase in UNEMPLOYMENT should theoretically weaken the exchange rate and decrease USD/EUR.

Omitted Variables

There are a few important variables omitted that ought to be noted for their potential ability to bias the results of this study. This study mostly analyzes the United States macroeconomic variables and how they influence the

exchange rate, but the regression does omit many macroeconomic variables on the European Union side of the exchange rate. While the European Union GDP is included, E.U. consumer price index and EU Unemployment rate are both omitted while the U.S. CPI and Unemployment rate are included. These variables were ultimately omitted due to time constraints and data limitations within the study, but they provide a direction for future research.

Following the standard equation for calculating bias:

$$BIAS_{X^2} = \beta_2 \alpha_1 \quad (5)$$

An increase in the E.U. Consumer Price Index indicates an increase in inflation for the E.U. This would lead to a decrease in the value of the Euro and an increase in the USD/EUR. In addition, it would indicate a weaker European Union economy and therefore influence a weaker EUGDP. Since in this case, 2 would be positive and 1 would be negative, omitting E.U. CPI would have a negative bias and therefore an underestimate of the coefficients.

E.U. unemployment would have a similar effect. A higher unemployment rate indicates a weaker economy, leading to a lower GDP and decreased demand for currency. Therefore, since 2 would also be positive and 1 would also be negative, omitting the European Union Unemployment rate would also lead to a negative bias and underestimated coefficients.

Results

While the results varied within each model, the final model indicated that there were a few variables including US GDP that were statistically significant predictors of the exchange rate between the dollar and the Euro. Table ?? represents a full table of univariate and multivariate regressions, both with and without lagged variables in order to reduce serial correlation. The main regression that will be analyzed in this study will be column four, the differenced regression including lag. The regression has a goodness of fit, or R2 value of 0.189.

The results of the equations fluctuated a lot between the models. In Model 4, the variables that proved to be significant were USGDP, USCPI, UNEMPLOYMENT, LAGUSGDP, and LAGUSCPI, with LAGUSCPI being the only variable that was significant at the 1% level. This makes sense and is in line with previous theoretical predictions; as the Consumer Price Index is a direct representation of inflation and the value of the dollar, it should make sense that it would have the most correlation with the exchange rate and demand for the currency. A 1-unit increase in the US Consumer Price Index is predicted to increase the difference in exchange rate between years by 0.00842, while a 1 unit increase in USCPI of the previous year is predicted to decrease the difference in exchange rate between the next two years by 0.0121. A 1-percentage point increase in the unemployment rate is predicted to increase the difference in exchange rate between years by 0.0221.

It appears that while the regression does confirm the initial hypothesis that United States GDP has a positive effect on the exchange rate, the significance of the primary predictor variable is at a lower level than other variables; USGDP is significant at the 10% level, and the model predicts that a 1 billion USD increase in USGDP is predicted to increase the difference in exchange rate between years by 9.74e-05. Converse to the hypothesis, though, LAGUSGDP was also significant at the 10% level and predicted that a 1 billion USD increase in the previous year's United States GDP would decrease the difference between exchange rates of the next two years by -9.59e-05.

These results appear similar to Gwarnicki and Bahmani's results, which used a very similar regression, yet had an R2 of 0.9965 and found USGDP to be significant at the 1% level. (Gwarnicki & Bahmani, 2021) However, Gwarnicki and Bahmani did not include lagged variables, though they did not find any evidence of autocorrelation with a Durbin-Watson statistic of 1.8860. While this model is only slightly different to the findings of Gwarnicki and Bahmani, the findings that macroeconomic indicators like Consumer Price Index and Unemployment are significant would disagree with Dal Bianco et al.'s assertion that macroeconomic indicators are not effective at predicting fluctuations in spot exchange rates. (Dal Bianco et al., 2012) Therefore, it appears that this study's model finds itself in the middle of previous research—while macroeconomic variables still are able to be predictors of spot exchange rate, the study finds that due to higher p-values and a low R^2 value there is still a high amount of random fluctuations unaccounted for by a regression utilizing only macroeconomic factors. While this could be due to lack of correlation between variables like Dal Bianco et al. assert, it is also likely that the differences in goodness of fit and variable significance could be due to errors in specification, which are discussed in the next section of this paper.

Post-Estimation

There are a few key post-estimation tests to consider – the three tests this study uses to affirm key classical assumptions are the Breusch-Pagan test for heteroskedasticity, the Durbin-Watson test for autocorrelation, and an analysis of Variance-Inflation Factor (VIF) scores for multicollinearity. The Breusch-Pagan test, shown in Table ??, returned a chi-squared value of 1.19 and a p-value of 0.2752, failing to reject the null hypothesis that states

there is constant variance. The Durbin-Watson test (Table ?? also indicates the presence of serial correlation. low d-score of 1.250408 provides evidence to reject the null hypothesis that there is no serial correlation within the data. Serial correlation can cause inaccurate beta coefficients leading to biased significance testing; this could be a possible explanation for the very different p-values between this study and that of Gwarnicki and Bahmani's. (Gwarnicki & Bahmani, 2021) Lastly, analysis of VIF scores (Table ??) indicates the model does appear to have extremely high presence of multicollinearity with a Mean VIF of 666.198. However, this is likely due to the presence of lagged variables which would clearly be associated with their non-lagged counterparts; therefore, while the presence of this multicollinearity should be noted it should not have significant effects on results. The risks of serial correlation and multicollinearity ultimately make the beta coefficients more imprecise, leading to biased significance testing. However, considering many of these variables were found to be significant despite biased coefficients, this indicates that these problems do not significantly affect the results of this study.

Conclusion

The regression models this study analyzes are most similar to that of Gwarnicki and Bahmani's study, with the same key independent and dependent variables as well as similar independent variables. Both models come to similar results as to the significance of macroeconomic variables in determining future exchange rates that go against Dal Bianco, Camacho and Quiros's findings on the superiority of the 'random walk' model; however, this study still has different findings to Gwarnicki and Bahmani. While Gwarnicki and Bahmani find much clearer results that United States GDP is a significant predictor of the spot exchange rate between the Dollar and the Euro, (Gwarnicki & Bahmani, 2021) this study has less goodness of fit and has higher p-values and lower coefficients, meaning that the conclusion as to the relationship between the key variables is less clear, and instead opts towards other macroeconomic variables like the Consumer Price Index and the United States unemployment rate as better predictors than GDP.

However, a more thorough and researched regression could prove different results. This study's model contained various violations of the classical assumptions including presence of serial correlation and multicollinearity, meaning this model's p-value and significance results could be called into question. Future research could include more omitted variables like EU interest rates, EU unemployment, or EU inflation to further reduce endogeneity; alternatively, it could include more lagged variables going back further in time than just one year to attempt to reduce serial correlation.

Ultimately, though, it appears this study finds that the best economic predictors for the USD/EUR exchange rate are macroeconomic statistics such as U.S. Unemployment and inflation for the previous year; therefore, investors into foreign exchange ought to continue to track these variables as an indicator as to where they may place their investments for the future.

References

- Dal Bianco, M., Camacho, M., & Quiros, G. P. (2012). Short-run forecasting of the euro-dollar exchange rate with economic fundamentals. *Journal of International Money and Finance*, 31(2), 377–396.
- Gwarnicki, J., & Bahmani, S. (2021). Eur/usd speculative significance study: Using an econometric model analyzing macroeconomic data announcements. *International Journal of Humanities Social Sciences and Education*, 8(2), 15-26.
- McCoy, E. (2020). *Euro-us dollar exchange rate dynamics at the effective lower bound* (Tech. Rep.). Directorate General Economic and Financial Affairs (DG ECFIN), European . . .

