Stanford Economic Review

Stanford Undergraduate Economics Journal

Volume 14

Editors and Staff

Note from the Editors

On behalf of the 2025-26 *Stanford Economic Review* Editorial Board, we are pleased to present the fourteenth volume of Stanford University's undergraduate economics journal.

This academic year, we received submissions from across the world, and this publication represents a phenomenal and diverse set of topics. From trade shocks and climate change to the impact of tourism and AI on employment, the breadth of research featured in this volume reflects the wide-ranging economic questions shaping our world today – and the kind of intellectual curiosity that defines the Stanford Economic Review.

As part of our mission to support meaningful inquiry, empirical excellence, and thoughtful economic scholarship, we have also published commentaries that cover topical and imperative issues such as the environmental and economic impacts of the AI boom and recent developments in central banking. We hope our commentaries and this journal serve as a source of inspiration for students and readers to think critically about how systems, incentives, and policy interact.

We would like to thank each author behind the pieces that we publish, without whom this journal could not exist. Their rigor and passion shines through not only in their research and writing but in their communication and diligence throughout the editorial process. We commend all new and returning editors within the Stanford Economic Review, who help ensure that our journal consistently represents the best of our submissions. Lastly, we would like to extend our gratitude to the Stanford Economics Association for its continued support of our mission.

Araha Uday and Kasha Tyranski
2025-26 Editors in Chief

Contents

The Hidden Cost of Climate Change: Examining the Rates of Property Insurance Coverage Among Low-Income Households

Rachel Costello

Tulane University

Abstract— **As destruction from weather events has increased, property insurance premiums have also risen. This article utilizes weather data from the National Oceanic and Atmospheric Administration and household survey data from IPUMS USA to answer the following question: "Is the increase in damaging weather events causing lower-income households to drop their property insurance policies at a disproportionate rate?". While this study finds that lower-income households are more likely to be uninsured, the results indicate that the total number of weather events an area experiences does not have a significant effect on whether a household has property insurance, and the interaction effect of income and weather events appears to have a very small, if any, effect on property insurance outcomes.**

I. Introduction

In recent years, climate change's effects have been prevalent in the American property insurance market. In 2024, there were 27 weather events that caused at least a billion dollars or more in damages, and insurer payouts that year exceeded $100 billion (Smith, 2025; Swiss Re, 2024). The increase in destructive weather events has caused insurance premiums to rise. The national average rate increase in 2024 was 10.4% (Woleben, 2025). As damage from weather events further intensifies and insurance premiums continue to increase, it is important to quantify the effects that these trends are having on homeowners. This paper examines whether the increase in damaging weather events has caused lower-income households to drop their property insurance policies at disproportionate rates.

A. Trends in Property Insurance Markets

First, it is important to understand how, and to what extent, climate change is impacting the property insurance market in order to understand how households might be affected. Overall, the literature supports the idea that insurance premiums are increasing; however, there is debate regarding which households are being exposed to these higher rates.

Keys and Mulder (2024) use mortgage escrow data from 2014-2023 to analyze current patterns in the insurance market. Mortgage escrow data is used since mortgage lenders require borrowers to purchase property insurance, specifically a HO-3 policy, and to make insurance payments through escrow accounts. The authors find that insurance premium rates remained stable from 2014-2017 and from 2019-2021. However, during the time periods of 2017-2019 and 2021 onwards, rates increased. From 2021-2023, Keys and Mulder (2024) find that the average premium increased by over $500 in the span of the three years. While the two find that higher disaster risk areas face higher premiums compared to lower disaster risk areas, the difference between what higher risk areas pay in premiums compared to lower risk areas has changed over time. From 2014-2020, the gap between the premium prices for higher disaster risk areas and lower risk areas remained stable; however, from 2020 onwards, premiums for those in riskier areas started to increase more, widening the gap.

Keys and Mulder (2024) also identify the increase in reinsurance costs as a factor contributing to higher property insurance premiums. Reinsurance is insurance that insurers (the primary insurer) buy from other insurers (the reinsurer) and is a way for the primary insurer to transfer some of their risk to another party (Insurance Information Institute, 2014). In the property insurance market, the primary insurer and reinsurer often enter into a proportional agreement in which, for large loss events (such as a hurricane), the primary insurer pays out claims up to a certain agreed-upon amount; any damages exceeding this amount are paid by the reinsurer (Insurance Information Institute, 2014). Although the premium increases that started in 2020 occurred throughout the United States, Keys and Mulder (2024) find that these increases primarily occurred in states in which primary insurers relied more heavily on reinsurance. In fact, the two estimate that the increase in reinsurance premiums explained over half of the overall insurance premium increases.

Unlike Keys and Mulder, Oh et al. (2025) identify a different pattern occurring with increasing premiums. Oh et al. (2025) examine how different regulatory environments in each state affect insurance pricing, and the authors classify each state as "high friction", "medium friction", or "low friction". High friction states are ones in which the regulatory environments are stricter, and thus insurers can make fewer adjustments to their premium pricing compared to when they operate in medium and low friction states. The authors specifically look at the pricing of insurers who offer policies in multiple states and find that since insurers who operate in high friction states are unable to adjust their premiums after facing increased losses, the companies offset these losses by increasing premiums in low friction states. This spillover effect from high friction states to low friction states is estimated to account for about 30% of the premium growth in the lowest friction states. Thus, while Oh et al. (2025) also document rate increases, they find that the households who see the largest premium increases do not live in the riskiest areas, but rather are residents of the lowest friction states.

B. *Property Insurance Policy Adaptation Behavior*

While the research regarding what is driving the current trends in the property insurance market is mixed, there is more consensus amongst the findings looking at property insurance policy adaptation behavior. The research indicates that people's willingness to pay for insurance is less than their risk and lower-income households are more likely to be underinsured.

Wagner (2022) estimates people's willingness to pay for flood insurance utilizing flood insurance policy and claims data, geographic data on flood zones, and data on housing characteristics. Although flood insurance is separate from property insurance since HO-3 policies do not cover damage due to flooding, the flood insurance market has undergone similar trends and seen an increase in claims in the last 20 years (Wagner, 2022). Wagner finds that on average, in high-risk flood zones, flood insurance premiums are only about two-thirds of a homeowner's expected payout, yet 40% of homeowners remain uninsured. The author attributes this to homeowners underestimating their house's risk of flooding, and she estimates that only about half of homeowners who live in high-risk flood zones are willing to pay a premium amount equal to their expected payout. While the property insurance market is separate from the flood insurance market, Wagner's (2022) finding that people underestimate their house's risk to natural disasters is important to consider when studying property insurance policy adaptation behavior.

Research by Cookson et al. (2025) supports the notion that underestimation behavior also occurs in the property insurance market. The authors use data from the Colorado Division of Insurance pertaining to insurance policyholders who were affected by the Marshall Fire that occurred in 2021. The data includes details regarding insurance contracts, property characteristics, and household characteristics. Cookson et al. (2025) estimate that 74% of the homeowners in their sample were underinsured and that 36% of these homeowners were so severely underinsured that their coverage was equal to less than three-fourths of their house's replacement cost. While after the fire, construction costs became inflated, the authors find that inflation alone could not explain the rates of underinsurance, as most of the inflated costs were covered by a policy's replacement cost provisions. Instead, the authors identify that the most likely explanation for these rates of underinsurance is that when shopping around for policies, homeowners are more likely to select the cheapest offer and to assume that the estimates for the necessary coverage amounts are the same for all policies.

Sastry et al. (2025) also document widespread rates of underinsurance in the U.S. property market. They find that 61% of American households are underinsured, with only 70% of their house's rebuilding costs being covered by their policy. The authors find that amongst higher-income households, those who live in areas that are more exposed to natural disasters purchase more coverage and are less likely to be underinsured. However, for lower-income households, who are more likely to be underinsured in general, they find that exposure to natural disasters does not correspond to increases in coverage. Instead, Sastry et al. (2025) find that in households in which the head of the household has an income that is less than $40,000 a year, the rates of severe underinsurance are similar, regardless of location. Additionally, for households in which the head of the household makes around $20,000 a year, the authors find that households in areas that have a higher exposure to natural disasters tend to have higher rates of underinsurance than those in low-risk areas. The authors also document a negative correlation between insurance premium prices and coverage, as households with the lowest insurance prices have an underinsurance rate of only 40%, while those with the highest prices have underinsurance rates of 80%.

C. *The Literature's Relevance to the Current Study*

Overall, previous research has documented an increase in property insurance claims and premiums in recent years. However, the literature disagrees on whether the increases in insurance premiums have mainly been experienced by those who live in more risk-prone areas or by those who live in low friction states. The two explanations have different implications for the current study. If households that live in areas with more intense weather events also face higher premiums, then it would be expected that lower-income households in these areas would be more likely to drop their property insurance coverage completely (Keys & Mulder, 2024). However, if the findings from Oh et al. (2025) are correct and households in low friction states face the largest increases in insurance premiums, then it would instead be expected that lower-income households in low friction states would be the most likely to drop their property insurance policies.

The literature provides evidence that people tend to underestimate both their house's risk of flooding and how much it will cost to repair their house in the wake of a natural disaster (Wagner, 2022; Cookson et al., 2025). Additionally, previous research also establishes that lower-income households are more likely to be underinsured (Sastry et al., 2025). Although it is important to note that the choice to have no property insurance is different than both the choice to have an inadequate amount of property insurance and the choice to opt out of flood insurance, the research does indicate a mismatch between households' evaluations of how much it will cost to repair their houses in the wake of a natural disaster and their willingness to pay for insurance.

D. *The Current Study*

As of the writing of this paper, there appears to be no research pertaining to households dropping property insurance coverage completely. Indeed, several of the aforementioned studies specifically look at households that still have mortgages, and as stated before, mortgage lenders typically require households to have property insurance (Keys & Mulder, 2024; Sastry et al., 2025). Yet, if the recent increases in insurance premiums (due to more expensive weather events) have widened the gap between insurers' perceived risks and expected payouts and a household's willingness to pay for insurance enough, then this could cause some households to drop their property insurance coverage completely. Additionally, if more households are starting to drop their property insurance, then the findings of Sastry et al. (2025) indicate that lower-income households would be more likely to drop their coverage, as these households are more likely to be underinsured.

Thus, this paper utilizes weather data from the National Oceanic and Atmospheric Administration's (NOAA) Storm Events Database and household survey data from IPUMS USA to study the following question: "Is the increase in damaging weather events causing lower-income households to drop their property insurance policies at a disproportionate rate?". In order to quantify the relationship, logistic, probit, and linear probability model regressions were run.

The paper proceeds as follows: Section 2 describes the data that was used, while Section 3 discusses the methodology. Section 4 presents the results, and Section 5 discusses the results. Section 6 concludes the paper.

II. Data

A. Weather Data

In order to analyze the effect that climate change has had on the property insurance market, weather data was needed. This paper used weather data from the National Oceanic and Atmospheric Administration's (NOAA) Storm Events Database, as it contained information on weather events that occurred throughout the country from 2004-2022, which corresponded to the period of interest. However, it is important to note that since some of the information contained in the dataset was reported to NOAA from external sources, the information for each event may not have always been 100% accurate (NOAA National Centers for Environmental Information, n.d.). Additionally, while the data included a variable for estimated property damage, most weather events were missing values for this. Thus, from the NOAA data, it was difficult to discern how costly a specific weather event was for insurers. The Storm Event Detail files covering weather events from January 2004 through December 2022 were utilized in the analysis, as the weather events were lagged behind a year from the household survey data to account for the fact that insurers need time to adjust their policy prices. For each weather event, the NOAA data provided information regarding what type of weather the event was, when it started and ended, and where it occurred.

B. Property Insurance Data

To answer the research question, information on household insurance coverage rates and household characteristics, including income, was also needed. This paper used data on insurance coverage along with household characteristics from IPUMS USA. Specifically, American Community Survey (ACS) data from 2005-2023 were used since, from 2005 onwards, all necessary variables for the analysis were available. The ACS variables used in this analysis were the following: annual property insurance premium, household income, house value, PUMA, mortgage status, year, and state. It is important to note that the ACS data is not panel data but instead is repeated cross-sectional data.

C. Data Linkage

The NOAA weather data was linked to IPUMS data by using Public Use Microdata Areas (PUMA). Although both datasets included state and county federal processing standards (FIP) codes, in recent years, IPUMS has not been able to identify most counties in its samples. Thus, PUMAS were the smallest geographic areas consistently used in the ACS, and each PUMA has at least 100,000 people in it.

To assign a PUMA to each weather event, crosswalk files from the Missouri Census Data Center's website were used. Since the data spans over three decades, three different PUMA systems had to be used. From 2002-2011, the 2000 PUMAs, which were based on 2000 Census data, were used, while for 2012-2021, the 2010 PUMAs, which were based on 2010 Census data, were utilized. From 2022 onwards, the 2020 PUMAs, which were based on 2020 Census data, were used.

This paper used the Missouri Census Data Center's Geographic Correspondence Engine (Geocorr) to download the necessary crosswalk files. For the 2000 and 2010 PUMAs, Geocorr 2018 was used to download files, each of which contained every PUMA and its corresponding state and county FIP. Additionally, as a county can be broken up into multiple PUMAs, the files contained weights that represented the proportion of a county's housing units that were in each PUMA. It is important to note that Geocor 2018's county boundaries and weights were based on 2010 Census data, thus making it an imperfect tool for obtaining the 2000 PUMA crosswalk file. However, Geocorr 2000, which was based on 2000 Census data, only provided a crosswalk file for the 2000 PUMA for the 5% sample, so Geocorr 2018 was used in its place. The 2020 PUMA crosswalk file was obtained using Geocorr 2020, which based county boundaries and weights on 2020 Census data.

III. Methodology

A. Weather Data

1) Weather Data Cleanup: The NOAA data consisted of weather episodes, which could contain multiple weather events. The following weather events were kept and used for analysis: Blizzard, Extreme Cold/Wind Chill, Hail, High Wind, Hurricane (Typhoon), Strong Wind, Thunderstorm Wind, Tornado, Tropical Storm, Wildfire, and Winter Storm. The other events contained in the dataset, such as Floods and Tropical Depressions, were deleted either because they are not strong enough to cause property damage or because a standard HO-3 homeowner's policy would not have coverage for the weather event (Insurance Information Institute, n.d.). Hail events in which the hail was reported as being less than one inch in diameter along with wind events for which the speeds were less than 40.84 knots were deleted as these events are not typically strong enough to cause property damage (U.S. Department of Commerce, NOAA, n.d.-a;

U.S. Department of Commerce, NOAA, n.d.-b). Similarly, weather events for which the property damage was estimated to be zero dollars were also removed.

2) Weather Data Calculations: In order to be able to count the total number of weather events, dummy variables for each type of weather event were created, where a value of one for dummy variable [X] indicated that the weather event was type [X]. Each weather observation included the county FIP code in which the weather event took place. The PUMA crosswalk files were then used to assign each weather observation its appropriate PUMA code(s). As the weather data would be lagged a year in the analysis, the 2000 Census-based PUMAs were used for weather events from 2004-2010, while the events from 2011-2020 used the 2010 PUMAs, and the observations from 2021-2022 used the 2020 PUMAs.

Some counties were split up into multiple PUMAs. For these counties, duplicate observations were created. For each weather-type dummy variable, a new weighted weather-type dummy variable was created by multiplying the weather-type dummy variable by the housing-unit weight. For example, suppose County A was split so that 30% of its housing units were in PUMA 1, while 70% of its housing units were in PUMA B. Then a hurricane observation that occurred in County A would take on a new weighted value of 0.30 for the observation in which it was assigned PUMA 1, and it would have a weighted value of 0.70 for the observation in which it was assigned PUMA 2. For counties that were located in only one PUMA, the weighted weather-type dummy variable was equal to the original weather-type dummy variable.

After these new weighted dummy variables were created, the total number of weighted weather events for each PUMA for each year was calculated by summing all of its weighted weather-type dummy variables.

B. Household Survey Data

1) Household Survey Data Cleanup: Only observations in which a person indicated that they lived in a household and that they owned the house were kept, as only these people would buy property insurance policies. Additionally, as the dataset included both person-level and household-level variables, only one observation per household was kept to prevent over-counting. A dummy variable, indicating whether a household had property insurance, was constructed. If annual property insurance costs were listed as zero dollars, then the dummy variable took a value of one; otherwise, it took a value of zero. Another dummy variable for mortgage status was created; if a household did not have a mortgage, then they were assigned a value of one, otherwise they were assigned a value of zero. As the data spanned from 2005-2023, the following variables that were used in the analysis were deflated to 1999 price levels using the CPI deflator: annual property insurance cost, house value, and total household income. Additionally, the variable for household income included negative values, which indicated that a household had a total income loss. For the purposes of this analysis, these values were recoded as being zero dollars.

2) Household Survey Calculations: In order to determine if lower-income households are dropping their property insurance coverage at disproportionate rates, both income quintile and decile ranks were assigned to each household. Additionally, the ACS data includes a variable called household weight that allows the sample to be scaled up to represent the entire United States population; household weight was used in the analysis.

C. Merging the Weather and Household Survey Data

Each household was assigned a total amount of weighted weather events by matching the two datasets using the PUMA and year variables.

D. Regressions

1) Primary Regressions: As this paper explores whether climate change is causing lower income households to drop their insurance coverage at disproportionate rates, the dependent variable was whether a household had property insurance, while the independent variables were the following interactions terms: first quintile x total weighted weather events, second quintile x total weighted weather events, third quintile x total weighted weather events, and fourth quintile x total weighted weather events. The dummy variables first quintile, second quintile, third quintile, and fourth quintile, along with total weighted weather events, were also included in the regression. Lastly, as previous economic research has looked at both physical house characteristics along with household characteristics in connection to property insurance, the following variables were also controlled for: mortgage status, inflation-adjusted house value, the decade in which the house was built, state fixed effects, and year fixed effects. Since the dependent variable is a binary variable, a logit, probit, and linear probability model regression (Regression 1) were run of the following form:

$$\begin{aligned} \mathit{NoPropertyInsur}_i = \beta_0 + \sum_{i=1}^{4} \mathit{Quantile}_i \\ + \mathit{TotalWeatherEvents}_i \\ + \sum_{i=1}^{4} (\mathit{Quantile}_i \times \mathit{Weather}_i) \\ + \gamma X_i + \epsilon_i \end{aligned} \tag{1}$$

Where γX_i is a vector of control variables previously mentioned.

These regressions were run using household weight as a probability weight and clustered at the PUMA level.

An additional linear probability model regression, Regression (2), was run. Regression (2) had the same independent variables and all of the same control variables, except for the state fixed effects, as Regression (1). For Regression (2), state fixed effects were replaced with PUMA fixed effects.

$$
\begin{aligned}
\textit{NoPropertyInsur}_i = \beta_0 + \sum_{i=1}^{4} \textit{Quantile}_i \\
+ \textit{TotalWeatherEvents}_i \\
+ \sum_{i=1}^{4} (\textit{Quantile}_i \times \textit{Weather}_i) \\
+ \gamma X_i + \epsilon_i
\end{aligned}
\tag{2}
$$

2) Secondary Regressions:

a) Regressions Using Income Decile: A logit, probit, and linear probability model variation of the primary regression were run in which, instead of using income quintiles as the income variables, income deciles were used. Thus, the regressions took the following form:

$$
\begin{aligned}
\textit{NoPropertyInur}_i = \beta_0 + \sum_{i=1}^{9} \textit{Decile}_i \\
+ \textit{TotalWeatherEvents}_i \\
+ \sum_{i=1}^{9} (\textit{Decile}_i \times \textit{Weather}_i) \\
+ \gamma X_i + \epsilon_i
\end{aligned}
\tag{3}
$$

Where γX_i is a vector of the same control variables from Regression (1).

Another linear probability model regression, Regression (4), was run. Regression (4) had the same independent variables and all of the same control variables, except for the state fixed effects, as Regression (3). For Regression (4), state fixed effects were replaced with PUMA fixed effects.

$$
\begin{aligned}
\textit{NoPropertyInur}_i = \beta_0 + \sum_{i=1}^{9} \textit{Decile}_i \\
+ \textit{TotalWeatherEvents}_i \\
+ \sum_{i=1}^{9} (\textit{Decile}_i \times \textit{Weather}_i) \\
+ \gamma X_i + \epsilon_i
\end{aligned}
\tag{4}
$$

b) Regressions With Property Insurance Premium as the Dependent Variable: In order to see if climate change was causing lower-income households' property insurance premiums to increase at a disproportionate rate, the following linear regressions were run:

$$
\begin{aligned}
\textit{PropertyPremium}_i = \beta_0 + \sum_{i=1}^{4} \textit{Quantile}_i \\
+ \textit{TotalWeatherEvents}_i \\
+ \sum_{i=1}^{4} (\textit{Quantile}_i \times \textit{Weather}_i) \\
+ \gamma X_i + \epsilon_i
\end{aligned}
\tag{5}
$$

$$
\begin{aligned}
\textit{PropertyPremium}_i = \beta_0 + \sum_{i=1}^{9} \textit{Decile}_i \\
+ \textit{TotalWeatherEvents}_i \\
+ \sum_{i=1}^{9} (\textit{Decile}_i \times \textit{Weather}_i) \\
+ \gamma X_i + \epsilon_i
\end{aligned}
\tag{6}
$$

Where, in both of the above equations, γX_i is a vector of the control variables. All of the control variables were the same as those from Regression (1); however, instead of state-fixed effects, PUMA-fixed effects were utilized.

These regressions were run using household weight as a probability weight and clustered at the PUMA level.

IV. Results

A. Descriptive Statistics

Table 1 provides the mean household income, property insurance premium, total amount of weather events, decile, and house value for each year. Note that the household income, property insurance premiums, and house value variables are presented in 1999 price levels. Additionally, the calculations for these means incorporate the household weights. Across the 19 years, the mean household income is $66,334.63, which corresponds to being in the fifth income decile. The mean annual property insurance premium is $692.11. During the time period, the average household experiences 12.3 weather events in the previous year and has a mean house value of $209,317.00. Over time, it appears that the average total number of weather events has fluctuated; however, property insurance premiums have steadily risen from an average of $638.36 in 2005 to a mean of $821.02 in 2023.

TABLE I: Mean Household Statistics

Year	Household Income	Insurance Premium	Weather Events	Decile	House Value
2005	62765.29	638.36	10.48	6.23	206961.30
2006	63638.26	666.08	10.50	5.85	218210.60
2007	65406.50	677.05	11.76	5.86	218204.70
2008	65075.82	628.75	10.82	5.86	215315.80
2009	64490.62	631.73	15.02	5.87	200776.80
2010	62383.34	631.66	11.95	5.90	191556.60
2011	61427.75	636.09	13.26	5.90	180652.20
2012	62055.54	638.17	18.17	5.90	175752.30
2013	63492.55	654.76	13.52	5.90	178089.10
2014	64127.35	672.13	11.24	5.90	183673.80
2015	66806.17	704.93	11.14	5.90	196055.80
2016	67738.77	715.58	11.01	5.88	202637.40
2017	68339.03	717.22	11.86	5.87	208077.60
2018	69152.51	719.77	12.76	5.86	213015.00
2019	71414.26	727.08	10.94	5.87	217509.50
2020	72180.41	755.67	13.72	5.84	224406.90
2021	69697.20	742.16	12.94	5.86	235798.40
2022	68601.81	731.81	10.65	5.84	246009.90
2023	69367.99	821.02	12.13	5.82	247690.40
Total	66334.63	692.11	12.30	5.90	209317.00

Table 2 and Table 3 break down the total number of weather events per year by event type and provide the mean number of each weather event type for each year. From 2005-2023, Thunderstorm Winds was the most common weather event, averaging 7.3 events per year, while Extreme Cold/Winter Chill and Blizzards were the least common, both having a mean of zero events per year.

B. Regressions

1) Results from Regression (1): Table 4 provides the results from Regression (1) after running a logit, probit,

TABLE II: Mean Weather Events Broken Down by Type of Event: Blizzard, Extreme Cold/Winter Chill, Hail, High Wind, and Hurricane

Year	Blizzard	Extreme Cold/ Winter Chill	Hail	High Winds	Hurricane	Strong Wind
2005	0.04	0.00	2.80	0.49	0.04	0.10
2006	0.09	0.00	2.94	0.45	0.13	0.06
2007	0.05	0.00	3.52	0.65	0.00	0.09
2008	0.12	0.00	2.84	0.44	0.01	0.10
2009	0.16	0.00	4.11	0.76	0.02	0.22
2010	0.26	0.00	3.65	0.60	0.00	0.15
2011	0.23	0.00	3.69	0.54	0.00	0.10
2012	0.17	0.00	6.00	0.51	0.00	0.09
2013	0.10	0.00	4.36	0.64	0.01	0.10
2014	0.16	0.00	3.32	0.32	0.00	0.09
2015	0.18	0.00	3.43	0.48	0.00	0.07
2016	0.07	0.00	3.28	0.36	0.00	0.06
2017	0.12	0.00	3.34	0.45	0.01	0.11
2018	0.07	0.00	3.67	0.50	0.00	0.10
2019	0.18	0.00	2.78	0.48	0.00	0.07
2020	0.20	0.00	3.31	0.62	0.00	0.18
2021	0.10	0.00	2.77	0.61	0.00	0.12
2022	0.04	0.00	2.22	0.77	0.00	0.07
2023	0.18	0.00	2.61	0.80	0.00	0.10
Total	0.13	0.00	3.38	0.56	0.01	0.10

TABLE III: Mean Weather Events Broken Down by Type of Event: Thunderstorm Winds, Tornado, Tropical Storm, Wildfire, Winter Storm

Year	Thunderstorm Winds	Tornado	Tropical Storm	Wildfire	Winter Storm
2005	5.87	0.96	0.15	0.03	0.00
2006	5.98	0.65	0.15	0.04	0.00
2007	6.75	0.56	0.04	0.09	0.00
2008	6.66	0.60	0.01	0.05	0.00
2009	8.54	0.99	0.15	0.07	0.00
2010	6.55	0.64	0.03	0.06	0.00
2011	7.89	0.75	0.01	0.06	0.00
2012	10.05	1.04	0.10	0.19	0.00
2013	7.67	0.51	0.04	0.10	0.00
2014	6.78	0.52	0.01	0.03	0.00
2015	6.40	0.52	0.01	0.05	0.00
2016	6.56	0.64	0.01	0.05	0.00
2017	7.18	0.51	0.06	0.07	0.00
2018	7.48	0.80	0.08	0.06	0.00
2019	6.67	0.62	0.07	0.08	0.00
2020	8.50	0.84	0.03	0.03	0.00
2021	8.52	0.60	0.14	0.07	0.00
2022	6.71	0.71	0.05	0.07	0.00
2023	7.68	0.63	0.04	0.08	0.00
Total	7.30	0.69	0.06	0.07	0.00

and linear probability model regression. The coefficients for the logit and probit models are the marginal probabilities, calculated at the averages. Overall, the results from all three regressions are quite similar.

The average probability of not having property insurance was 0.109 or 10.9%. The results indicate that the probability that a household has no property insurance increases as its income quintile decreases. For example, the logit regression finds that households that are in the first quintile are 6.8% more likely not to have property insurance than those in the fourth quintile, while the probit model estimates this difference to be 7.4%. However, there is no statistically significant effect of the total number of weather events that a household experienced the previous year on their probability of having a property insurance policy in the following year. Similarly, while the coefficients on the quintile x weather terms are statistically significant, they are extremely small values, indicating that there is a very small interaction effect of income quintile and total number of weather events experienced on the probability that a household has a property insurance policy.

TABLE IV: Regression 1 Results

Variable	Logit	Probit	Linear Probability
in_firstquint	0.073***	0.083***	0.102***
	(0.001)	(0.001)	(0.002)
in_secondquint	0.039***	0.045***	0.038***
	(0.001)	(0.001)	(0.001)
in_thirdquint	0.019***	0.024***	0.013***
	(0.001)	(0.001)	(0.001)
in_fourthquint	0.005***	0.009***	0.001***
	(0.001)	(0.001)	(0.000)
total_wtd_weather	-0.009	-0.000	-0.000**
	(0.000)	(0.000)	(0.000)
firstqin_weather	0.000***	0.000***	0.001***
	(0.000)	(0.000)	(0.000)
secondqin_weather	0.000***	0.000***	0.001***
	(0.000)	(0.000)	(0.000)
thirdqin_weather	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)
fourthqin_weather	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)
Intercept	-3.174***	-1.893***	0.032***
	(0.036)	(0.018)	(0.004)
Average Probability (No Insurance)	0.109	0.109	0.109
Pseudo R^2/R^2	0.133	0.126	0.089

*** p<0.01, ** p<0.05, * p<0.10

2) Results from Regression (2): The results from Regression (2) are very similar to Regression (1); however, the magnitude and significance of some of the coefficients slightly differ. Figure 1 in the Appendix provides the results for Regression (2).

3) Results from Regression (3): Table 5 provides the logit, probit, and linear probability model results from regression (2), where the coefficients for the logit and probit models are the marginal probabilities for each coefficient, calculated at the averages.

The average probability of not having property insurance was 0.109 or 10.9%. Overall, the results are quite similar to those from Regression (1) and indicate that as income increases, the probability of not having property insurance decreases. The logit model finds that those in the first income decile are 8.1% more likely not to have property insurance than those in the ninth decile, while the probit model estimates this number to be 10.2%. It appears that it is around the seventh income decile at which the marginal effect of income decile on having property insurance changes from being a negative to a positive effect. As was the case with the Regression (1) results, the amount of weather events in a previous year does not appear to have a significant effect on the probability of having property insurance in the following year, and the interaction effect of the number of weather events and income appears to be essentially zero.

4) Results from Regression (4): The results from Regression (4) are similar to Regression (3). Unlike Regression (3), all of the income decile coefficients are positive; however, for the higher deciles (which were in Regression (3)), these coefficients are very small. Additionally, in Regression (4), many of the interaction term coefficients change from 0.001 to 0.000, which is a very small decrease. As was the case with Regression (1) and (2), some of the significance levels of the coefficients have changed. Figure 2 in the Appendix provides the results for Regression (4).

5) Results from Regression 5: Table 6 provides the results for Regression (5). It is important to note that 3,903 obser-

TABLE V: Regression 3 Results

Variable	Logit	Probit	Linear Probability
in_firstdec	0.079***	0.094***	0.122***
	(0.001)	(0.001)	(0.003)
in_seconddec	0.053***	0.065***	0.072***
	(0.001)	(0.001)	(0.001)
in_thirddec	0.036***	0.046***	0.041***
	(0.001)	(0.001)	(0.001)
in_fourthdec	0.024***	0.033***	0.022***
	(0.001)	(0.001)	(0.001)
in_fifthdec	0.014***	0.022***	0.009***
	(0.001)	(0.001)	(0.001)
in_sixthdec	0.006***	0.014***	0.002
	(0.001)	(0.001)	(0.001)
in_seventhdec	-0.001	0.007***	-0.004***
	(0.001)	(0.001)	(0.001)
in_eighthdec	-0.007***	-0.000	-0.008***
	(0.001)	(0.001)	(0.001)
in_ninthdec	-0.012***	-0.008***	-0.011***
	(0.001)	(0.001)	(0.001)
total_wtd_weather	-0.00**	-0.000	-0.000***
	(0.000)	(0.000)	(0.000)
weather_firstdec	0.000***	0.000***	0.001***
	(0.000)	(0.000)	(0.000)
weather_seconddec	0.000***	0.000***	0.001***
	(0.000)	(0.000)	(0.000)
weather_thirddec	0.000***	0.000***	0.001***
	(0.000)	(0.000)	(0.000)
weather_fourthdec	0.000***	0.000***	0.001***
	(0.000)	(0.000)	(0.000)
weather_fifthdec	0.000***	0.000***	0.001***
	(0.000)	(0.000)	(0.000)
weather_sixthdec	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)
weather_seventhdec	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)
weather_eightdec	0.000***	0.000***	0.000***
	(0.000)	(0.000)	(0.000)
weather_ninthdec	0.000	0.000	0.000***
	(0.000)	(0.000)	(0.000)
Intercept	-3.079***	-1.863***	0.039***
	(0.036)	(0.018)	(0.004)
Average Probability (No Insurance)	0.109	0.109	0.109
Pseudo R^2/R^2	0.134	0.127	0.091

*** p¡0.01, ** p¡0.05, * p¡0.10

TABLE VI: Regression (5) Results

Variables	Coefficient
in_firstquint	-145.230***
	(2.191)
in_secondquint	-111.609***
	(1.722)
in_thirdquint	-93.258***
	(1.445)
in_fourthquin	-72.194***
	(1.242)
total_wtd_weather	0.368***
	(0.089)
firstqin_weather	-0.507***
	(0.098)
secondqin_weather	-0.340***
	(0.082)
thirdqin_weather	-0.209***
	(0.069)
fourthqin_weather	-0.083
	(0.059)
Intercept	769.891***
	(1.474)
Average Premium	$691.99
R^2	0.321

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

vations were dropped since they were the only observation from their respective PUMA.

The results indicate that lower-income households pay less in insurance premiums on average. The average annual premium was $691.99. Additionally, the total number of weather events experienced in the previous year has a very small positive effect on insurance premium prices in the following year, as experiencing one additional weather event per year is expected to increase insurance premium prices by $0.37. Although the coefficients for the quintile x weather events interaction terms are statistically significant for all quintiles except the fourth, the interaction effects on premium prices are quite small.

6) Results from Regression (6): Table 7 provides the results for Regression (6). It is important to note that 3,903 observations were dropped since they were the only observation from their respective PUMA.

The results from Regression (6) are very similar to those from Regression (5) as those in lower deciles tend to pay less in insurance premiums. Unlike Regression (5), the total number of weather events a household experienced in the previous year does not have a significant effect on insurance premium prices in the following year. Additionally, the decile x weather event interaction term coefficients are very small and, for the most part, not statistically significant.

V. Discussions

The results indicate that lower-income households are more likely to pay less in property insurance premiums and are less likely to have property insurance in general. However, there appears to be no statistically significant effect of the total number of weather events a household experienced in the previous year on their likelihood of having insurance, and very little to no effect on their property insurance premium prices the following year. Nor does there appear to be a sizable interaction effect of household income and total amount of weather events experienced on the price of insurance or on the probability of having insurance. Overall, the results indicate that while there is an effect of income on the price a household pays for insurance as well as on the probability of having an insurance policy, climate change-related weather events are not an underlying mechanism driving these effects. There are several plausible explanations for these findings.

Previous economic literature seems to be in alignment with the findings regarding the relationship between income and property insurance. Wagner (2022) and Cookson et al. (2025) both observe that people tend to underevaluate their risks in the insurance market and, as a result, are often underinsured. Additionally, Sastry et al. (2025) find that lower-income households have higher rates of underinsurance, as well as observe a negative correlation between insurance premiums prices and coverage. Sastry et al. (2025) attribute the higher rates of underinsurance amongst lower-income households to the fact that these households are more financially constrained and thus struggle more to pay when premium prices increase. Although the choice between being underinsured and uninsured is not the same, insurance policies do have a

TABLE VII: Regression 6 Results

Variable	Coefficient
in_firstdec	-266.028***
	(3.436)
in_seconddec	-241.350***
	(3.241)
in_thirddec	-224.085***
	(3.028)
in_fourthdec	-213.819***
	(2.887)
in_fifthdec	-204.671***
	(2.787)
in_sixthdec	-196.173***
	(2.678)
in_seventhdec	-184.579***
	(2.621)
in_eighthdec	-172.956***
	(2.523)
in_ninthdec	-141.797***
	(2.216)
total_wtd_weather	0.155
	(0.113)
weather_firstdec	-0.356***
	(0.132)
weather_seconddec	-0.260**
	(0.126)
weather_thirddec	-0.175
	(0.116)
weather_fourthdec	-0.086
	(0.117)
weather_fifthdec	-0.038
	(0.105)
weather_sixthdec	0.039
	(0.106)
weather_seventhdec	0.069
	(0.101)
weather_eighthdec	0.187
	(0.096)
weather_ninthdec	0.319***
	(0.071)
Intercept	875.439***
	(2.649)
Average Premium	$691.99
R^2	0.322

*** p<0.01, ** p<0.05, * p<0.10

baseline price to get a minimum amount of coverage. Thus, if this baseline price is too high for a household, especially for a lower-income household, these households might decide to forgo coverage completely. However, one thing to note is that Sastry et al. (2025) find that for higher-income households, living in a higher disaster risk area was associated with purchasing more insurance coverage. Yet the results from Regression (5) indicate a small negative effect of the income x weather event interaction terms for higher income quintiles, while the results from Regression (6) fail to find a statistically significant positive effect for these interaction terms, which is contrary to Sastry et al.'s finding.

Although the results find that weather events have no effect on the probability of having property insurance and have little to no effect on insurance premium prices, which is contrary to Keys and Mulder's (2024) findings and the hypotheses of this study, there are several possible explanations for this.

First, most of the previous research on the topic has looked at disaster risk, rather than the total number of weather events, in relation to property insurance. While this study utilizes NOAA weather data with the assumption that the total number of weather events experienced in an area would be highly related to disaster risk, it is not a perfect substitute measure for disaster risk. Additionally, while the weather events that were retained for the analysis were ones that are typically costlier to insurers, the NOAA data makes it hard to decipher how destructive a particular weather event is. As it is likely that the total amount of property damage from weather events in an area, rather than the total amount of weather events itself, is a better predictor of premium increases, this could be one potential reason why the results failed to find a meaningful correlation between total amount of weather events and insurance policies. Indeed, the most common weather event type experienced was Thunderstorm Winds, which, on average, would be expected to cause significantly less damage than Hurricanes, which were a lot less common, leading credence to this idea that a high concentration of weather events does not necessarily mean a high risk of natural disaster.

Second, this study uses an imperfect method to assign each PUMA a total amount of weather events for a given year. As weather is independent of housing densities, the weighting method could have both overcounted the amount of weather events in some PUMAs while undercounted the amount of weather events in others.

Third, it is possible that the lack of connection between the total amount of weather events and insurance outcomes can be explained by Oh et al.'s (2025) findings that higher insurance premiums are associated with living in a low friction state, rather than living in a high-disaster-risk area. As Oh et al.'s methodology for classifying high versus low friction states is quite labor-intensive, it was outside of the scope of this study to conduct these classifications. However, in attempts to see how plausible Oh et al.'s findings might be in explaining this study's results, supplementary regressions were run. For each state, the total amount of weather events in the previous year was regressed on the change in the state's average property insurance premium in the following year. Figure 1 in the appendix shows the coefficients and 95% confidence interval for each state. Overall, it appears for these results that across all states, total weather events did not have a large impact on premium price changes. Although this method is a lot less precise than Oh et al.'s, it seems that the high friction versus low friction state theory cannot completely explain this study's findings, as the coefficients for each state were quite similar. If the high friction versus low friction explanation were to hold, it would have been expected that low friction states would have had noticeably higher, positive coefficients than high friction states. Given the overall lack of a significant effect of total weather events on property insurance outcomes, it is not surprising that the income x weather interaction terms do not appear to have a notable effect on the dependent variables.

This study adds to the economics literature by finding that lower-income households are more likely not to have property insurance coverage and that the amount of weather events a household experiences in the previous year does not appear to be predictive of property insurance outcomes. However, as mentioned previously, there are several limitations of this study, and thus, more research that uses different methodologies for estimating climate change-related impacts should be conducted in order to draw stronger conclusions regarding the relationship between climate change and property insurance policies.

VI. Conclusion

This paper studies whether the increase in damaging weather events due to climate change has caused lower-income households to drop their property insurance policies at disproportionate rates. Although this study finds that lower-income households are more likely to be uninsured, the results indicate that the total amount of weather events an area experiences does not have a significant effect on whether a household has property insurance, and the interaction effect of income and weather events appears to have a very small, if any, effect on property insurance outcomes.

References

[1] Background on: Reinsurance. (2014, January 12). *Insurance Information Institute*. https://www.iii.org/publications/insurance-handbook/regulatory-and-financial-environment/background-on-reinsurance.

[2] Cookson et al. (2024). *Coverage Neglect in Homeowners Insurance*. SSRN. https://doi.org/10.2139/ssrn.5057551.

[3] Hurricanes, Severe Thunderstorms and Floods Drive Insured Losses above USD 100 Billion for 5th Consecutive Year, Says Swiss Re Institute — Swiss Re. (2024, December 5). *Swiss Re*. https://www.swissre.com/press-release/Hurricanes-severe-thunderstorms-and-floods-drive-insured-losses-above-USD-100-billion-for-5th-consecutive-year-says-Swiss-Re-Institute/f8424512-e46b-4db7-a1b1-ad6034306352.

[4] Keys, B., & Mulder, P. (2024). *Property Insurance and Disaster Risk: New Evidence From Mortgage Escrow Data (Working Paper No. 32579)*. National Bureau of Economic Research. https://doi.org/10.3386/w32579.

[5] Oh et al. (2025). *Pricing of Climate Risk Insurance: Regulatory Frictions and Cross-subsidies*. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.3762235.

[6] Sastry, P. et al. (2024). *Climate Risk and the U.S. Insurance Gap: Measurement, Drivers and Implications**. SSRN. https://doi.org/10.2139/ssrn.4909444.

[7] Smith, A. B. (2020). *U.S. Billion-dollar Weather and Climate Disasters, 1980–present (NCEI Accession 0209268)*. NOAA National Centers for Environmental Information. https://doi.org/10.25921/STKW-7W73.

[8] Storm Events Database - FAQ — National Centers for Environmental Information. (n.d.). *NOAA National Centers for Environmental Information*. Retrieved April 5, 2025, from https://www.ncdc.noaa.gov/stormevents/faq.jsp.

[9] 2024: An active year of U.S. Billion-dollar Weather and Climate Disasters — NOAA Climate.gov. (2025, January 10). *NOAA Climate.gov*. https://www.climate.gov/news-features/blogs/beyond-data/2024-active-year-us-billion-dollar-weather-and-climate-disasters.

[10] U.S. Department of Commerce, NOAA. (n.d.-a). *Estimating Wind*. NOAA's National Weather Service. Retrieved April 13, 2025, from https://www.weather.gov/pqr/wind.

[11] U.S. Department of Commerce, NOAA. (n.d.-b). *Hail Threat Defined*. NOAA's National Weather Service. Retrieved April 5, 2025, from https://www.weather.gov/mlb/hail_threat.

[12] U.S. Homeowners Rates Rise by Double Digits for 2nd Straight Year in 2024. (2025, January 21). *S&P Global Market Intelligence*. https://www.spglobal.com/market-intelligence/en/news-insights/articles/2025/1/us-homeowners-rates-rise-by-double-digits-for-2nd-straight-year-in-2024-87061085.

[13] Wagner, K. R. H. (2022). Adaptation and Adverse Selection in Markets for Natural Disaster Insurance. *American Economic Journal: Economic Policy*, 14(3), 380–421. https://doi.org/10.1257/pol.20200378.

[14] Which Disasters are Covered by Homeowners Insurance? — Insurance Information Institute. (n.d.). *Insurance Information Institute*. Retrieved April 5, 2025, from https://www.iii.org/article/which-disasters-are-covered-by-homeowners-insurance.

Appendix

Variables	Coefficient
in_firstquint	0.049***
	(0.001)
in_secondquint	0.013***
	(0.000)
in_thirdquint	0.006***
	(0.000)
in_fourthquin	0.004***
	(0.000)
total_wtd_weather	-0.000
	(0.000)
firstqin_weather	0.000***
	(0.000)
secondqin_weather	0.000***
	(0.000)
thirdqin_weather	0.000
	(0.000)
fourthqin_weather	-0.000*
	(0.000)
Intercept	0.097***
	(0.000)
Average Probability of Having No Insurance	0.109
R^2	0.199

*** p<0.01, ** p<0.05, * p<0.10

Fig. 1: Regression 2 Results

Variables	Coefficient
in_firstdec	0.073***
	(0.002)
in_seconddec	0.037***
	(0.001)
in_thirddec	0.020***
	(0.001)
in_fourthdec	0.013***
	(0.001)
in_fifthdec	0.009***
	(0.001)
in_sixthdec	0.008***
	(0.000)
in_seventhdec	0.007***
	(0.001)
in_eighthdec	0.007***
	(0.000)
in_ninthdec	0.004***
	(0.000)
total_wtd_weather	0.000*
	(0.000)
weather_firstdec	0.000***
	(0.000)
weather_seconddec	0.000***
	(0.000)
weather_thirddec	0.000**
	(0.000)
weather_fourthdec	0.000
	(0.000)
weather_fifthdec	-0.000*
	(0.000)
weather_sixthdec	-0.000*
	(0.000)
weather_seventhdec	-0.000***
	(0.000)
weather_eighthdec	-0.000***
	(0.000)
weather_ninthdec	-0.000***
	(0.000)
Intercept	0.094***
	(0.000)
Average Probability of Having No Insurance	0.109
R^2	0.200

*** p<0.01, ** p<0.05, * p<0.10

Fig. 2: Regression 4 Results

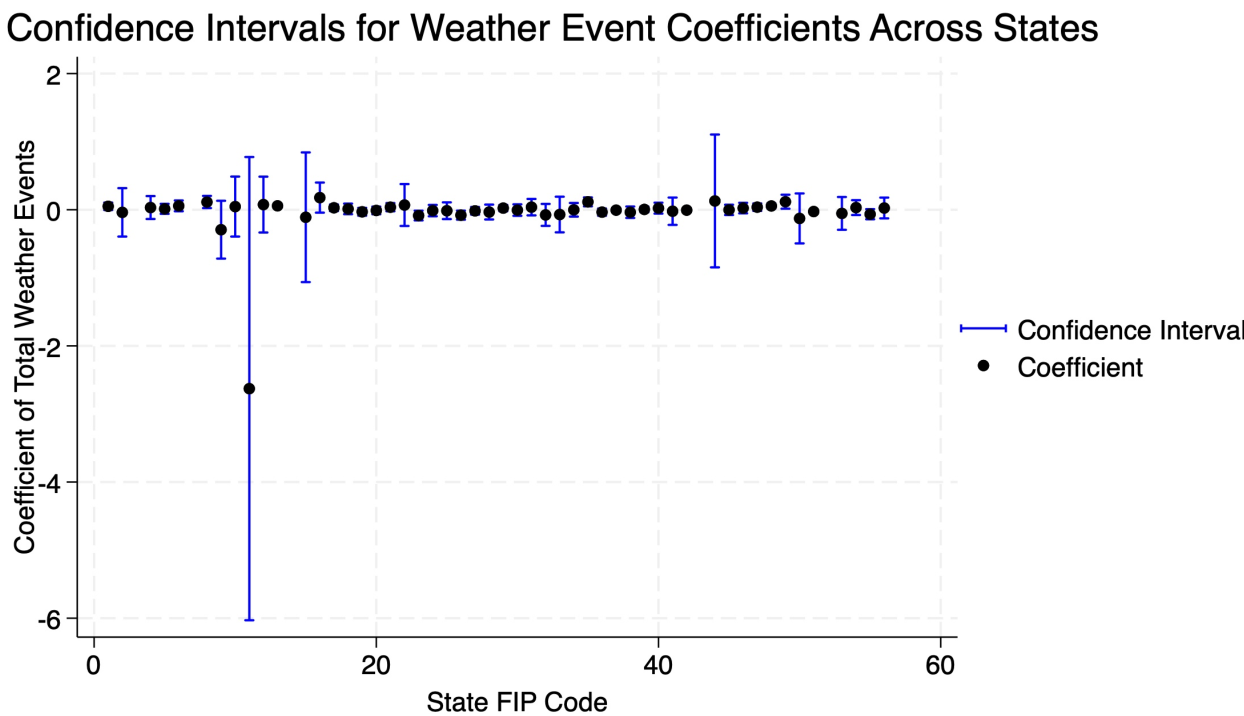

Fig. 3: Coefficients for Each State From Regressing Total Number of Weather Events in a State on Change in Premiums

The Impact of Artificial Intelligence on the Labour Market in Guangzhou

Rou Yu (Roanne) Yap

University of Oxford

Abstract— **This paper estimates the impact of Artificial Intelligence (AI) on the labour market in Guangzhou. This is done by studying the effects of past disruptive technologies on the city's workforce and incorporating theories and evidence from the existing literature. A linear regression model was developed to analyse the relationship of R&D expenditure and Routine Task Index (RTI) of the labour market in Guangzhou from 1995-2022. The model is run with different time lags (0, 1, 2, 3, 5, 7, and 10 years) and includes three control variables: real GDP per capita, the interest rate, and the total population of the city's workforce. According to the model calculations, there is a significant negative correlation between R&D expenditure and RTI of the labour market during the first three years. The regression results, along with evidence from academic literature, were used to assess the plausibility of three theories regarding the potential impact of AI on the labour market. Based on the evaluation of the theories, we tentatively predict that AI will not cause very damaging effects on the workforce in Guangzhou due to the creation of new jobs elsewhere in the economy and the flexible nature of the labour market. These findings can help assess the challenge of the rise of AI in Guangzhou from a policy perspective.**

I. Introduction

A. Background

Guangzhou is the capital and largest city of Guangdong province in southern China. It is located at the heart of the Guangdong-Hong Kong-Macao Greater Bay Area (GBA), the most populous metropolitan area in the world. This crucial location makes Guangzhou a key economic hub of China, with direct routes to Hong Kong, Shenzhen, Macao and Taiwan. The consular district encompasses a combined population of more than 220 million people and is responsible for nearly 20% of China's GDP (Burt, 2024). According to a study by Oxford Economics, Guangzhou is projected to be among the world's top 10 cities by gross domestic product (GDP) in 2035 (Ghosh, 2019). Known as the "Factory of the World," the city is an advanced manufacturing powerhouse with three pillar industries, namely the automobiles, electronics and petrochemical industry (GBA, 2023; Burt, 2024). Guangzhou is also a hub for trade and commerce. As a Peart River port, it is the host of the annual Canton Fair, the oldest and largest trade fair in China (Jin & Weber, 2008).

In the past decade, Guangzhou has intensified its efforts to establish itself as the centre for technology innovation in China (O'Neill, 2018). Zhang Xiaobo, director of the industry and information technology commission of Guangzhou, said: "Guangzhou will facilitate the development of the new generation of information technology, artificial intelligence (AI), and biological medicine" (O'Neill, 2018). However, the current wave of technological change based on advancements in AI has created widespread fear of job loss and further rises in inequality (Ernst et al., 2019). This development has spurred researchers to investigate what the future of labour markets and economics will look like with the rise in AI. Many analysts warn that advances in AI over the next few decades could lead to significant job losses or job polarisation, hence widening income and wealth disparities (Korinek & Stiglitz, 2017; Méda, 2016). Other researchers, however, reach much less dramatic conclusions (Arntz et al., 2016, 2017). Moreover, most existing studies focus mainly on the future impacts of technology and AI on developed countries in the West, while the impacts on developing countries, including China, remain under-studied. This makes it unclear to what extent AI will disrupt the Chinese workforce and how to assess the challenge of automation from a policy perspective. Given that advanced manufacturing is a major industry in Guangzhou, AI has the potential to disrupt the workforce and, consequently, the growth of this key economic region in China. Thus, we believe that it is vital to question how the rise in AI will affect the future labour market in Guangzhou.

Researchers have already been theorising about the future impact of AI on labour markets. This paper will focus on three main theories. The first theory is that AI will complement workers' jobs, thus creating a productivity J-curve (Brynjolfsson et al, 2021). This is based on the observation of past eras of technological change, where there were times in which productivity lulled, then followed by a period of acceleration (Brynjolfsson et al, 2021). For example, the steam technologies of the US industrial revolution took nearly half a century to show visible productivity effects (Brynjolfsson et al, 2021). There was also a productivity slump in the first 25 years following the invention of the electric motor and combustion engine before the pace of productivity exploded in 1915 (Brynjolfsson et al, 2021). This is mainly because adoption of new technologies requires fundamental changes in business processes and workflows, co-invention of new products and business models, and investment in human capital (Dickson, 2022). Workers may need training and reskilling to employ the new technologies

I would like to thank Professor Jake Organ for his continued advice and support throughout the project. I am also grateful to Immerse Education for providing the scholarship and facilitating the Online Research Programme.

(Dickson, 2022). These procedures often take years and millions of dollars without contributing to the company's output in the short term (Dickson, 2022). After a period of decreasing productivity, the changes will result in a sudden rebound, as shown in Figure 1. The intangible investments start to pay off, generating output that can be consumed and measured, increasing workers' productivity (Brynjolfsson et al, 2021). Most jobs and industries are only partially exposed to automation and thus are more likely to be complemented rather than substituted by AI (Stropoli, 2023).

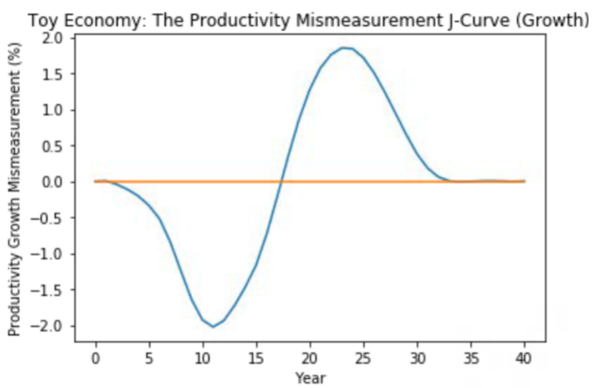

Fig. 1: The J-curve (Brynjolfsson et al, 2021)

Many sources argue that, like disruptive technologies in the past, AI will displace some existing jobs but create many new ones, which is the second theory. A report by Nexford University states that jobs in customer service, accounting, sales, research and analysis, and retail are more likely to be automated in the future (Rostrom, 2024). As discussed in the report, AI may pose a bigger threat to higher-skilled workers than to physical labourers. The in-depth development of digital technologies, such as deep learning and big data analysis, will possibly replace some complex, cognitive and creative jobs that are currently considered irreplaceable (Shen & Zhang, 2024). In addition, men are more likely to be affected in the long run as autonomous vehicles and other machines replace many manual tasks where their share of employment is higher (Hawksworth et al., 2018). However, during the first and second waves of AI, women could be at greater risk due to their higher representation in clerical and other administrative functions (Hawksworth et al., 2018). It is equally important to note that despite displacing many existing jobs, AI will also create new jobs. The World Economic Forum (WEF) states that previous predictions of fewer jobs in the future have generally proven to be false because the total number of people employed has risen worldwide, as shown in Figure 2 (Ekelund, 2024). This implies that previous disruptive technologies have created more jobs than it has destroyed. When the task evolves from "cooperation for all" to "cooperation between man and machine," it results in fewer production constraints and maximises total factor productivity, thus creating more jobs and generating novel collaborative tasks (Balsmeier & Woerter, 2019; Duan et al., 2023). Materialised AI technology can improve the total factor production efficiency and market efficiency, driving the expansion of production scale of enterprises (Liu et al, 2022). This synergy will promote the synchronous growth of labour demand involving various skills, resulting in a creative effect (Liu et al., 2022). Moreover, the WEF predicts for AI to create jobs in three key areas: trainers, explainers and sustainers (Shine, 2023). Trainers develop AI, including programmers, electrical engineers and systems administrators (Shine, 2023). Explainers design interfaces that enable people to interact with AI, such as personalised AI assistants and tutors (Shine, 2023). The three main types of sustainers, who ensure AI systems are being used in the best way possible, are content creators, data curators and ethics and governance specialists (Shine, 2023). Therefore, the net impact of AI on jobs will be neutral or slightly positive rather than very negative as a result of new jobs being created elsewhere in the economy (Gries & Naudé, 2018).

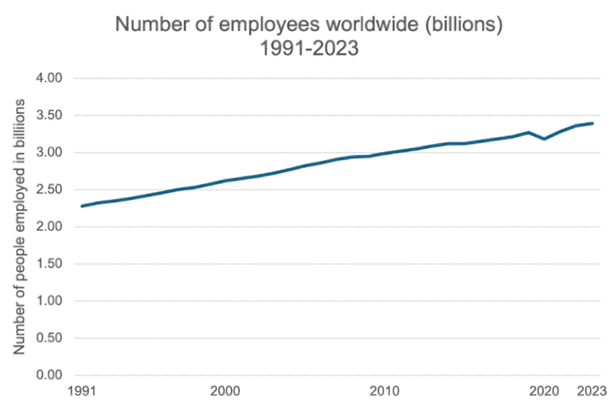

Fig. 2: Number of Employees Worldwide (billions) (Ekelund, 2024)

The third and final theory discussed in this paper regarding the future effects of AI on the labour market is that the extent of the impact depends on the level of freedom and skilled workers in the labour market, according to the WEF (Ekelund, 2024). A more flexible and free labour market should allow more rapid moves between sectors and businesses as the nature of jobs change. Skilled workers will be able to adapt faster to a change in their tasks. This means that nations that are less attractive to entrepreneurs and investors because of lower skilled labour and heavily regulated labour markets may suffer from more unemployment (Ekelund, 2024). Overall, economies that have fewer restrictions in the labour market and more accessible information on job opportunities and training programs are more likely to have higher rates of employment.

This study makes a contribution by providing insight on the plausibility of these three theories to tentatively predict how AI will impact the Guangzhou labour market in the future. First, we will explore the effects of technological change on the Guangzhou labour market in the past by creating a linear regression model. This will provide a deeper understanding of how the Guangzhou labour market responds to disruptive technologies, which will then be used to assess the three theories and make a final prediction.

The next section reviews the literature on three areas: existing predictions about the future impacts of AI on the labour market (not specific to Guangzhou), concepts related to technological change in the labour market, and the background context of technology and AI in Guangzhou. Section 3 introduces the data sources and construction of variables used in the study, while Section 4 and 5 presents the results and its analysis. Lastly, this paper ends with the discussion and conclusion.

II. Literature Review

A. Relevant Literature

There are studies that have predicted the general impacts of AI on the labour markets as well as the impacts in certain countries, many of which have results similar to the three theories that will be evaluated in this paper. A study that modelled the impacts of AI on the labour market in Czechia found that automation will change the nature of work and will cause organisational changes in companies (Fatun & Pazour, 2021). Tasks in production and services will be outsourced, and companies will prefer a direct and flexible networking instead of a hierarchical structure (Fatun & Pazour, 2021). This conclusion relates to the first theory of a productivity J-curve, because the initial changes in business workflows will cause a productivity slump in the short run. Many researchers' results support the second theory, that the net impact of AI on jobs is likely to be neutral or positive. Empirical evidence suggests that the jobs at risk from automation will be balanced by new jobs being created elsewhere in the economy (Dauth et al., 2017; Berriman & Hawksworth, 2017). It is likely that AI creates new jobs by increasing consumer demand, which may be the result of AI creating new or cheaper products, boosting labour demand indirectly (Gries & Naudé, 2018). In addition, Ernst et al. (2019) discovered that the impact of a large-scale introduction of AI will depend on three factors: the price elasticity of supply of capital vs. the elasticity of labour, the substitution elasticity between capital and labour, and the direction of technical change induced by AI, that is, whether AI is capital or labour augmenting (Ernst et al., 2019). The more inelastic the supply of AI, the higher the substitution elasticity between AI and jobs (Ernst et al., 2019). The more labour-saving AI-based innovations are, the higher will be the extent of technological unemployment and the lower will be any wage gains (Ernst et al., 2019). This prediction builds on the final theory mentioned in the previous section, which states that the extent of the impact of AI depends on certain features of the labour market.

This paper aims to either support the predictions mentioned above or present a new possibility for the future of Guangzhou's workforce. In order to predict the effects of AI on the labour market, we need to understand how other disruptive technologies have impacted it in the past. Based on a review of existing studies, there are two basic approaches that explain how technology has changed labour demand. The first is skill biased technological change (SBTC), initially proposed by Griliches in 1969 and Welch in 1970, where technology increases demand for highly skilled and educated workers by increasing workers' productivity (Griliches, 1969; Welch, 1970; Violante, 2016). Low-skilled workers would experience lower demand, resulting in wage inequality. However, in the 2000s, there has been a change in trends. Many economies experienced job polarisation, where demand for high and low-skilled labour increased while demand for middle-skilled labour decreased. Thus, it has been necessary to follow a different approach to SBTC. Routine biased technological change (RBTC) is the second approach, proposed by Autor, Levy and Murnan in 2003 (Autor et al., 2003). The basis of this theory is that technology is biased towards replacing routine tasks that can be described with programmed rules, and complementing non-routine tasks that require analytic and interactive skills that cannot be described with programmed rules (Mihaylov & Tijdens, 2019). In recent decades, RBTC has succeeded in explaining the job polarisation in various countries, including the US, with a task-based approach (Dağli & Kösekahyaoğlu, 2021). A crucial element in the RBTC approach is the measurement of routine and non-routine tasks in occupations, also known as the routine task intensity (RTI) (Mihaylov & Tijdens, 2019). This measures the repetitive content of professions using data from occupational surveys. According to the RBTC approach, new technologies will replace routine tasks, and thus occupations that have a high-RTI value.

A key factor that impacts the speed of SBTC and RBTC - how quickly new technology affects an economy and its labour market - is the rate of technology transfer. Technology transfer is the process by which new inventions and innovations created in research institutions are turned into products and commercialised (TTC, n.d). This is usually done in two ways: through patent licensing and the creation of start-up companies (TTC, n.d). There has been evidence for strong state-led technology transfer in Guangzhou in the past. The city is a major Asia-Pacific research and development hub with a high level of scientific research output, ranking 8th globally and 4th in the Asia-Pacific in 2023 (Nature, 2023). Guangzhou was already home to most of the research centres in its province in 2016, including 79 universities, 141 research organisations and 19 national key laboratories (Fei & Qing, 2016). This indicates that technological innovations are developed rapidly, with only the commercialization and application of these new technologies remaining to complete the process of technology transfer. The local and national government have introduced multiple policies to achieve this. For example, the Guangzhou Nansha Plan proposes to promote joint organisation and implementation of a batch of scientific and technological innovation projects by Guangdong-Hong Kong-Macao research institutions (China News Guangdong, 2023). This plan will actively undertake the transfer and transformation of innovative achievements in fields such as electronic engineering, computer science, marine science, artificial intelligence, and smart cities, establishing a highland for the transfer and transformation of scientific and technological achievements in South China (China News Guangdong, 2023). Moreover, the 12th Five-Year Plan (2011-2015) made it clear that advancing high-technology industry was a national priority, naming 7 strategic industries that will be “nurtured” with fiscal and tax policies to “improve industry core competitiveness” and increase the number of tech start-ups (Nowak, 2012). The AI industry is no exception to receiving significant state support, with many national and local policies to promote its development, as shown in Table 1.

TABLE I: Several Support Policies for AI industry (Huang & Fan, 2020)

Released Time	Policy
Aug, 2019	Guidelines for the Construction of the National New Generation Artificial Intelligence Innovation Development Experimental Zone
Mar, 2019	Guidance on Promoting the Deep Integration of Artificial Intelligence and the Real Economy
Dec, 2017	Three-year Action Plan to Promote the Development of the New Generation of Artificial Intelligence Industry
Jul, 2017	New Generation Artificial Intelligence Development Plan
Nov, 2016	”13th Five-Year” National Strategic Emerging Industry Development Plan
Apr, 2016	Robot Industry Development Plan (2016-2020)
May, 2015	Made in China 2025

Furthermore, the Chinese and local Guangzhou government has shown strong determination to develop technology and AI in the future. According to a development plan for new-generation AI, China aims to become the world’s major AI innovation centre by 2030, with the scale of its AI core industry exceeding 1 trillion yuan and the scale of related industries exceeding 10 trillion yuan (Xinhua, 2024). China has also issued a New Generation Artificial Intelligence Development Plan, which proposes a strategic goal and plan for AI development in 2030 (Huang & Fan, 2020). Moreover, the Guangdong provincial government plans to establish a national hub for the AI industry by 2025 and achieve the largest scale of intelligent computing power (Qiu, 2023). They plan for the number of AI enterprises in Guangdong to exceed 2,000 by 2025, with the core AI industry’s value surpassing 300 billion yuan (Qiu, 2023). The National University of Singapore Guangzhou Research Translation and Innovation Institute (NUS GRTII) launched in 2024 will provide business incubation for technology and AI enterprises (NUS, 2024). One startup that will be incubated is Yimiji Technology, an AI-based medical image processing analysis platform and smart surgical robotic technology, enabling diagnosing and treating medical conditions during perioperative period (NUS, 2024). Therefore, state-led technology transfer is likely to remain strong in the future.

The rate of technology transfer led by the Guangzhou private market is also strong. In 2021, industrial enterprises and businesses in the service sector above the designated size increased their investment in R&D by over 20 percent and 10 percent, respectively (Dezshira, n.d). In the same year, Guangzhou attracted investment from 330 Fortune Global 500 companies, with an average annual growth of 10 percent, and an actual use of FDI amounting to RMB 224 billion (Dezshira, n.d). A reason why private enterprises are so keen on investing in R&D is because, as research has shown, there is a stronger direct link between private R&D expenditure and firm productivity compared to government R&D expenditure (Hu, 2001). Government R&D only contributes indirectly to productivity by promoting private R&D (Hu, 2001). Hence, providing incentives for enterprises to invest in R&D may be a better alternative than providing R&D grants directly (Hu, 2001). A private market example of technology transfer is the Guangzhou Metaverse Innovation Association (GMIA), established in 2022. The metaverse is a type of new technology that involves an immersive virtual world where digital representations of people can interact with each other (Cao, 2022). Made by five local tech companies, including GrowthEase and 37Games, the GMIA now has 23 Guangzhou-based tech companies and research institutions, as well as one of the world’s top AI system developers iFlyTek (Cao, 2022). It is a good example of successful collaboration between private enterprises and local governments in emerging technologies. The concept of the metaverse has become increasingly popular with Chinese companies and investors, as evidenced by the number of metaverse-related trademark applications reaching 16,000 in the first two months of 2022 (Cao, 2022). Another strong example of private-led technology transfer is the animal vaccine-producer Guangdong Winsun Bio-pharmaceutical Corporation (GWBC). Originally a state-owned enterprise established in 1958, the GWBC was restructured into a private holding company with no in-house R&D capabilities in 2002 (Fei & Qing, 2016). After restructuring, it established connections with the China Institute of Veterinary Drug Control, the South China Agricultural University, the Sun Yat-sen University, the Guangdong Ocean University and the East China University of Science and Technology to develop and test new vaccines (Fei & Qing, 2016). Now it invests about 8% of its sales value in R&D every year (Fei & Qing, 2016). Combined with the government’s efforts through policies and subsidies, the eagerness of private enterprises to invest in R&D and collaborate with research institutes to develop new technologies for production provides strong evidence for a rapid rate of technology transfer in Guangzhou. Therefore, the expenditure in R&D should translate into the development and application of new technologies in a short period of time.

However, the nature of technologies has changed over time. A key difference with recent technology such as generative AI is that they diffuse much faster than previous disruptive technology (Stropoli, 2023). A report from Swiss investment bank UBS found that within two months of its initial release, ChatGPT had 100 million monthly active users globally. By comparison, TikTok took nine months and Instagram about 2.5 years to reach the same amount of users. This suggests that, in the future, the utilisation of AI in the market will be faster in contrast with the adoption of disruptive technology in the past.

Overall, based on the approaches mentioned in existing literature and the compelling evidence for rapid technology transfer and substantial R&D expenditure, our hypothesis is as follows: there will be a negative relationship between R&D expenditure and the RTI of the Guangzhou labour market. We expect R&D expenditure to be translated into innovation of new technology, replacing jobs with high-RTI values and decreasing the overall RTI value of the workforce. Testing this hypothesis will provide a better understanding of how disruptive technology affects the Guangzhou labour market. Along with evidence from research papers, this data analysis will help the evaluation of the three theories about the potential impacts of AI on Guangzhou's workforce.

III. Data

A. Data Source

The main data source in the model is the collection of Guangzhou Statistical Yearbooks from 1984-2022, published annually by the Guangzhou Statistics Bureau. The database provided information to build the variables in the model, including: the average number of employed staff and workers in urban non-private units by industrial sectors, total number of employed staff and workers in urban non-private units, average wages of employed staff and workers in urban non-private units, annual CPI (1978 as the base year), expenditure in research and development (R&D), lending interest rate, and GDP per capita. There are 19 industrial sectors that workers were classified into:

1) Agriculture, forestry, animal husbandry and fishery
2) Mining
3) Manufacturing
4) Production and Supply of Electricity, Heat, Gas and Water
5) Construction
6) Transport, Storage and Post
7) Wholesale and Retail Trade
8) Hotels and Catering Services
9) Information Transmission, Software and Information Technology
10) Financial Intermediation
11) Service to Households, Repair and Other Services
12) Real Estate
13) Leasing and Business Services
14) Management of Water Conservancy, Environment and Public Facilities
15) Scientific Research and Technical Services
16) Education
17) Health and Social Work
18) Culture, Sports and Entertainment
19) Public Management, Social Security and Social Organizations

B. Variable Construction

1) Independent Variable: R&D Expenditure: To predict the effect of AI on the future labour market in Guangzhou, the model explores the effect of past technological change on the Guangzhou labour market. Since there is no standard measurement for technological change, adjusted expenditure in R&D was set as a proxy for the independent variable. It is a suitable proxy because R&D expenditure will increase innovation and development of new technology, driving technological change. Moreover, as explained in the literature review, technological transfer and innovation has been strongly encouraged by the Chinese government and private enterprises in Guangzhou through funding (and deregulation). This maximises the proportion of R&D expenditure that is translated into new technology.

2) Dependent Variable: RTI of the Labour Market: Using the RBTC approach, it is expected for new technology to replace high-RTI skills in occupations. This changes the nature or skill content of jobs to lower-RTI skills that are not replaced by technology. Thus, the ideal dependent variable would be the change in RTI of each industrial sector over time. Generally, researchers use different occupational and survey-based data sources to measure RTI. Mihaylov and Tijdens (2019) used the ILO's International Standard Classification of Occupations 2008 (ISCO-08), a four-level hierarchical system for classifying jobs worldwide into 436 unit groups, 130 minor groups, 43 sub-major groups and 10 major groups (Mihaylov & Tijdens, 2019) (ILO, 2008). Marcolin et al. (2016) used the OECD's Program for the International Assessment of Adult Competencies (PIAAC) (Marcolin et al., 2016) (OECD 2014, 2017, 2019), high and middle income countries excluding China. Lewandowski et al. (2020) combined three different data sources: PIAAC, the World Bank's Skills toward Employment and Productivity (STEP) surveys (World Bank 2016a, 2016b, 2018), conducted in middle- and low-income countries excluding China, and the China Urban Labor Survey (CULS) collected by the Institute of Population and Labor Economics of Chinese Academy of Social Science (2017) which included a module based on STEP (Lewandowski et al., 2020).

Out of the different data sources, this model uses the RTI of each occupation measured by Mipahylov and Tijdens (2019) using the ISCO-08 system (Mihaylov & Tijdens, 2019). This is because it is the basis for the many national occupational classifications (WHO, 2019). In addition, China is a member state of the International Labour Organization (ILO) that published ISCO-08, whereas the other data sources (except CULS) do not include data from China. Furthermore, Marcolin et al. (2016) had a routine index that is negatively correlated with the skill-level of occupations, suggesting that it not only measures the routine content of occupations, but also partially captures the skill-level of occupations (Marcolin et al., 2016: Mihaylov & Tijdens, 2019).

However, it is extremely difficult to calculate the change in RTI of each industrial sector over time because all of the data sources mentioned above only measure the fixed RTI of an occupation and not measure the change in RTI over time. Therefore, the change in RTI of the whole labour market will be used as a proxy. This can be done by combining the fixed RTI of each sector with the change in sectoral composition of the Guangzhou labour market over time. The calculation of the fixed RTI of each sector requires multiple steps. Firstly, the 427 ISCO-08 occupations are classified into one or more (usually one) of the 19 industrial sectors by analysing their occupation descriptions from the ILO and the Open Risk Manual. Secondly, the RTIs of each occupation measured by Mipahylov and Tijdens (2019) are added. Originally, the RTI value ranged from -1 to 1. To aid the regression analysis, the range was changed to 0 to 2 by adding 1 to the original values. Next, the RTI values were weighted based on the proportion of the occupation in its respective sector. This was done in one of two ways. There was no available data on the proportion of each occupation in its sector for Guangzhou, instead we used the Hong Kong's Census and Statistics Department's table of employed persons by industry and occupation of main employment as an estimate. The weighted RTI value was found by multiplying the RTI value ranging from 0 to 2 and the proportion of the occupation in its sector. However, there were four sectors excluded from the table, namely agriculture, forestry, animal husbandry and fishery, mining, production and supply of electricity, heat, gas and water, and scientific research and technical services. For these sectors we assumed that each occupation had equal proportions in the sector. Then, the weighted RTI values of all the occupations in a sector are added up to produce a total fixed RTI of the sector from 0 to 2.

TABLE II: RTI of Individual Sectors (RTI_i)

Industrial Sector	**RTI**
Agriculture, forestry, animal husbandry and fishery	0.2682987706
Mining	0.4320707167
Manufacturing	0.9871191077
Production and Supply of Electricity, Heat, Gas and Water	0.62095238
Construction	0.48595964
Wholesale and Retail Trade	0.2522632123
Transport, Storage and Post	0.6054595786
Hotels and Catering Services	0.4015444909
Information Transmission, Software and Information Technology	0.4907146673
Financial Intermediation	0.7704807878
Real Estate	0.6279826691
Leasing and Business Services	0.7693711009
Scientific Research and Technical Services	0.5960266867
Management of Water Conservancy, Environment and Public Facilities	0.1626908691
Service to Households, Repair and Other Services	0.3301721118
Education	0.2612396285
Health and Social Work	0.2228520921
Culture, Sports and Entertainment	0.3379882028
Public Management, Social Security and Social Organizations	0.2837289066

After finding the fixed RTI of each sector (RTI_i), the proportion of each sector in the labour market each year (s_{it}) is found by dividing the average number of workers in each sector by the total number of workers in the workforce. Lastly, the dependent variable, RTI of the labour market in year t is calculated by using the fixed RTI of each sector and proportion of each sector in the labour market using this formula:

$$\mathrm{RTI}_t = \frac{\sum_{i=1}^{19} (\mathrm{RTI}_i \cdot s_{it})}{19} \tag{1}$$

TABLE III: X and Y variables

Year	**Adjusted R&D Expenditure (10 million yuan)**	**RTI of labour market**
1995	20.76	0.037622064
1996	21.83	0.037082828
1997	27.29	0.036743326
1998	33.86	0.0369175
1999	41.34	0.036553363
2000	40.67	0.036186261
2001	46.22	0.035626328
2002	50.55	0.036336111
2003	53.99	0.032544775
2004	63.40	0.032791201
2005	66.37	0.032469089
2006	70.49	0.032705473
2007	133.21	0.033428524
2008	161.85	0.033028838
2009	219.75	0.033362508
2010	239.61	0.033176061
2011	280.96	0.03441091
2012	301.21	0.032964085
2013	326.19	0.032752378
2014	364.64	0.033570055
2015	408.04	0.032853991
2016	478.11	0.032601125
2017	543.94	0.031992126
2018	598.79	0.029615506
2019	656.47	0.029524064
2020	731.51	0.028889246
2021	823.34	0.029002906
2022	901.29	0.028956328

3) Control Variables: There were four control variables in this model: real GDP per capita, total population of the Guangzhou workforce, lending interest rate, and the average real wages. There is a possibility that the increase in R&D expenditure is simply due to economic growth and an increase in GDP. Economic growth may also result in decrease in RTI of the labour market as the higher RTI sectors shrink (agriculture, manufacturing, mining etc.) and the lower RTI sectors expand (service, education etc.) Thus, real GDP per capita is a necessary control variable. Furthermore, the decrease in RTI could be caused by demographic changes. Workers in high RTI jobs may have retired or moved out of Guangzhou, and new workers entering from outside Guangzhou may be employed in low RTI jobs. This does not represent the internal shift of workers from high to low RTI jobs that we want to observe. Therefore, to account for the change in numbers in the workforce, the total population of the labour market is chosen as a control variable. Lending interest rates and average real wages make the last two control variables because they represent the price of capital and labour respectively. Employers may have implemented more capital than labour in certain sectors (and vice versa) because it is cheaper to do so and not because of the new technology created by R&D centres. This will affect the RTI of the labour market, for example, more capital and fewer workers in higher RTI sectors will decrease the overall RTI of the labour market.

TABLE IV: Control Variables

Year	Lending interest rate (%)	Real GDP per capita (yuan)	Total population of the labour market	Average real wages (yuan)
1995	12.06	2484	2080207	1580
1996	10.08	2561	2027943	1672
1997	8.64	2742	1994095	1816
1998	6.39	3030	1977548	2029
1999	5.85	3326	1904490	2331
2000	5.85	3605	1823317	2672
2001	5.85	4062	1751570	3133
2002	5.31	4719	1820544	3653
2003	5.31	5595	1889518	4091
2004	5.58	6579	1928522	4420
2005	5.58	7601	1965337	4751
2006	6.12	8634	2034522	4983
2007	7.47	9325	2182444	5332
2008	5.31	9667	2216165	5683
2009	5.31	10210	2271728	6324
2010	5.81	10781	2393766	6786
2011	6.56	11002	2959425	6783
2012	6.00	10949	3089317	7305
2013	6.00	11641	3028729	7783
2014	5.60	11739	3104549	8105
2015	4.35	11921	3064510	8713
2016	4.35	11852	3109354	9312
2017	4.35	11856	3146946	10075
2018	4.35	11824	3308859	11158
2019	4.35	12728	3713798	11962
2020	4.35	12775	3854110	12758
2021	4.35	14038	4016567	13474
2022	4.35	14009	4038599	13891

4) Summary Statistics: The summary statistics are displayed in Appendix Table VIII. Overall, there has been a large increase in adjusted R&D expenditure, adjusted GDP per capita, and the population of the workforce, as shown by their ranges of 8,805.3 million yuan, 11,554 yuan, and 2,287,029 workers, respectively. Furthermore, the average wages have nearly grown tenfold from the low minimum of 1,580 yuan in 1995. This data indicates the rapid growth of Guangzhou over the past twenty years. However, there is a large standard deviation for adjusted R&D expenditure, adjusted GDP per capita, and the population of the workforce, showing that there was considerable variation around the mean. The interest rates also varied considerably, but the low mean indicates that the interest rates stayed low for most of the time period. On the other hand, there is a much smaller variation in RTI, implying that the routine task content of jobs in Guangzhou did not change significantly.

IV. Methodology

A. Model Specification

In this model, the x variable is the change in R&D expenditure, the y variable is the change in RTI of Guangzhou's labour market, and the four control variables are the interest rate, population of the labour market, real GDP and average real wages. We ran an ordinary least squares regression in Excel using a 1, 2, 3, 5, 7, and 10 year time lag. However, we realised that average wages is not a suitable control variable as there is reverse causation, where the y variable (change in RTI) will cause a change in average wages. Therefore, there are only three variables, but it should be noted that this might limit the accuracy of our results.

TABLE V: Regression Results for RTI and R&D Expenditure (0-2yrs)

	Current yr	*1 yr*	*2 yrs*
Intercept	0.0319*** (0.002)	0.0338*** (0.002)	0.0348*** (0.002)
Adjusted expenditure in R&D (10 million yuan)	−0.00001** (0.000004)	−0.000009** (0.000004)	−0.000008** (0.000004)
Lending interest rate (%)	0.0001 (0.0002)	−0.00001 (0.0002)	−0.00004 (0.0002)
Total population of labour market	0.000000002 (0.000000002)	0.000000003* (0.000000001)	0.000000003* (0.000000001)
Real GDP per capita (yuan)	−0.0000004*** (0.0000001)	−0.0000005*** (0.0000001)	−0.0000006*** (0.0000001)

Note: Standard errors are reported in parentheses.
*,**,*** indicates significance at the 90%, 95% and 99% level, respectively.

TABLE VI: Regression Results for RTI and R&D Expenditure (3-5yrs)

	3 yrs	*5 yrs*
Intercept	0.0223** (0.011)	0.0158 (0.010)
Adjusted expenditure in R&D (10 million yuan)	0.00001 (0.00002)	0.00003 (0.00002)
Lending interest rate (%)	0.0034*** (0.0009)	0.0056*** (0.0009)
Total population of labour market	−0.000000002 (0.000000007)	−0.0000000008 (0.000000007)
Real GDP per capita (yuan)	−0.0000006 (0.0000006)	−0.000002*** (0.0000006)

Note: Standard errors are reported in parentheses.
*,**,*** indicates significance at the 90%, 95% and 99% level, respectively.

TABLE VII: Regression Results for RTI and R&D Expenditure (7-10yrs)

	7 yrs	*10 yrs*
Intercept	0.0319* (0.016)	0.080*** (0.013)
Adjusted expenditure in R&D (10 million yuan)	−0.0000004 (0.00003)	0.00002 (0.00002)
Lending interest rate (%)	0.0022 (0.001)	−0.0002 (0.001)
Total population of labour market	0.00000001 (0.00000001)	0.000000005 (0.000000009)
Real GDP per capita (yuan)	−0.000004*** (0.0000009)	−0.000006*** (0.0000007)

Note: Standard errors are reported in parentheses.
*,**,*** indicates significance at the 90%, 95% and 99% level, respectively.

V. Results

The regression shows a strong correlation between change in R&D expenditure and change in RTI of the labour market during the first three years. This fast rate of technology transfer provides evidence for the growing support from the government and the private market for technology and AI. In addition, real GDP per capita is found to have a negative relationship with the RTI of the labour market. This is a logical result as workers tend to shift to high-skilled, low-RTI jobs when the GDP per capita increases and the economy grows. However, this correlation could be spurious, where both are happening naturally without direct relation to each other. Fortunately, the chance of this spurious relationship is lowered because the x and y variables are significant in under two years instead of throughout all the time lags. Our hypothesis that there will be a negative relationship between R&D expenditure and the RTI of the Guangzhou labour market is not rejected.

VI. Discussion

Based on our regression results and current literature, we can assess the three theories regarding the future impact of AI and tentatively form a final prediction.

The regression results cannot infer much about the first theory on a productivity J-curve directly as there is no information about the productivity of the Guangzhou workforce. However, some predictions can be made if we assume that the theory is true and that a J-curve will exist. Data has shown that the interest rate is decreasing or staying low and the business services sector is growing. The local and national government has also encouraged businesses to use new technologies or collaborate with research institutes by providing subsidies. This increases the incentive for businesses to invest in human capital and to change their business models to better adapt to AI. In addition, the education sector in Guangzhou is growing, which suggests that more training programmes for workers will be available to provide them with skills to use AI. All of this implies that the productivity J-curve will reach the turning point faster and productivity will start to increase in fewer years compared to the past.

The second theory, which states that AI will create as many or more jobs than it destroys, is plausible. Based on the information collected while creating this model, in the past 40 years there is a trend where certain sectors in Guangzhou have shrunk and other sectors have grown. The sectors that have shrunk are mostly high-RTI and low-skilled, including agriculture, mining, manufacturing and construction. On the other hand, sectors that are growing are low-RTI, high-skilled and more technology-based, such as information transmission, education, and scientific research and technical services. The correlation between change in R&D expenditure and change in RTI of the labour market imply that technology and AI displaces high-RTI jobs. Since most high-RTI sectors in Guangzhou have been shrinking, it is plausible that there will be fewer workers displaced from high-RTI jobs and more workers being employed in the expanding low-RTI sectors. Nevertheless, there are other factors that may change in the future and affect the plausibility of this theory, including demographic changes. For example, as shown in Figure 3 and 4, the median age in China has been increasing and is projected to increase even further. The increase in median age results from a decrease in young workers entering the workforce, decreasing the overall supply of labour. According to the basic laws of supply and demand, this will increase the average wages. This trend is supported by the theory, as AI will replace the lower paying, high-RTI jobs and create as many or more technical and low-RTI jobs. As a result, average wages will rise.

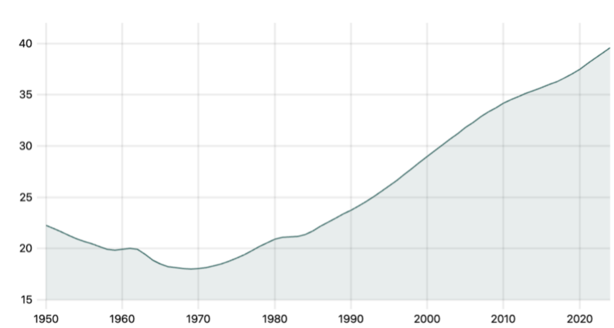

Fig. 3: Historic Median Age of China (1950-2024) (Database.earth, n.d)

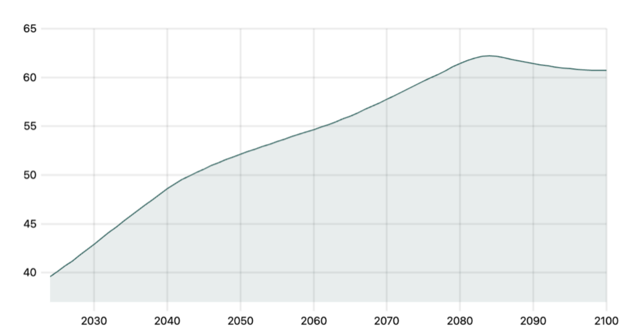

Fig. 4: Future Median Age of China (2024-2100) (Database.earth, n.d)

Lastly, the third theory which argues that a freer labour market will suffer less unemployment from disruptive technology, including AI, could be plausible. The strong correlation between change in R&D expenditure and change in RTI of the labour market for the first 3 years confirm that the rate of technology transfer is high, as technological innovations are commercialised in under 3 years. This implies that Guangzhou has a free labour market and a highly trained, efficient workforce. The third theory is supported by the fact that registered unemployment in Guangzhou has dropped almost continuously from 2009 to 2019, with a dramatic increase only after 2019 due to the COVID-19 pandemic. This suggests that the workforce has not been severely disrupted by the technological changes in the past years. Thus, based on this theory, we can tentatively predict that the impact of AI will be lessened. AI is unlikely to have a very disruptive effect on the Guangzhou labour market because the market is flexible and open. The workforce can quickly adapt to the changes brought by AI, creating and filling in new jobs to reduce unemployment. Moreover, switches between sectors and business can be made swiftly in the open labour market. This flexibility will attract investors and more economic activity, further expanding Guangzhou's economy and workforce.

There are limitations to this study. Firstly, the predictions made above are based on a few assumptions. We assumed that the rise in AI is fundamentally similar to past technological change, and that the state and private market support for AI development is similar to past support for technological development in Guangzhou. These assumptions may be incorrect, as AI is a unique branch of technology and it is extremely difficult to predict the nature and features of AI in the coming decades. Secondly, while the state and private market support for AI development in the past suggests that future support will stay constant or even increase, there is no guarantee that it will definitely be the case as unexpected factors may come into play. Thirdly, the regression results cannot be applied to all three theories. While the correlation between R&D expenditure and RTI of the labour market can lead to suggestions about the second and third theories, more research and analysis is required for an evaluation of the feasibility of the first theory. Furthermore, there are econometrical limitations when the regression model was being built. The y variable, change in RTI of the Guangzhou labour market, could be calculated in different ways. For example, this paper used the RTIs of all the ISCO-08 occupations from the paper by Mihaylov & Tijdens (Mihaylov & Tijdens, 2019). Even though effort was put into determining the best data sources and calculations for the RTI values, the values may differ if an alternate source or method was used. There is also the possibility of a missing variable. Although we have tried to consider all potential control variables, there is still a danger of some unidentified variable that affects both the x and y variable. This possibility is reduced by the fact that the regression results are not consistently significant throughout all the time lags, but it is still there. Lastly, there was no satisfactory method to control for time in the construction of the regression model, leaving the possibility of a spurious relationship between R&D expenditure and RTI of the labour market.

VII. Conclusion

In summary, this paper studies the effects of previous disruptive technology in the Guangzhou workforce by analysing the relationship between R&D expenditure and RTI of the workforce. The results of the OLS regression suggest that the two variables have a significant negative correlation during the first three years, which supports our hypothesis based on the RBTC approach. In addition, the results indicate a fast rate of technology transfer in Guangzhou, which is supported by evidence from the literature. Combined with information from research studies, we can tentatively predict that AI might not have very damaging effects on Guangzhou's labour market, contrary to what many have feared. The overall net impact of AI on jobs may be neutral, due to the creation of new jobs elsewhere in the economy and the rapid switches between sectors facilitated by an efficient workforce and a flexible labour market. In practice, these findings can help assess the challenge of the rise of AI in Guangzhou from a policy perspective.

Going forward, more research should be conducted into other theories of how AI might affect the labour market as this paper only discusses three. Moreover, researchers should explore different methods of calculating RTI values, compare the values to see how they vary for each method, and establish a method that produces the most accurate results. Lastly, exploring the demographic changes that might change the structure of the labour market and, therefore, the way AI affects it will help bring this research forward.

References

[1] Arntz, M., T. Gregory, U. Zierahn (2016) The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis. *OECD Social, Employment and Migration Working Paper No. 189*. https://doi.org/10.1787/1815199X

[2] Arntz, M. (2017) Revisiting the Risk of Automation. *Economic Letters* 159, 157-160. https://doi.org/10.1016/j.econlet.2017.07.001

[3] Autor, D.H., Levy, F., & Murnane, R.J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4), 1279-1333. https://doi.org/10.1162/003355303322552801

[4] Balsmeier, B., & Woerter, M. (2019). Is this time different? How digitalization influences job creation and destruction. *Research Policy* 48(8):103765. https://doi.org/10.1016/j.respol.2019.03.010

[5] BayArea Gov. (2023) Guangzhou: Major Social and Economic Indicators. *BayArea Gov*. https://www.bayarea.gov.hk/en/about/guangzhou.html

[6] Berriman, R., & Hawksworth, J. (2017). Will Robots Steal our Jobs? The Potential Impact of Automation on the UK and other Major Economies. *PwC UK Economic Outlook*, March.

[7] Brynjolfsson, E., Rock, D., & Syverson, C. (2021). The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal: Macroeconomics*, 13(1): 333–72.

[8] Burt, E. (2024) Guangzhou: Sustaining a 240-year trade relationship with China's economic and cultural powerhouse. *ConGen Guangzhou*.

[9] Cao, A. (2022). Big Tech companies in Guangzhou heat up China's metaverse race through new local government-backed industry alliance. *South China Morning Post*.

[10] China News Guangdong. (2023). Guangzhou Nansha Plan strengthens its capabilities to build the international science and technology innovation centre in the Greater Bay Area. *China News Guangdong*.

[11] Dağli, I., & Kösekahyaoğlu, L. (2021). Skill-Biased Technological Change and Job Polarisation and Evidences from the Turkish Labor Market. *Journal of Economics, Business and Political Researches*, 6(16), 348-359. https://doi.org/10.25204/iktisad.901039

[12] Database.Earth (n.d) Median Age of China. *Database.Earth*. https://database.earth/population/china/median-age

[13] Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2017). German Robots - The Impact of Industrial Robots on Workers. *CEPR Discussion Paper No. 12306*. CEPR Press, Paris & London.

[14] Dezan Shira & Associates (n.d) Guangzhou: Industry, Economics and Policy. *China Briefing*.

[15] Dickson, B. (2022). AI's J-curve and Upcoming Productivity Boom. *AI Business, TechTalks*.

[16] Duan, S.X., Deng, H., & Wibowo, S. (2023). Exploring the impact of digital work on work-life balance and job performance: a technology affordance perspective. *Information Technology & People* 36(5):2009–2029.

[17] Ekelund, H. (2024) Why there will be plenty of jobs in the future - even with artificial intelligence. *World Economic Forum*.

[18] Ernst, E., Merola, R., & Samaan, D. (2019) Economics of Artificial Intelligence: Implications for the Future of Work. *IZA Journal of Labor Policy*, 9:4.

[19] Fatun, M., & Pazour, M. (2021) Modelling the impact of Artificial Intelligence on the labour market in Czechia. *Central European Journal of Public Policy*, 15:2 (31-42).

[20] Fei, Yu., & Qing, Fu. (2016) Spotlight on Guangzhou. *Springer Nature*.

[21] Ghosh, I. (2019). These will be the most important cities by 2035. *World Economic Forum*.

[22] Gries, T., & Naudé, W. (2018) Artificial Intelligence, Jobs, Inequality and Productivity: Does Aggregate Demand Matter? *IZA Discussion Paper No. 12005*.

[23] Griliches, Z. (1969). Capital-skill complementarity. *Review of Economics and Statistics*, 51, 465-468.

[24] Guangzhou Statistics Bureau (1984-2022) Guangzhou Statistics Yearbooks. *Guangzhou Statistics Bureau*.

[25] Hawksworth, J., Berriman, R., & Goel, S. (2018). Will robots really steal our jobs? An international analysis of the potential long term impact of automation. *PriceWaterhouseCoopers*.

[26] Hu, A.G. (2001). Ownership, Government R&D, Private R&D, and Productivity in Chinese Industry. *Journal of Comparative Economics*, 29(1): 136-157.

[27] Huang, Weilong., & Fan, Zhenguo. (2020) Innovation and Coordination: a SWOT Analysis of Artificial Intelligence Industry in Guangdong Province. *E3S Web of Conferences*, 235.

[28] ILO (2008). International Standard Classification of Occupations 2008. *International Labour Office*, Geneva.

[29] Institute of Population and Labor Economics of Chinese Academy of Social Science (2017). The China Urban Labor Survey (CULS). *Peking University Open Research Data Platform*.

[30] Jin, X., & Weber, K. (2008). The China Import and Export (Canton) Fair: Past, Present, and Future. *Journal of Convention & Event Tourism*, 9(3), 221-234.

[31] Korinek, A.; J. Stiglitz (2017) Artificial Intelligence and its Implications for Income Distribution and Unemployment. *NBER Working Paper No. 24174*.

[32] Lewandowski, P., Park, A., & Schotte, S. (2020). The Global Distribution of Routine and Non-Routine Work. *IZA Institute of Labor Economics, IZA DP No.13384*.

[33] Liu, N., Gu, X., & Lei, C.K. (2022). The equilibrium effects of digital technology on banking, production, and employment. *Finance Research Letters*, 49:103196.

[34] Marcolin, L., Miroudot S., & Squicciarini, M. (2016). The Routine Content of Occupations: New Cross-country Measures Based on PIAAC. *OECD Science, Technology and Industry Working Papers*.

[35] Méda, D. (2016) The Future of Work: The Meaning and Value of Work in Europe. *ILO Research Paper No. 18*.

[36] Mihaylov, E., & Tijdens, K.G. (2019). Measuring the Routine and Non-Routine Task Content of 427 Four-Digit ISCO-08 Occupations. *Tinbergen Institute Discussion Paper 2019-035/V*.

[37] Nature Index (2023). Leading 200 Science Cities. *Springer Nature*.

[38] Nowak, Spike. (2012) On the Fast-Track: Technology Transfer in China. *China Briefing*.

[39] NUS. (2024) New NUS Research Institute in Guangzhou partners industry leaders in research, enterprise and talent development. *NUS News*.

[40] OECD (2014). Programme for the International Assessment of Adult Competencies (PIAAC) Survey, Round 1. *OECD*.

[41] OECD (2017). Programme for the International Assessment of Adult Competencies (PIAAC) Survey, Round II. *OECD*.

[42] OECD (2019). Programme for the International Assessment of Adult Competencies (PIAAC) Survey, Round III. *OECD*.

[43] O'Neill, M. (2018). Innovation push elevates Guangzhou's role in Greater Bay Area project. *MacaoMagazine*.

[44] Qiu, Q.L. (2023) Guangdong to establish national hub for the AI industry by 2025. *China Daily Information Co.*

[45] Rostron, M.T. (2024). How Will Artificial Intelligence Affect Jobs 2024-2030. *Nexford University*.

[46] Shen, Y., & Zhang, X. (2024). The impact of artificial intelligence on employment: the role of virtual agglomeration. *Humanities & Social Sciences Communications*, 11, 122.

[47] Shine, I. (2023) We often hear that AI will take our jobs. But what jobs will it create? *World Economic Forum*.

[48] Stropoli, R. (2023) AI Is Going to Disrupt the Labor Market. It Doesn't Have to Destroy It. *Chicago Booth Review*.

[49] Tech Transfer Central. (n.d). What is technology transfer? *TechPipeline*.

[50] Violante, G.L. (2008). Skill-Biased Technical Change. *The New Palgrave Dictionary of Economics*. Palgrave Macmillan.

[51] Welch, F. (1970). Education in production. *Journal of Political Economy*, 78, 35-59.

[52] WHO (2019). Classifying Health Workers: Mapping Occupations to the International Standard Classification. *World Health Organization*.

[53] World Bank (2016a). STEP Skills Measurement Household Survey, Wave 1. *World Bank*.

[54] World Bank (2016b). STEP Skills Measurement Household Survey, Wave 2. *World Bank*.

[55] World Bank (2018). STEP Skills Measurement Employer Survey, Wave 3. *World Bank*.

[56] Xinhua. (2024) China accelerates AI development to build AI innovation Center. *The State Council of the People's Republic of China*.

VIII. Appendix

A. OLS Regression Output Spreadsheet

The subheading contains the link for the regression results for all 8 regressions: one with only the x and y variables, another with the x, y and the control variables, and the rest with a 1, 2, 3, 5, 7, and 10 time lag.

TABLE VIII: Summary Statistics

	RTI	**Adj. R&D expenditure (10 million yuan)**	**Interest rate**	**Population**	**Adj. GDP per capita**	**Average wage**
Mean	0.033347	275.2029	5.91	2596302.82	8616.25	6520.535714
Standard Error	0.000491	51.06001	0.33934	140339.659	742.729085	711.8021199
Median	0.032996	190.8	5.59	2243946.5	9938.5	6003.5
Standard Deviation	0.002597	270.1842	1.795618	742607.675	3930.152901	3766.502784
Range	0.008733	880.53	7.71	2287029	11554	12311
Minimum	0.028889	20.76	4.35	1751570	2484	1580
Maximum	0.037622	901.29	12.06	4038599	14038	13891
Count	28	28	28	28	28	28

The Impact of Tourism on Detroit City Employment

Salmaan Seraji-Bozorgzad

Wayne State University

Abstract— **Following the 2013 bankruptcy, Detroit has pursued a redevelopment strategy which significantly focuses on downtown tourism; however, there are concerns as to whether these investments benefit the city's resident population or primarily serve regional consumers and the surrounding metropolitan and statistical areas. Given the economic spatial disconnect between the concentration of investment in the downtown core and the surrounding city neighborhoods, there is a risk that regional aggregate growth may overlook local disparities. This paper empirically analyzes the short-run and long-run economic impacts of downtown tourism on Detroit's resident labor market. By specifically utilizing resident-based employment and unemployment series, the analysis isolates the labor market experiences of Detroiters. The analysis employed cointegration tests, vector autoregressive models, and error correction models. The results confirm a stable long-run equilibrium, demonstrating that resident labor market conditions are structurally linked to downtown visitation. In the short-run, hotel occupancy growth acts as a robust, uni-directional, leading indicator, predicting employment growth and unemployment decline in the following period. In contrast, monthly visitation growth functions as a significant contemporaneous driver of change in the labor market. Furthermore, the error correction models reveal that employment gains exhibit persistence while unemployment decline is immediate but short-lived without further developments. These findings validate tourism as a driver of local economic development for Detroit residents.**

I. Introduction

In recent history, Detroit has been the object of fiscal and labor market pressures. From 2007 to 2009, the Great Recession resulted in unprecedented demand for charitable support and stabilization programs (Allard, Wathen, and Danziger, 2015). Directly after the Recession, the city's mounting obligations such as service costs, pensions, and borrowing forced the city to file for bankruptcy in 2013 (Naglick, Lavelle, and Moore, 2020). Since exiting bankruptcy, Detroit has pursued redevelopment aimed at broadening its service base and attracting private investment. Investments have targeted neighborhood small-business corridors and the downtown core, including public realm improvements and funding new tech-oriented workspaces (City of Detroit Planning and Development Department, 2024a; Rahal, 2025). The city has also prioritized visitor-focused assets by revitalizing downtown areas such as Grand Circus Park, advancing East Riverfront parks and properties, and building new sports and entertainment venues to attract visitors, new residents, and capital investment (Downtown Detroit Partnership, 2024; Community Foundation for Southeast Michigan, 2024). National travel media, including *Condé Nast Traveler* and *AFAR*, have highlighted Detroit's visitor appeal (Chester, Bansal, and Cohen, 2023; CNT Editors, 2023), and in 2024, the city hosted the NFL Draft, drawing roughly 775,000 attendees downtown and surpassing the previous record of roughly 600,000 people (Lage, 2024). The city's investment into visitor-facing development along with its organization and hosting of major special events motivates this analysis of the economic impact of tourism on Detroit.

Tourism has become one pillar of Detroit's broader redevelopment plan (Detroit Metro Convention & Visitors Bureau, 2020; Visit Detroit, n.d.a). While a growing body of literature analyzes tourism's role in emerging and developing economies (Li et al., 2025; Bakker et al., 2023; Alcala-Ordonez and Segarra, 2025), its impact at the municipal level remains limited. Municipalities like Detroit have complex spatial structures that complicate standard economic frameworks, particularly regarding the distribution of investment, goods, and services. A primary concern regarding Detroit's tourism projects is that they are often focused in the downtown/Midtown core (Visit Detroit, 2025a; Williams, 2025), whereas the majority of city residents live in outlying neighborhoods. This spatial disconnect, risks exacerbating economic disparities between the smaller downtown population and broader residential base (Owens, Rossi-Hansberg, and Sarte, 2020). However, the Detroit municipal government has actively sought to mitigate this divide through policy interventions, such as mandating local hiring requirements for new development projects (Duggan, 2024). Given this context, this paper hypothesizes that growth in downtown visitor activity and hotel occupancy generates a positive economic effect leading to broader resident employment growth and lower unemployment for city residents.

To empirically test for these effects, the analysis examines the short-run and long-run dynamics between two Detroit-specific tourism indicators, the average monthly downtown visitors (AMV) and downtown hotel occupancy recovery (HOR) gathered from the Downtown Detroit Partnership, and the resident employment and unemployment series from the Michigan Labor Market Information database (MILMI).

II. Brief Background and Review

Detroit's pre-2013 trajectory featured compounding fiscal and labor market stress that culminated in filing for Chapter 9 bankruptcy. In 2007, prior to the Great Recession, the city's poverty rate (33.8%) was more than twice the U.S. average (13.0%), and unemployment (14.6%) was nearly

Thank you to Dr. Liang Hu for all your helpful guidance, detailed feedback, and constant encouragement.

four times the national rate (4.1%) (Allard, Wathen, and Danziger, 2015). During the recession, home prices in the Detroit area fell by nearly 50% as tax and mortgage foreclosures surged, amplifying demand for public and charitable assistance (McDonald, 2014). At the same time, Detroit's large geographic footprint sustained high service demands (fiscal and medical support, policing, fire response to increasing vacancies) even as the taxpayer base continued to decline. Despite deep workforce reductions over prior decades, legacy liabilities, including pensions and retiree health, and bonded debt continued to rise through 2013, outpacing an eroding revenue base and straining service delivery (Naglick, Lavelle, and Moore, 2020). Broader shocks, including the auto bankruptcies, pushed unemployment from 14.1% in 2005 to 25.0% in 2009 (State of Michigan DTMB, 2015), while vacancies grew and conventional home-purchase lending fell from close to 8500 loans annually pre-crisis to only around 400 by 2014 (Sugrue, 2014; Phillips, 2021). Concluding that revenues were insufficient to cover operating costs and growing obligations, the State of Michigan initiated a financial review in December 2011, declared a financial emergency in March 2012, and entered a consent agreement with the city (Naglick, Lavelle, and Moore, 2020). After updated reviews in early 2013, Kevyn Orr was appointed emergency manager in March 2013. The emergency manager's office submitted a preliminary financial and operating plan and published a creditor proposal with significant cuts and the suspension of unsecured-debt payments, then in July 2013, the city petitioned for bankruptcy protection (Naglick, Lavelle, and Moore, 2020).

Following the 2013 filing, the 2014 Plan of Adjustment, supported by the "Grand Bargain" and state oversight via the Financial Review Commission, restructured legacy obligations and financed basic service restoration (City of Detroit, 2014). The Grand Bargain, coordinated outside standard litigation channels, combined philanthropic and state contributions with negotiated creditor concessions and transactions involving city owned assets (including measures to protect the Detroit Institute of Arts collection) to support pension settlements and facilitate exit from bankruptcy. These actions allowed Detroit to exit bankruptcy on December 10, 2014 (Naglick, Lavelle, and Moore, 2020).

Since exiting bankruptcy, the City of Detroit has pursued different strategies to redevelop and grow with a lot of downtown growth being associated with investments by the Rock Ventures family of companies (Naglick, Lavelle, and Moore, 2020). Although there has been significant controversy surrounding the city's financial backing of the companies' projects (Ernsthausen and Elliott, 2019), Rocket Companies has become one of the city's largest employers amassing over 17,000 Detroit employees in 2015 and heavily investing in city revitalization and development efforts (City of Detroit Mayor's Office, 2021; Naglick, Lavelle, and Moore, 2020). In addition to new employers and private investment, the city has turned to tourism to support local growth and economic development. An early example of tourism oriented development was the construction of Little Caesars Arena (LCA), which began in 2014. The new arena replaced Joe Louis Arena, the former home of the Red Wings, and the Palace of Auburn Hills, the former home of the Detroit Pistons, bringing more sports fans downtown and expanding event related employment, including local labor commitments in development agreements (though the enforcement of these commitments has been contested) (Motor City Electric Co., n.d.; Aguilar, 2017). Taken together, Detroit's fiscal restructuring and subsequent redevelopment created the context for renewed private investment and an explicit turn to tourism oriented assets and growth.

A. *Tourism Supported Development in Detroit*

A substantial literature examines tourism's impact on key economic indicators. One notable impact of tourism is its ability to generate employment given that it forms a labor intensive sector (Meyer and Meyer, 2016). Detroit's LCA and the East-Riverfront parks and properties illustrate distinct channels through which tourism development can generate jobs. First, both projects created construction-phase employment for city residents (Chambers, 2017; Hartig, 2018). Second, they require ongoing staffing, LCA especially, for operations, management, and maintenance (Aguilar, 2017; Aguilar, 2024; Detroit Riverfront Conservancy, n.d.). Third, by increasing foot traffic these developments can help spur private investment in nearby retail and services, supporting additional employment in adjacent areas (Chambers, 2017; City of Detroit Planning and Development Department, n.d.; Noble, 2018; Skidmore, Owings & Merrill, n.d.; City of Detroit Planning & Development Department, 2017).

To grow visitation and support these employment channels, the Special Event Management Team (SEMT) in the Detroit Mayor's Office along with the Downtown Detroit Partnership (DDP), and other non-profit organizations, organize downtown events, including concerts, fashion shows, art expos, and, notably, the NFL Draft (Downtown Detroit Partnership, 2025a; City of Detroit Mayor's Office, n.d.). The 2024 Draft underscores DDP's capacity to execute large-scale events. The three-day program drew more than 775,000 attendees, and the economic impact was estimated to be around $160 million according to analysis from economic consulting firm Anderson Economic Group (Rahal, 2024). Furthermore, Visit Detroit, the region's destination-marketing organization, amplifies this strategy by promoting local events and developments, partnering with businesses and sponsors to increase national—and, increasingly, international (Visit Detroit, 2025b)—visibility while highlighting investment opportunities (Visit Detroit, n.d.b). Complementing these efforts, recent and planned hotel investments align with this strategy to support multi-day events and sustained visitation. For example, the new AC Hotel on Woodward Avenue is positioned to serve attendees for LCA, Ford Field, and Comerica Park, all within or adjacent to the downtown–Midtown core (Williams, 2025). In addition, projects such as the Michigan Central Station redevelopment led by the Ford Motor Company include public-facing components that can draw visitors while anchoring new employment activity

(Williams, 2025). Together, destination marketing, event-production capacity, and expanded lodging supply constitute the institutional and physical infrastructure for tourism-driven demand.

To quantify the regional economic impact of these activities, Tables I and II summarize visitor spending and tourism-supported employment and labor income for Wayne County. The pattern aligns with the broader narrative of steady expansion through 2019, with a sharp decrease in 2020, and a partial recovery by 2023. From 2016 to 2019, total visitor expenditure rose about 9.7%, with lodging and food and beverage up roughly 11.8% and 13.6%, respectively; from 2020 to 2023, total expenditure rebounded about 65%, remaining roughly 11.5% below the 2019 peak.

Tourism-supported labor metrics show a similar arc. From 2016 to 2019, direct and total employment each increased by roughly 4%, while direct labor income and total labor income rose about 11% (Tourism Economics, n.d.). The pandemic produced a marked drop in 2020; however, by 2023, total labor income slightly exceeded its 2019 value, and direct labor income was just below it, while employment remained below 2019 levels (by roughly 14–18%). From 2020 to 2023, labor income growth ($\sim$40% in both direct and total measures) outpaced employment growth ($\sim$28–33%), consistent with rising earnings per job or shifts toward higher-value positions (Tourism Economics, n.d.).

While these regional aggregates indicate a strong recovery, the broader literature cautions that tourism-led growth does not guarantee equitable distribution. A large body of work examines and demonstrates short-run causal effects from tourism to economic activity. While much of this literature is at the national level, evidence from Cape Town, South Africa, indicates that tourism can be causally related to poverty alleviation and economic growth (Garidzirai and Nguza-Mduba, 2020). Additionally, the literature demonstrates evidence of long-run effects with a review of more than thirty studies reporting unidirectional causality from tourism to economic development in a majority of cases (Alcala-Ordonez and Segarra, 2025); however, distributional concerns remain persistent. For example, in an analysis of Peruvian departments, Llorca-Rodriguez and Carmen Maria found that although tourism reduces poverty overall, the poorest residents do not receive the full benefits (Llorca-Rodriguez, Casas-Jurado, and Garcia-Fernandez, 2017). Similar equity concerns have been raised regarding Detroit's concentration of investment in the downtown/Midtown "7.2" square miles (Moskowitz, 2015).

The city has, however, initiated planning efforts aimed at more inclusive neighborhood investment. Recent plans for the Greater Warren–Conner area and for the Brightmoor neighborhood outline strategies for commercial revitalization that account for existing urban structures and facilities. Proposed actions include establishing local retail and commercial nodes and activating underutilized properties with food-truck vendors, local markets, event programming, infrastructure upgrades, and landscaping enhancements (City of Detroit Planning and Development Department, 2024a). These economic development plans are consistent with present literature that discusses equitable strategies to develop the city's neighborhoods (Owens, Rossi-Hansberg, and Sarte, 2020). The city's plans also emphasize parks, open-space improvements, and youth development. In Brightmoor, the city has incorporated resident input directly into project selection and design (City of Detroit Planning and Development Department, 2024b). While not all elements are strictly "tourism," several features—such as re-purposing vacant buildings and lots for restaurants and local retail—are consistent with visitor-facing development. Such initiatives align with policy approaches that seek to coordinate minimum investment in targeted neighborhoods to move them toward a developed equilibrium (Owens, Rossi-Hansberg, and Sarte, 2020). Together, these efforts signal attempts to balance downtown-oriented attractions with neighborhood-scale investments that serve residents while capturing some tourism-related gains.

Ultimately, while county aggregates provide regional context and the neighborhood plans proposed by the city outline long-term development goals, neither captures the immediate high-frequency dynamics of Detroit's local labor market. For these reasons, the Detroit-focused measures used in this paper offer a distinct analytical advantage. Unlike the annual aggregates in Tables I and II, which do not capture within-year variations that matter for how month-to-month changes in tourism relate to employment, the average monthly visitors (AMV) and hotel occupancy recovery (HOR) datasets track changes in downtown Detroit foot traffic and lodging at monthly frequency. Importantly, these are analyzed with resident-based employment and unemployment series which specifically describe the local labor market conditions for city residents, rather than the broader metropolitan statistical area. By isolating the resident population, this approach ensures that the analysis does not conflate regional prosperity with local economic well-being. This specific scope directly addresses the concerns regarding the marginalization of Detroiters in development discourse by placing their actual labor market experiences at the center of the empirical inquiry.

III. Methodologies

A. Data Description

The AMV time series reported by the DDP reflects the monthly estimates of visitor activity within the defined downtown area. On the other hand, HOR is a monthly percentage measure benchmarked to 2019, defined as the current month's occupancy rate relative to the corresponding month in 2019 (2019=100) (Downtown Detroit Partnership, 2025b). Both HOR and AMV are the tourism indicators used in the analysis, however, they have different sample periods. While AMV spans from January 2019 to June 2025, HOR spans from January 2020 to December 2024. This requires adjustments to the Detroit labor series so that the time series analysis is aligned at monthly frequency. The Detroit-city employed persons and unemployed persons series are split into separate datasets to match each of the respective sample periods. The labor series each reflect the number of employed

TABLE I: Visitor Expenditures in Wayne County (millions of dollars)

Year	Lodging	Food and Beverages	Retail	Recreation	Transport	Total
2016	$966.61	$952.31	$562.01	$2,075.91	$2,533.23	$7,090.07
2017	$1,045.15	$1,011.70	$580.59	$2,110.46	$2,608.99	$7,356.90
2018	$1,103.97	$1,044.40	$590.17	$2,225.55	$2,702.40	$7,666.50
2019	$1,080.25	$1,082.07	$612.71	$2,235.36	$2,769.00	$7,779.40
2020	$556.25	$682.49	$474.99	$1,074.64	$1,374.42	$4,162.80
2021	$783.80	$877.40	$562.30	$1,453.40	$1,822.30	$5,499.20
2022	$946.30	$1,035.40	$617.60	$1,658.60	$2,120.60	$6,378.50
2023	$1,021.60	$1,180.80	$672.10	$1,764.90	$2,247.20	$6,886.60
2024	$958.90	$1,262.40	$709.20	$1,932.50	$2,471.10	$7,334.10

Source: Tourism Economics, an Oxford Economics Company; figures published by the State of Michigan.

TABLE II: Tourism-Supported Employment and Labor Income (Wayne County)

Year	Direct Employment	Labor Income†	Total Employment	Total Labor Income
2016	42,951	$2,018.42	66,219	$3,412.39
2017	43,508	$2,102.78	67,200	$3,533.90
2018	44,377	$2,149.42	68,190	$3,621.02
2019	44,833	$2,242.77	68,982	$3,794.46
2020	27,794	$1,511.02	46,407	$2,774.39
2021	30,700	$1,764.10	50,591	$3,158.70
2022	35,164	$2,005.10	56,611	$3,561.80
2023	36,973	$2,144.80	59,523	$3,891.80
2024	37,496	$2,236.60	60,365	$4,138.30

† Direct labor income (millions of dollars). Total labor income includes indirect and induced effects. Income values are in millions of dollars. Source: Tourism Economics, an Oxford Economics Company; figures published by the State of Michigan.

and unemployed residents in the Detroit-city area defined by the state of Michigan.

Since both labor series and AMV are level time series while HOR is an index relative to a corresponding month in 2019, analysis in levels can yield coefficients with unclear units leading to uninterpretable results. Additionally, since all the series include data from the COVID-19 pandemic era the data follows a visual pattern of exponential growth quickly after the quarantine restrictions were eased and eventually lifted. To improve the interpretability of the results and account for the observed exponential behavior of the data, logarithmic transformations are used for each of the series.

B. Unit-root Tests

The first steps taken in the analysis were testing each of the series for unit roots using the augmented Dickey-Fuller (ADF) test and the Phillips-Perron (PP) test. The ADF test is commonly used in cases where the process X_t is not necessarily AR(1) and may have more dependence. Consider the process,

$$\begin{aligned}\nabla X_t &= a_0 + a_1 t + \delta X_{t-1} \\ &\quad + \delta_1 \nabla X_{t-1} + \cdots + \delta_{p-1} \nabla X_{t-p+1} + w_t. \end{aligned} \tag{1}$$

The null hypothesis of the ADF test is $H_0 : \delta = 0$ with a test statistic,

$$T := \frac{\hat{\delta}}{\mathrm{SE}(\hat{\delta})} \tag{2}$$

so rejecting the null hypothesis would mean the process has $\delta < 0$ implying there does not exist a unit root. The PP test is similar to the ADF test, but it does not use lags, $\nabla X_{t-1}, ..., \nabla X_{t-p+1}$. Instead the test directly accounts for autocorrelation within the error terms, w_t. In both tests, the null hypothesis is that the series has a unit root, and is therefore non-stationary. Due to the limited sample size of the data, the Elliott, Rothenberg and Stock (ERS) unit root test is used to supplement both unit root tests as both the ADF and PP tests rely on larger sample sizes. To improve the power of the unit root test, ERS uses an ADF-type test applied to the detrended data without intercept. Since ERS has been shown to exhibit superior power properties in small samples compared to ADF and PP, when the ERS test conflicts with the ADF and PP tests, ERS is treated as the deciding test. All three tests are crucial for analyzing long-run dynamics using cointegration and modeling the relation between each of the series using dynamic regression models.

C. Phillips-Ouliaris Cointegration Test

The Phillips-Ouliaris (PO) cointegration test was used to determine whether the series are cointegrated in levels. The PO test determines cointegration between two time variables by first regressing the dependent variable (Y_t) on the independent variable (X_t) using

$$Y_t = \beta X_t + \varepsilon_t, \tag{3}$$

and then running an ADF test on ε_t. However, since ε_t is the residual series, corrections are made on the distribution of the test statistic T in the ADF test (2). One of the assumptions of the PO test is that both X_t and Y_t are non-stationary. The null hypothesis (H_0) of the test is that X_t and Y_t are not cointegrated, and this is evaluated using the adjusted ADF test on ε_t. If the residual series are stationary, then H_0 is rejected and the series are said to be cointegrated, meaning there exists non-zero constants a and b such that

$$aX_t + bY_t, \tag{4}$$

is stationary. Equivalently, cointegration can be interpreted as the existence of a long-run equilibrium relationship between X_t and Y_t. One of the challenges with using the PO test is that the sample size significantly influences the power of the test. Since the sample size of the series in this analysis have inherent limitations, failing to reject the null hypothesis at certain significance levels is taken under scrutiny by further examining the power of the test result. This is done by estimating the cointegration regression, fitting an ARMA/AR model to the residual series ε_t, and simulating new series using the observed $\hat{\beta}$ and the fitted residual process.

D. Cross Covariance and Cross-Correlation Function

Cross-correlation analysis is used to determine whether movements in visitors or hotel occupancy, X_t, lead movements in employment or unemployment levels, Y_t. That is, whether X_{t-k} is highly correlated with Y_t for negative lags $k < 0$. The cross covariance $\gamma_k(X, Y)$ and cross-correlation function (CCF) $\rho_k(X, Y)$ are defined as

$$\gamma_k(X, Y) = E[(X_{t+k} - \mu_x)(Y_t - \mu_y)] \tag{5}$$

$$\rho_k(X, Y) = \frac{\gamma_k(X, Y)}{\sigma_x \sigma_y} \tag{6}$$

and the estimators of $\gamma_k(X, Y)$ and $\rho_k(X, Y)$ are given by

$$c_k(X, Y) = \frac{1}{n} \sum_{t=1}^{n-k} (X_{t+k} - \bar{X})(Y_t - \bar{Y}) \tag{7}$$

$$r_k(X, Y) = \frac{c_k(X, Y)}{\sqrt{c_0(X, X)}\sqrt{c_0(Y, Y)}} \tag{8}$$

A negative lag $k < 0$ measures the correlation between X_{t+k} and Y_t, so that X is said to lead Y when $\rho_k(X, Y)$ is large for $k < 0$.

E. Dynamic Regression Model

Motivated by the CCF results, dynamic regression models are used to determine the predictive power and directional effect of one variable on the other. Suppose it is concluded that X_t leads Y_t by one period ($k = 1$), then the dynamic regression between X_t and Y_t is represented as,

$$Y_t = \alpha + \beta_1 X_{t-1} + w_t \tag{9}$$

where X_{t-1} is a lagged regressor for Y_t and w_t is the error term. Additionally, to account for the possible contemporaneous relation between Y_t and X_t (9) can be expanded,

$$Y_t = \alpha + \beta_0 X_t + \beta_1 X_{t-1} + w_t \tag{10}$$

and then both estimates, $\hat{\beta}_0$ and $\hat{\beta}_1$, can be interpreted as elasticities for the contemporaneous and lagged regressors X_t and X_{t-1}, respectively. Serial correlation is assessed using the Box-Ljung test on the residuals. When residual autocorrelation is detected, inference is reported using heteroskedasticity and autocorrelation consistent (HAC) standard errors calculated using Newey-West estimators. The dynamic regression models illustrated serve as outlines for the regressions that are used to determine the predictive power and directional effects of HOR and AMV on the employed and unemployed persons series.

F. Vector Autoregressive Models and Causality Tests

Building off the results of the dynamic regression modeling, the vector autoregressive (VAR) model allows for simultaneous analysis of X_t and Y_t in a joint system. In the dynamic regression model, X_{t-k} acts as a regressor while Y_t acts as the dependent variable. The VAR model estimates both equations of X_t and Y_t together, treating each series as endogenous and allowing for cross-lag effects in addition to contemporaneous correlations. In general, VAR is an extension of the autoregressive model

$$X_t = aX_{t-1} + w_t \tag{11}$$

to a multidimensional vector X_t and a multivariate white-noise w_t. (X_t, Y_t) is called a VAR(1) process if and only if

$$\begin{bmatrix} X_t \\ Y_t \end{bmatrix} = \begin{bmatrix} \theta_{11} & \theta_{12} \\ \theta_{21} & \theta_{22} \end{bmatrix} \begin{bmatrix} X_{t-1} \\ Y_{t-1} \end{bmatrix} + \begin{bmatrix} w_{x,t} \\ w_{y,t} \end{bmatrix} \tag{12}$$

Given this definition, the VAR(1) process can be rewritten using

$$\begin{aligned} Z_t &:= \begin{bmatrix} X_t \\ Y_t \end{bmatrix}, \quad \Theta_1 := \begin{bmatrix} \theta_{11} & \theta_{12} \\ \theta_{21} & \theta_{22} \end{bmatrix}, \\ Z_{t-1} &:= \begin{bmatrix} X_{t-1} \\ Y_{t-1} \end{bmatrix}, \quad w_t := \begin{bmatrix} w_{x,t} \\ w_{y,t} \end{bmatrix} \end{aligned} \tag{13}$$

in the same form as (11)

$$Z_t = \Theta_1 Z_{t-1} + w_t \tag{14}$$

The $\mathrm{VAR}(p)$ model is similarly defined; however, for the purposes of this paper $\mathrm{VAR}(1)$ is sufficient according to the Akaike Information Criterion $(\mathrm{AIC}(n))$ and Portmanteau test on the residuals of the VAR model. After estimating the VAR(1) model, the bivariate white-noise w_t is summarized by reporting the standard deviation of its components, $\mathrm{sd}(w_{X,t})$ and $\mathrm{sd}(w_{Y,t})$, along with the contemporaneous correlation, $\mathrm{cor}(w_{X,t}, w_{Y,t})$. The standard deviations reflect the magnitude of period shocks to each series that are not explained by the VAR lag structure, and the correlation indicates the extent to which unexpected shocks to X_t and Y_t tend to occur in the same period. A positive correlation indicates that positive shocks to X_t are more likely to coincide with positive shocks to Y_t.

By using the VAR model, formal tests of predictive and contemporaneous relations are established using the Granger causality and instantaneous causality tests. Granger causality acts as a restriction on the VAR lag coefficients where in the Y_t equation of a given $\mathrm{VAR}(p)$ model, X_t fails to Granger-cause Y_t if all coefficients on $X_{t-1}, \ldots, X_{t-p}$ are jointly zero. This corresponds to an F-test of the null hypothesis H_0 :

$\theta_1 = \ldots = \theta_p = 0$ wherein X_t is said to Granger-cause Y_t if there exists $\theta_k \neq 0$ for some $k \in \{1, \ldots, p\}$. Instantaneous causality is assessed by testing whether the contemporaneous covariance between the bivariate white-noise innovation vectors is zero. The null hypothesis in the instantaneous causality test is defined as $H_0 : \text{cov}(w_{X,t}, w_{Y,t}) = 0$ where rejection of H_0 indicates that shocks to X_t and Y_t co-move within the same period, consistent with the contemporaneous association not accounted for by lagged dynamics in the VAR model. The findings of the VAR models, Granger-causality and instantaneous-causality tests point to the presence of both directional predictive links through lagged effects and strong contemporaneous co-movement in the residuals, indicating whether overall visitation and hotel occupancy contain meaningful short-run information for fluctuations in the labor series while also sharing within-period shocks with the employed and unemployed persons series.

G. Error-Correction Models

Building off the PO cointegration test results, the error correction model (ECM) serves as a framework to combine short-run dynamics in differences with the long-run equilibrium relation implied by cointegration. The ECM incorporates the long-run equilibrium into a short-run regression by including an error correction term (ξ_{t-1}), defined as the lagged cointegrating residual,

$$\xi_{t-1} := \varepsilon_{t-1} = Y_{t-1} - \beta X_{t-1}. \tag{15}$$

Intuitively, ξ_{t-1} measures the deviation from the long-run equilibrium in the previous period. If Y_{t-1} is unusually high relative to the equilibrium value implied by X_{t-1}, then $\xi_{t-1} > 0$ represents an "overvaluation" of Y relative to X. If it is unusually low, then $\xi_{t-1} < 0$. The bivariate ECM for ∇Y_t can be written as

$$\nabla Y_t = \alpha \xi_{t-1} + \sum_{k=0}^{n} \beta_k \nabla X_{t-k} + \sum_{k=1}^{\tilde{n}} \theta_k \nabla Y_{t-k} + \varepsilon_t, \tag{16}$$

where ∇ denotes the first-difference operator. The terms in differences, ∇X_{t-k} and ∇Y_{t-k}, capture short-run responses and persistence, while the coefficient α on ξ_{t-1} captures long-run adjustment. When the cointegrating relation is normalized as $\varepsilon_t = Y_t - \beta X_t$, a negative estimate $\hat{\alpha} < 0$ is consistent with mean reversion toward the equilibrium. This means that if Y_{t-1} is above its equilibrium value (positive disequilibrium), then ∇Y_t tends to be negative, bringing Y_t down to the equilibrium. Conversely, if Y_{t-1} is below equilibrium, then ∇Y_t tends to be positive, restoring equilibrium. The magnitude $|\hat{\alpha}|$ is interpreted as the speed of adjustment, indicating the fraction of the prior-period disequilibrium that is corrected within one period.

The ECM serves as a progression after establishing cointegration. The cointegration test provides evidence of a long-run equilibrium relation in levels, while the ECM explicitly embeds that equilibrium into the model for growth rates. The ECM therefore answers how contemporaneous and lagged changes in X_t transmit to changes in Y_t in the short run via the β_k coefficients, along with whether, and how quickly, deviations from the long-run equilibrium are corrected with the adjustment parameter α. In the context of this paper's overall analysis, the ECM framework allows the short-run predictive results from fitting dynamic regression, VAR models, and running VAR-based tests to be interpreted alongside long-run equilibrium forces implied by the PO test findings.

As with the dynamic regression models, inference in ECM estimation depends on the properties of the residual process ε_t. The Box-Ljung test for residual autocorrelation is used to assess whether OLS standard errors are appropriate. In cases where there is residual autocorrelation, inference is reported using HAC standard errors calculated by Newey-West estimators. These corrections ensure valid statistical conclusions about both the short-run dynamics and the long-run adjustment mechanism in (16).

IV. Empirical Results

A. Unit-Root Test Results

The adjusted HOR, AMV, and respective labor series were tested for the presence of a unit-root using the ADF test defined in (2). The AMV series failed to reject the null hypothesis with the p-value $p = 0.29$ indicating that the series is non-stationary after logarithmic transformation. The HOR series also failed to reject the null hypothesis and the p-value was even larger with $p = 0.7154$. Both labor series following the AMV sample period also reported insignificant p-values indicating that both the employed and unemployed persons datasets are non-stationary. However, running the same ADF test on the logarithmic adjusted employed persons series following the HOR sample period resulted in the test rejecting the null hypothesis, indicating the series is stationary. This was initially unexpected because the data, after additive decomposition, follows a positive trend and does not exhibit any mean reversion over time. On the other hand, the logarithmic adjusted unemployed persons series following the HOR sample period, failed to reject the null hypothesis of the ADF test with the p-value $p = 0.65$ exceeding all standard significance levels.

To further analyze the results in the unit-root testing of the employed persons series following the HOR sample period, both the ERS and PP tests were conducted and compared to the ADF results. Both tests corroborate the initial ADF results with the ERS test statistic being more negative than the critical value at the 1% level, and the PP test strongly rejecting at the 1% level, so it was concluded that the logarithmic adjusted employed persons series following the HOR sample period is stationary. Due to these results, the logarithmic adjusted unemployed persons series following the HOR sample period was tested for stationarity using the PP and ERS tests. The PP test results were consistent with the ADF findings while the ERS test rejects the null hypothesis of a unit-root at the 1% level. Since the ERS test has higher finite-sample power than ADF or PP and removes deterministic components using a generalized least squares (GLS) detrending step, both labor series under the HOR sample period were determined as stationary series over time.

This concern regarding inconsistent unit-root test results due to sample size limitations was exercised for the HOR series, AMV, and labor series under the AMV sample period. The ERS test failed to reject the null hypothesis for each of these series while the PP test gave mixed results. Based on these findings it was concluded that both tourism indicators are non-stationary after logarithmic adjustment along with both logarithmic adjusted labor series under the AMV sample period.

B. *Cointegration Test Results*

Since the AMV series and the respective labor series are non-stationary, cointegration analysis is applied to determine if a stable long-run equilibrium exists between the variables. Let V_t be defined as the logarithmic adjusted AMV series, U_t be the logarithmic adjusted unemployed persons series with the AMV sample period, and L_t be the logarithmic adjusted employed persons series with the same sample period. The PO test rejects the null hypothesis at the 10% level for V_t and U_t with the p-value of the test being $p = 0.0598$. To analyze the significance of the cointegration between V_t and U_t, the power of the test was calculated by estimating the cointegration regression (3), fitting an AR(p) model to the residual series, and running simulations using the observed estimate $\hat{\beta}$ and the fitted residual process. It was found that the PO test has moderate power (0.676) but is conservative in size (0.036 at the 5% level). These results imply that $p = 0.0598$ can occur even when cointegration is present and that failure to reject at 5% should not be taken as strong evidence against a long-run equilibrium relationship.

On the other hand, the PO test strongly rejects the null hypothesis at the 1% level for V_t and L_t with the p-value of the test being $p < 0.01$. This result strongly indicates there exists a long-run equilibrium relationship between V_t and L_t. Both initial cointegration results will be further analyzed using error correction models;however, notably there appears to be a long-run equilibrium relationship between both labor series and AMV including the COVID period.

C. *First Differencing and Cross-Correlation Analysis*

The cointegration results motivate further long-run analysis of AMV and the labor series; however, short-run dynamics between both tourism series and the labor series must also be considered to describe the timing or adjustment path through which deviations are corrected. Additionally, the short-run dynamics between HOR and the labor series can be examined since the testing for cross-correlation and modeling dynamic regression do not necessitate the series be non-stationary. Prior to analyzing the data using cross-correlation the series must all be stationary which requires further adjustments to the data and subsequent unit-root tests to confirm stationarity. The first differences of the log-adjusted AMV, HOR, and labor series are used so that the variables are in comparable units and so that all the series are stationary. The adjusted AMV and respective labor series are defined as

$$v_t = \nabla V_t, \qquad u_t = \nabla U_t, \qquad \ell_t = \nabla \tilde{L}_t \tag{17}$$

where ∇ is the first difference operator. Similarly, the adjusted HOR and respective labor series are defined as

$$h_t = \nabla H_t, \qquad \tilde{u}_t = \nabla \tilde{U}_t, \qquad \tilde{\ell}_t = \nabla \tilde{L}_t \tag{18}$$

where H_t represents the logarithmic adjusted HOR series, $\tilde{L}_t$ represents the logarithmic adjusted employed persons series following the HOR sample period, and $\tilde{U}_t$ represent the logarithmic adjusted unemployed persons series following the same sample period. After first differencing each of the series, the ADF, PP, and ERS tests were run on each individual series and in all cases the tests strongly rejected the null hypothesis at the 1% significance level. Both $\tilde{u}_t$ and u_t are interpreted as the growth rate of unemployment, while $\tilde{\ell}_t$ and ℓ_t are interpreted as the growth rate of employment. The first difference of the log adjusted HOR series, h_t, is best understood as the seasonality adjusted growth rate of hotel occupancy, and v_t is understood as the growth rate of downtown visitor activity.

The cross-correlation analysis between v_{t+k} and u_t yielded strongly significant negative contemporaneous correlation (-0.735) and significant negative correlation between v_{t-1} and u_t (-0.312) indicating that growth in downtown visitor activity leads declines in unemployment. The strong contemporaneous correlation suggests that visitor growth and unemployment declines often occur within the same month, consistent with common shocks in the short run. In the case of v_{t+k} and ℓ_t there is no significant leading or lagging effect, however, there is a strong observed contemporaneous correlation (0.806) indicating that changes in visitor activity and employment growth move together within the same month, suggesting either a rapid within-month adjustment or the influence of common shocks simultaneously affecting both series.

In comparison, the cross-correlation analysis between h_{t+k} and $\tilde{\ell}_t$ yielded significant contemporaneous correlation (0.433) and stronger correlation between h_{t-1} and $\tilde{\ell}_t$ (0.555). This leading and contemporaneous relation was also seen in the analysis of h_{t+k} and $\tilde{u}_t$ where the contemporaneous correlation was significantly negative (-0.458) and the first lag h_{t-1} was significantly correlated with $\tilde{u}_t$ (-0.592). However, in both cases all other cross-correlation values fell below the significance levels. Overall, these cross correlation results for the adjusted HOR series indicate that changes in hotel occupancy tend to move in the same month as changes in employment and unemployment. Additionally, hotel occupancy changes lead changes in employment and unemployment with higher occupancy growth preceding higher employment growth and lower unemployment growth. Each of the four tests lacked significance at other lags suggesting the relationship between the tourism indicators and labor series are short-lived and concentrated at the contemporaneous and first-lag periods rather than spread across many months. However, the magnitudes of these correlations raises further questions about how much of the variation in the respective labor series can be explained by changes in visitation and hotel occupancy.

D. Dynamic Regression Models

The CCF results indicate strong within-month co-movement between each of the tourism and labor series. However, for h_{t+k} the largest cross-correlation occurs at lag $k = -1$. Consistent with the CCF results, the first step for the dynamic regression analysis is fitting a contemporaneous model in which the labor series is regressed on v_t, and fitting a lagged model where the labor series is regressed on h_{t-1}. These models are fit using (9)

$$u_t = \alpha_u + \beta_u v_t + w_t^{(u)} \qquad \ell_t = \alpha_\ell + \beta_\ell v_t + w_t^{(\ell)} \tag{19a}$$

$$\tilde{u}_t = \tilde{\alpha}_u + \beta'_u h_{t-1} + \tilde{w}_t^{(u)} \qquad \tilde{\ell}_t = \tilde{\alpha}_\ell + \beta'_\ell h_{t-1} + \tilde{w}_t^{(\ell)} \tag{19b}$$

and the value of the estimated slope coefficients $\hat{\beta}_u$ and $\hat{\beta}_\ell$ summarize the marginal association between v_t and u_t, and between v_t and ℓ_t, respectively. Similarly, the value of the estimated slope coefficients $\hat{\beta'}_u$ and $\hat{\beta'}_\ell$ summarize the leading association between h_{t-1} and $\tilde{u}_t$, and between h_{t-1} and $\tilde{\ell}_t$, respectively. Fitting the models defined in (19a) yielded statistically significant results for both $\hat{\beta}_u$ ($p < 0.001$) and $\hat{\beta}_\ell$ ($p < 0.001$).

$$u_t = -0.5434 v_t + w_t^{(u)} \qquad \ell_t = 0.109 v_t + w_t^{(\ell)} \tag{20}$$

The intercepts α_u and α_ℓ were not significant and near zero so they have been omitted from the results. Since the Box-Ljung test indicates residual autocorrelation, coefficient inference is based on HAC standard errors calculated using Newey-West estimators. From these initial results, it's clear that higher visitor growth is significantly associated with higher employment growth and declining unemployment in the same period, reaffirming the results from the CCF. The explanatory power of the ℓ_t model is $R^2 = 0.65$ while the explanatory power of the u_t model is $R^2 = 0.54$. This indicates that within-month growth of downtown visitors explains about 65% of the variability in the employment growth rate and 54% of the variability in the unemployment growth rate. Economically, these results also indicate that a 1% growth in visitors is associated with about a 0.11% increase in the employment growth rate and a 0.54% decrease in the unemployment growth rate within the same month (see Figures 1 and 2 for plots of the fitted regression lines against the observed visitor growth data). Fitting the models defined in (19b) yielded similar estimates for both $\hat{\beta'}_u$ ($p < 0.001$) and $\hat{\beta'}_\ell$ ($p < 0.001$).

$$\tilde{u}_t = -0.6747 h_{t-1} + \tilde{w}_t^{(u)} \qquad \tilde{\ell}_t = 0.1222 h_{t-1} + \tilde{w}_t^{(\ell)} \tag{21}$$

Again, the intercepts $\tilde{\alpha}_u$ and $\tilde{\alpha}_\ell$ were found to not be statistically significant and close to zero. However, the Box-Ljung test failed to reject the null hypothesis for both model specifications with the p-value far exceeding standard significant levels in both cases ($p \geq 0.5$) indicating there is no evidence of residual serial correlation and standard inference is appropriate. These dynamic regression results strongly indicate that growth in hotel occupancy precedes declines in unemployment growth and higher employment growth. The explanatory power of the h_{t-1} models is weaker than the visitation models where $R^2 = 0.31$ in the $\tilde{\ell}_t$ model and $R^2 = 0.35$ in the $\tilde{u}_t$ model. This difference in explanatory power is likely due to the exclusion of the contemporaneous occupancy term, as the model solely relies on lagged data; nonetheless, growth in hotel occupancy acts as a significant leading indicator for changes in both unemployment and employment growth. Economically, these results also indicate that a 1% growth in the previous month's hotel occupancy is associated with a 0.12% increase in the current month's employment growth rate and a 0.68% decrease in the current month's unemployment growth rate (refer to Figures 3 and 4 for visual representations of these lagged relationships).

To account for the contemporaneous relation exhibited in the CCF results, the term h_t is included in the (19b) models and the results were statistically significant ($p < 0.001$ for all terms) and similar to the original estimations of $\hat{\beta'}_u$ and $\hat{\beta'}_\ell$.

$$\tilde{u}_t = -0.4752 h_t - 0.6415 h_{t-1} + \tilde{w}_t^{(u)} \tag{22}$$

$$\tilde{\ell}_t = 0.0912 h_t + 0.1158 h_{t-1} + \tilde{w}_t^{(\ell)} \tag{23}$$

Consistent with the initial dynamic regression model, the Box-Ljung test failed to reject the null hypothesis for both model specifications with the p-value for testing the residual serial correlation in $\tilde{u}_t$ being 0.87 and the p-value for the test on $\tilde{\ell}_t$ being 0.35. Both results for the new $\tilde{u}_t$ and $\tilde{\ell}_t$ models indicate that hotel occupancy growth has significant predictive power on current unemployment and employment growth. In both cases the lagged effect is larger than the contemporaneous relation, but including the contemporaneous effect increases the explanatory power of each model. For both $\tilde{u}_t$ and $\tilde{\ell}_t$ the R^2 value was found to be around 0.5 meaning the specified models for employment and unemployment explain about 50% of the variability in each of the labor growth rates. The total economic effect over the two periods indicates that a 1% spike in the hotel occupancy growth rate increases the employment growth rate by 0.207% and decreases the unemployment growth rate by 1.117% (see Figures 7 and 8 for the fitted values of the combined contemporaneous and lagged hotel occupancy models).

The cross-correlation analysis for v_{t+k} and the labor series also indicated some lag relation between v_{t-1} and u_t, so the dynamic regression model was refit including the lagged visitors term v_{t-1} in the (19a) model specifications. The Box-Ljung test rejected the null hypothesis of no residual serial correlation, so as with the initial regression, coefficient inference is based on HAC standard errors calculated using Newey-West estimators.

$$u_t = -0.5273 v_t - 0.1856 v_{t-1} + w_t^{(u)} \tag{24}$$

$$\ell_t = 0.1079 v_t + 0.0128 v_{t-1} + w_t^{(\ell)} \tag{25}$$

The estimated coefficient of the lagged visitor growth term in the u_t model is statistically significant ($p < 0.001$; however, in the fitted ℓ_t model the p-value of the estimated coefficient for the lagged visitor growth term exceeded the 10% significance level. The explanatory power of the u_t model increased ($R^2 = 0.60$) suggesting that including the previous periods visitor growth helps in further explaining the variability in the unemployment growth rate. The total economic effect over the two periods indicates that a 1% change in the downtown visitor growth rate decreases the unemployment growth rate by 0.713% (refer to Figures 5 and 6 for the fitted values of the models incorporating the lagged visitor growth term). On the other hand, since the estimated coefficient in the ℓ_t model was not statistically significant, the economic impact of visitor growth on employment growth is characterized as strictly contemporaneous in the dynamic regression model specification.

In total, the dynamic regression results confirm and clarify the results of the CCF analysis. From each of the dynamic regression models, hotel occupancy acts as a significant indicator for within-month and near-term changes in both employment and unemployment growth. Meanwhile, while visitation growth in both the current and prior periods strongly predicts near-term changes in unemployment, but its predictive power for employment is limited to the current month. To further analyze the short-run dynamics of the tourism indicators and Detroit-labor series the VAR models are fit and causality tests are performed to formally distinguish between lagged predictive relationships and instantaneous co-movement.

E. Vector Autoregressive Models and Causality Test Results

Given the distinct structural behaviors identified in the single-equation analysis, the VAR framework was applied to formally model the endogenous interactions between the variables. The VAR(1) model was initially fit for v_t and ℓ_t and the findings reflected contrasting results to what was observed in the dynamic regression.

$$v_t = 0.3150v_{t-1} - 2.1023\ell_{t-1} + w_{v,t} \tag{26}$$

$$\ell_t = 0.1046v_{t-1} - 0.7574\ell_{t-1} + w_{\ell,t} \tag{27}$$

Since the VAR(1) estimates both v_t and ℓ_t simultaneously, each series is treated as endogenous allowing for a two-way analysis of the visitor and employment growth rates. For the visitor growth rate equation the p-values for each of the terms were statistically insignificant ($p > 0.10$), demonstrating that visitor growth is neither significantly predicted by its own past values nor influenced by lagged employment growth. The large standard deviations associated with these estimates confirm this lack of significance.

In comparison, the employment equation yields robust results with $R^2 = 0.23$ and both coefficients being significant at the 1% level. The lagged employment growth term, ℓ_{t-1}, is negative indicating that employment growth is mean-reverting, oscillating over time likely due to macroeconomic noise and some seasonality. In contrast to the dynamic regression results, the predictive power of visitor growth is greater in magnitude suggesting that visitor growth acts as a robust leading indicator for employment growth in the following period. The standard deviations of the estimated coefficients are small relative to the estimated values reinforcing the strength of the observed relationship between visitation and employment growth. While the model explains approximately 23% of the variance only using lagged values, the strong positive correlation between the residuals of the visitor and employment growth equations ($\mathrm{cor}(w_{v,t}, w_{\ell,t}) = 0.83$) indicates a high degree of contemporaneous co-movement, reinforcing the previously observed contemporaneous correlation between visitation and employment growth.

The dynamic regression results for visitation and unemployment growth suggest a strong negative relationship where visitation growth indicates lower unemployment growth; however, the regression model does not account for persistence in the unemployment growth series which could obscure the real effect of visitation growth. The VAR(1) model was fit for u_t and v_t with the results for the u_t model generally supporting the initial findings in the dynamic regression analysis.

$$v_t = 0.4418v_{t-1} + 0.6552u_{t-1} + w_{v,t} \tag{28}$$

$$u_t = -0.3448v_{t-1} - 0.2098u_{t-1} + w_{u,t} \tag{29}$$

First looking at the visitor growth equation, while both estimated coefficients are statistically significant ($p < 0.01$), this result diverges from the previous fitted VAR, where visitor growth behaved as an exogenous process. Furthermore, the positive estimated coefficient for the lagged unemployment term implies a counterintuitive dynamic where rising unemployment predicts tourism growth. This result sharply contrasts with the previously discussed literature and suggests this specific correlation may be idiosyncratic to the unemployment series rather than a structural feature of the tourism economy.

On the other hand, the unemployment growth equation demonstrates a relationship consistent with the literature and previous findings. Both estimated coefficients are negative confirming that growth in the tourism sector acts as a leading indicator for declines in unemployment. However, the estimated coefficient for the lagged unemployment term is not statistically significant ($p > 0.10$) with a large standard deviation relative to the coefficient estimate. This lack of significance in the autoregressive term indicates that, unlike employment, the unemployment growth rate does not exhibit strong internal persistence or mean-reversion when modeled with visitation growth. While the model explains approximately 12% of the variance using lagged values ($R^2 = 0.12$), the residuals of the v_t and u_t equations exhibit a strong negative correlation ($\mathrm{cor}(w_{v,t}, w_{u,t}) = -0.749$) suggesting strong contemporaneous co-movement. These results reinforce the findings of the CCF analysis, suggesting that while the lagged predictive power is significant, the primary mechanism connecting visitor growth to unemployment declines occurs within the same month.

Further testing the two-way short-run relationship between tourism growth and the labor series, the VAR(1) model was fit for the hotel occupancy and respective employment growth rate series.

$$h_t = 0.0492h_{t-1} + 0.2174\tilde{\ell}_{t-1} + w_{h,t} \quad (30)$$

$$\tilde{\ell}_t = 0.1659h_{t-1} - 0.4588\tilde{\ell}_{t-1} + w_{\tilde{\ell},t} \quad (31)$$

Consistent with the visitation findings, growth in hotel occupancy recovery is primarily driven by external factors rather than employment growth or its own past values. Both estimated coefficients for the lagged hotel occupancy and employment growth terms exceed the 10% significance level ($p = 0.738, p = 0.745$) and the R^2 value for the equation is near zero, solidifying the conclusion that neither its own past values nor employment growth have predictive power for current hotel occupancy growth.

The employment growth equation reinforces previous empirical findings. First, the findings are consistent with dynamic regression and CCF analysis which suggested that hotel occupancy acts as a significant leading indicator for future employment growth. Second, the negative lagged employment growth term is consistent with the observations made in the previous ℓ_t equation. Overall, the results are statistically significant ($p < 0.001$) and the explanatory power is more than double that of the employment-visitor model ($R^2 = 0.48$) suggesting that hotel occupancy growth may be a stronger predictor of employment growth rather than overall growth in visitation. Additionally, the correlation between the residuals ($\text{cor}(w_{h,t}, w_{\tilde{\ell},t}) = 0.5862$) is moderately strong, confirming that while lagged predictive power is strong, there is also positive co-movement between both series.

Since hotel occupancy growth, even after accounting for mean-reversion in the employment series, demonstrated strong predictive power for employment growth, the VAR(1) model was fit for the hotel occupancy and respective unemployment growth rate series.

$$h_t = 0.1050h_{t-1} + 0.0673\tilde{u}_{t-1} + w_{h,t} \quad (32)$$

$$\tilde{u}_t = -0.7263h_{t-1} - 0.0990\tilde{u}_{t-1} + w_{\tilde{u},t} \quad (33)$$

The fitted h_t equation does not exhibit the idiosyncratic behavior of the visitation and unemployment growth model. Instead the results are consistent with the previously observed h_t equation, where neither the autoregressive term nor past unemployment growth have any predictive power for current hotel occupancy growth ($R^2 = 0.01$, $p = 0.48$, $p = 0.61$).

The fitted unemployment growth equation, on the other hand, provides the strongest evidence yet for tourism's role as a leading indicator. The estimated coefficient for lagged hotel occupancy growth is negative and large in magnitude, and statistically significant at the 0.1% level. This indicates that a 1% growth in hotel occupancy is associated with a 0.73% decline in the unemployment growth rate in the following month. Furthermore, consistent with the previous specifications, the autoregressive term for unemployment is insignificant ($p = 0.417$), demonstrating that unemployment growth does not exhibit internal persistence when modeled with hotel occupancy growth. The model explains approximately 36% of the variance ($R^2 = 0.36$) which is three times that of the visitor-unemployment model. In addition, the correlation between the residuals ($\text{cor}(w_{h,t}, w_{\tilde{u},t}) = -0.5067$) is similar in magnitude to that of the correlation found between the hotel-employment residuals suggesting that both the lagged predictive power is strong along with the co-movement between the series. While these structural estimates suggest clear leading relationships and contemporaneous co-movement, formal Granger causality and instantaneous causality tests are used to rigorously confirm the directionality and statistical validity of these short-run interactions.

For visitor and employment growth, uni-directional causality was identified with strong evidence for visitor growth Granger-causing employment growth. The null hypothesis of no Granger causality was rejected ($p < 0.01$), implying that past values of visitor growth significantly improve the prediction of employment growth. Conversely, the test failed to reject the null hypothesis that employment growth does not Granger-cause visitor growth ($p > 0.1$). This indicates that visitor growth acts as an exogenous shock to the labor series, where past movements in visitors help predict future employment trends. Regarding visitor and unemployment growth, the Granger causality tests initially indicate a bi-directional causal relation, with the null hypothesis being rejected in both directions ($p < 0.01$). However, given the positive coefficient observed in the VAR equation, where an increase in unemployment growth appeared to predict an increase in visitor growth, this specific directionality likely reflects idiosyncratic features of the unemployment series rather than a robust structural mechanism. So, despite the bi-directional test results, the overall evidence aligns more strongly with the established literature, supporting the conclusion that visitor growth is the primary driver of declines in unemployment rather than the reverse. In both cases, the instantaneous causality test strongly rejected the null hypothesis ($p < 0.01$) formally verifying the significant contemporaneous co-movement observed in the VAR model residuals.

Next, the causality analysis for the hotel occupancy models provided definitive evidence of directionality. For both employment and unemployment growth, strict uni-directional causality was determined with hotel occupancy growth Granger-causing both labor series. The evidence from the Granger causality test was substantially stronger than in the visitor VAR model. For example, the null hypothesis that hotel occupancy growth does not Granger-cause employment growth was rejected with extreme significance ($p < 2.4 \times 10^{-10}$). Crucially, when testing whether changes in each of the labor series Granger-cause hotel occupancy growth the test failed to reject the null hypothesis for both variables ($p > 0.6$). The absence of feedback from unemployment to hotel occupancy growth resolves the ambiguity found in the previous model, substantiating the claim that the earlier bi-directionality was an artifact of the general visitor data. In total, past values of hotel occupancy growth are robust predictors of changes in unemployment and employment, while the labor series themselves do not predict changes in hotel occupancy. Finally, the instantaneous causality tests again rejected the null hypothesis ($p < 0.01$), confirming that shocks to hotel occupancy growth impact unemployment and employment growth simultaneously within the current month, as well as in subsequent periods.

F. Error-Correction Models

The empirical analysis has so far established two distinct relationships between the tourism and labor series. First, the cointegration results confirmed that logarithmic adjusted visitor series shares a long-run equilibrium with both the logarithmic adjusted employment and unemployment series. Second, utilizing the CCF analysis, dynamic regression, and VAR models, there are established short-run dynamics and causal links that drive monthly changes in both employment and unemployment growth rates. To further analyze the long-run relationship between downtown visitation and the Detroit-labor series ECMs are estimated for both employment and unemployment. This approach follows the PO cointegration test results, which imply that a valid error correction representation exists for both the employment and unemployment variables. This long-run analysis estimates the rate at which the labor series return to equilibrium following a shock, while controlling for the short-run dynamics captured in the VAR and regression analyses.

The results for the employment ECM demonstrate a robust mechanism of adjustment and significant short-run sensitivity to changes in visitor growth. The error correction term, ξ_{t-1} is statistically significant at the 1% level with a coefficient of -0.2929. This implies that approximately 29% of any disequilibrium between employment and visitation is corrected in the following month, ensuring the system reverts to its long-run trend. The estimated coefficients of the contemporaneous and lagged visitor growth terms, v_t and v_{t-1}, are statistically significant and describe the short-run dynamics of the model. The full model is expressed as

$$\ell_t = -0.2929\xi_{t-1} + 0.101v_t + 0.054v_{t-1} - 0.3996\ell_{t-1} + \epsilon_t \tag{34}$$

with the estimated coefficients for both visitor growth terms indicating that a 1% shock in visitor growth raises employment by 0.10% in the same month, with a persistent positive effect carrying over into the next month. The model explains a substantial portion of the variance in employment changes, indicated by an adjusted R^2 value of 0.77. Notably, diagnostic testing using the Box-Ljung test rejected the null hypothesis of no autocorrelation in the residuals. To address this and ensure the results are robust, HAC standard errors were calculated using Newey-West estimators. The estimated standard errors reaffirm that the coefficients for the error correction term and visitor variables are statistically significant at the 1% level, confirming the reliability of the estimates despite the residual structure.

Similarly, the results of the unemployment ECM indicate a stable long-run relationship, although, with distinct short-run characteristics compared to the employment series. The error correction term is statistically significant at the 1% level with a coefficient of -0.2067, confirming that unemployment also follows a mean-reverting process, with approximately 21% of any deviation from the long-run equilibrium being corrected in the subsequent month. While this result verifies the validity of the cointegration, the rate of adjustment is slightly slower than that observed in the employment model, suggesting that the unemployment level takes longer to revert to the equilibrium following a shock or change.

In the short run, the estimated coefficients of contemporaneous visitor growth is larger than in the fitted employment ECM, suggesting that unemployment is more sensitive to changes in visitor growth. The full model is expressed as

$$\begin{aligned} u_t = &-0.2067\xi_{t-1} - 0.5486v_t - 0.0031v_{t-1} \\ &+ 0.2274u_{t-1} + \epsilon_t \end{aligned} \tag{35}$$

with the estimated coefficient of v_t indicating that a 1% increase in visitor growth generates a 0.55% decrease in unemployment in the same month. While the coefficient of the contemporaneous term is highly significant at the 1% level, the lagged visitor growth term is statistically insignificant with $p = 0.97$. This implies that while tourism shocks provide a powerful immediate reduction in unemployment growth, the effect does not actively drive reductions in the following month. The overall fit of the model is strong explaining approximately 64% of the variance in unemployment growth.

V. Conclusions and Future Works

The empirical results establish the short-run and long-run relationship between Detroit's downtown tourism growth and resident labor market. From establishing stationarity and cointegration to modeling dynamic causality, the analysis concludes that tourism acts as a robust, exogenous driver of employment growth and unemployment decline. The analysis establishes two primary mechanisms. First, there exists a stable long-run equilibrium where the labor market corrects deviations to align with visitor growth. Second, there are established short-run dynamics where hotel occupancy acts as a leading indicator while visitation acts as a significant contemporaneous shock.

The distinction between the predictive power of hotel occupancy growth and the immediate impact of visitation growth highlights structural differences in how these distinct tourism measures interact with the local labor series. Both the dynamic regression and VAR models consistently identified hotel occupancy growth as a leading indicator, predicting subsequent period changes in employment growth and unemployment decline. In contrast, visitation's primarily contemporaneous relationship suggests that both employment and unemployment immediately react to fluctuations in downtown visitation. Intuitively, these results reflect the economic context of the variables. Since hotel occupancy is representative of long-term stays, these extended periods of visitation require hospitality reservations and reflect periods of greater consumer expenditure thereby resulting in adaptive hiring of local residents in the near following period. On the other hand, since visitation encapsulates both long-term and short-term visits, the primarily contemporaneous, but in some cases leading, results of the short-run dynamics highlight how the labor market also responds through immediate mechanisms such as gig work or spontaneous hirings.

Beyond the short-run fluctuations, the most significant finding for Detroit's development strategy is the confirmation of cointegration. Despite the extreme volatility of the COVID-19 pandemic, the resident labor series remains linked to changes in visitation. The statistically significant error correction terms in both ECM models confirm that these variables share a long-run equilibrium, validating the claim that investment in attracting visitors to the city raises the structural level of local labor demand. Importantly, the results indicate an asymmetry in this adjustment where employment growth exhibits persistence, or lagged significance, suggesting that jobs created by tourism tend to remain. Whereas, the reduction in unemployment is immediate, but temporary, implying that while tourism is an effective shock for within-month unemployment decline, sustained visitor growth is required to maintain those structural gains.

In total, this paper has demonstrated how Detroit visitation, both long-term and overall, impacts the employment and unemployment of Detroit-city residents. The empirical results illustrate that in the short-run higher overall downtown visitation and downtown hotel occupancy indicate periods of increasing employment and decreasing unemployment, and these results can serve as justification for further expanding tourism-based developments. Given the immediate and predictive effects of visitation and hotel occupancy on the labor market, policy makers may consider more development plans like the ones for the Greater Warren-Conner area and the Brightmoor neighborhood with an emphasis on attracting business investment. Based on the short-run causal analysis, this type of investment could help spur jobs for Detroit residents and immediately help neighborhood communities further develop and keep up, economically, with the downtown area. Over the long-run, developing the neighborhoods from passive recipients of visitors into active visitor destinations could structurally aid local labor demand and ensure that the economic momentum of the city center is effectively transmitted to residents across Detroit. These analytic implications for growing tourism developments across the neighborhoods need to be carefully considered because of concerns like population displacement that could result from gentrification. Before any policy suggestions are made, the possibilities and effects of neighborhood specific developments need to be analyzed to ensure that the structural gains in resident employment are not negated by the displacement of the very communities intended to benefit from this economic growth.

REFERENCES

[1] Aguilar, Louis. LCA contractors fined $2.9M over Detroit hires. *The Detroit News* (August 23, 2017). `https://www.detroitnews.com/story/news/local/detroit-city/2017/08/23/little-caesars-arena-construction-fines/104902524/`

[2] Aguilar, Louis. Detroit Riverfront Conservancy announces more management changes; names Rachel Frierson vice president of operations and programming. *The Detroit News* (July 15, 2024). `https://www.detroitnews.com/story/news/local/detroit-city/2024/07/15/detroit-riverfront-conservancy-rachel-frierson-riverwalk-detroit-operations/74415493007/`

[3] Alcala-Ordonez, Alejandro, and Veronica Segarra. Tourism and economic development: A literature review to highlight main empirical findings. *Tourism Economics* 31.1 (2025): 76-103.

[4] Allard, S. W., M. V. Wathen, and S. K. Danziger. Bundling Public and Charitable Supports to Cope with the Effects of the Great Recession. *Social Science Quarterly* 96.5 (2015): 1348-1362.

[5] Bakker, Martine, Rene van der Duim, Karin Peters, and Jeroen Klomp. Tourism and Inclusive Growth: Evaluating a Diagnostic Framework. *Tourism Planning & Development* 20.3 (2023): 416-439.

[6] Chambers, Jennifer. Construction jobs soar with Little Caesars Arena work. *The Detroit News* (May 2, 2017). `https://www.detroitnews.com/story/business/2017/05/02/detroit-construction-jobs/101201020/`

[7] Chester, Tim, Sarika Bansal, and Betsy Cohen. Where to Go in 2024. *AFAR Media* (November 28, 2023). `https://www.afar.com/magazine/where-to-go-2024`

[8] City of Detroit Mayor's Office. City of Detroit and Rocket Community Fund announce $2.2 million in investment to advance housing stability network in Detroit. News release, May 6, 2021, `https://detroitmi.gov/news/city-detroit-and-rocket-community-fund-announce-22-million-investment-advance-housing-stability`

[9] City of Detroit Mayor's Office. Special Events. City of Detroit, `https://detroitmi.gov/government/mayors-office/special-events`

[10] City of Detroit Planning and Development Department. *Greater Warren/Conner Neighborhood Framework Plan: Executive Summary.* City of Detroit, May 2024, `https://detroitmi.gov/`

[11] City of Detroit Planning & Development Department. *Detroit East Riverfront Framework Plan: Report Appendix.* City of Detroit, April 14, 2017.

[12] City of Detroit Planning and Development Department. *Brightmoor Economic Development Focus Group.* Slide deck, City of Detroit Planning and Development Department, April 2024, `https://detroitmi.gov/brightmoor`

[13] City of Detroit Planning and Development Department. District Detroit — Community Benefits Ordinance: Past CBO Engagements. City of Detroit, `https://detroitmi.gov/departments/planning-and-development-department/community-benefits-ordinance/past-cbo-engagements/district-detroit`

[14] City of Detroit, Michigan. *Plan for the Adjustment of Debts of the City of Detroit.* United States Bankruptcy Court, Eastern District of Michigan, Case No. 13-53846, Doc. 2708, February 21, 2014.

[15] CNT Editors. The Best Places to Go in North America & the Caribbean in 2024. *Condé Nast Traveler* (November 15, 2023). `https://www.cntraveler.com/story/best-places-to-go-in-north-america-and-the-caribbean-2024`

[16] Community Foundation for Southeast Michigan. Community Foundation Secures $35M Fund to Complete Detroit Riverfront Revitalization. Press release, June 24, 2024, `https://cfsem.org/community-foundation-secures-35m-fund-to-complete-detroit-riverfront-revitalization/`

[17] Detroit Metro Convention & Visitors Bureau. Metro Detroit's Hospitality Industry Employment Hits Record Highs Before COVID-19. Press release, July 7, 2020, `https://visitdetroit.com/media/press-releases/metro-detroits-hospitality-industry-employment-hits-record-highs-before-covid-19/`

[18] Detroit Riverfront Conservancy. Jobs. Detroit Riverfront Conservancy, `https://www.detroitriverfront.org/inside-the-conservancy/jobs`

[19] Downtown Detroit Partnership. *Downtown Detroit Dashboard.* Downtown Detroit Partnership, 2025, https://downtowndetroit.org/downtown-dashboard/
[20] Downtown Detroit Partnership. *Events.* Downtown Detroit Partnership, 2025, https://downtowndetroit.org/events/
[21] Downtown Detroit Partnership. *Urban Innovation.* Downtown Detroit Partnership, 2024, https://downtowndetroit.org/about-the-ddp/impact/program-initiatives/urban-innovation/
[22] Duggan, Michael E. *Executive Order No. 2025-02: Utilization of Detroit Residents on Publicly-Funded Construction and Demolition/Rehab Projects.* City of Detroit, Office of the Mayor, July 2024.
[23] Ernsthausen, Jeff, and Justin Elliott. How a Tax Break to Help the Poor Went to NBA Owner Dan Gilbert. *ProPublica* (October 24, 2019). https://www.propublica.org/article/how-a-tax-break-to-help-the-poor-went-to-nba-owner-dan-gilbert
[24] Garidzirai, Rufaro, and Bulelwa Nguza-Mduba. Does tourism contribute to local economic development (LED) in the City of Cape Town Municipality? A time series analysis. *African Journal of Hospitality, Tourism and Leisure* 9.1 (2020): 1-13.
[25] Hartig, John. *From Cleanup of the Detroit River to Revitalization of the Waterfront.* Case study, International Association for Great Lakes Research, 2018.
[26] Lage, L. 2024 NFL Draft beats attendance record with 700,000 fans in Detroit. *AP News* (April 27, 2024). https://apnews.com/article/nfl-draft-detroit-record-b475546a37aad54c48c4199cc275971f
[27] Li, Meiyi, Yanmei Fan, Cheng Guo, and Xiaokun Li. Tourism prosperity and high-quality economic development. *International Review of Economics & Finance* 101 (2025): 104246.
[28] Llorca-Rodriguez, Carmen Maria, Amalia Cristina Casas-Jurado, and Rosa Maria Garcia-Fernandez. Tourism and poverty alleviation: An empirical analysis using panel data on Peru's departments. *International Journal of Tourism Research* 19.6 (2017): 746-756.
[29] McDonald, John F. What happened to and in Detroit? *Urban Studies* 51.16 (2014): 3309-3329.
[30] Meyer, Daniel Francois, and Natanya Meyer. The Relationship between the Tourism Sector and Local Economic Development (LED): The Case of the Vaal Triangle Region, South Africa. *Journal of Environmental Management and Tourism* 7.3 (15) (2016): 466-472.
[31] Moskowitz, Peter. The two Detroits: a city both collapsing and gentrifying at the same time. *The Guardian — Cities* (February 5, 2015). https://www.theguardian.com/cities/2015/feb/05/detroit-city-collapsing-gentrifying
[32] Motor City Electric Co. Little Caesars Arena — Detroit. Motor City Electric Co., https://www.mceco.com/projects/little-caesars-arena-detroit/
[33] Naglick, John, Martin Lavelle, and John A. Moore. "The Grand Bargain": Detroit's Financial Fall to Bankruptcy and Rise to New Possibilities. *Essays in Economic & Business History* 38 (2020): 310-334.
[34] Noble, Breana. Google to open Detroit office near Little Caesars Arena. *The Detroit News* (November 8, 2018). https://www.detroitnews.com/story/business/2018/11/08/google-detroit-office-little-caesars-arena/1852102002/
[35] Owens, Raymond, III, Esteban Rossi-Hansberg, and Pierre-Daniel Sarte. Rethinking Detroit. *American Economic Journal: Economic Policy* 12.2 (2020): 258-305.
[36] Phillips, Richard. Reviving a mortgage market through financial inclusion? Experimental housing governance and alternative home loan programmes in Detroit, Michigan. *International Journal of Housing Policy* 21.4 (2021): 559-591.
[37] Rahal, Sarah. The inside story: Everything Detroit did to prepare for the NFL Draft. *The Detroit News* (April 25, 2024). https://www.detroitnews.com/story/news/local/detroit-city/2024/04/25/the-inside-story-everything-detroit-did-to-prepare-for-the-nfl-draft/73414495007/
[38] Rahal, Nour. Detroit gets $19.8M from Knight Foundation to boost art, technology and community projects. *Detroit Free Press* (September 8, 2025). https://www.freep.com/story/news/local/michigan/detroit/2025/09/08/knight-foundation-detroit-neighborhoods/85957224007/
[39] Skidmore, Owings & Merrill (SOM). Detroit East Riverfront Framework Plan. Skidmore, Owings & Merrill, https://www.som.com/projects/detroit-east-riverfront-framework-plan/
[40] State of Michigan DTMB. *Demographic and Labor Market Profile: Detroit City.* State of Michigan, Department of Technology, Management and Budget, April 2015, https://milmi.org/Portals/198/publications/Detroit_City_Demographic_and_Labor_Mkt_Profile.pdf
[41] Sugrue, Thomas J. *The Origins of the Urban Crisis: Race and Inequality in Postwar Detroit.* Updated ed., Princeton University Press, 2014.
[42] Tourism Economics. Tourism Economic Impact — Region and County. Michigan.org — Industry: Research & Reports, https://www.michigan.org/industry/researchandreports
[43] Visit Detroit. *About Us.* Detroit Metro Convention & Visitors Bureau, https://visitdetroit.com/about-us/
[44] Visit Detroit. *The Guide to Metro Detroit Major Annual Events.* Detroit Metro Convention & Visitors Bureau, 2025, https://visitdetroit.com/events/annual-events/
[45] Visit Detroit. PASSION MEETS INNOVATION: Visit Detroit to Serve as Juventus Front-of-Shirt Partner. Press release, May 12, 2025, https://visitdetroit.com/media/press-releases/passion-meets-innovation-visit-detroit-to-serve-as-juventus-front-of-shirt-partner/
[46] Visit Detroit. Partnership. Detroit Metro Convention & Visitors Bureau, https://visitdetroit.com/membership/
[47] Williams, Candice. Hotel construction boom drives Detroit's push to attract major events. *Signature Associates — Insights* (March 26, 2025). https://www.signatureassociates.com/hotel-construction-boom-drives-detroits-push-to-attract-major-events/

VI. Appendix

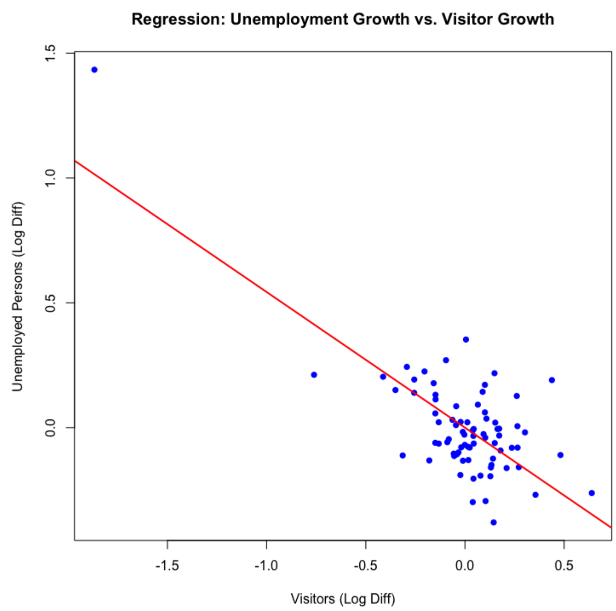

Fig. 1: Unemployment vs. Visitation Growth Scatter Plot and Fitted Line via Dynamic Regression

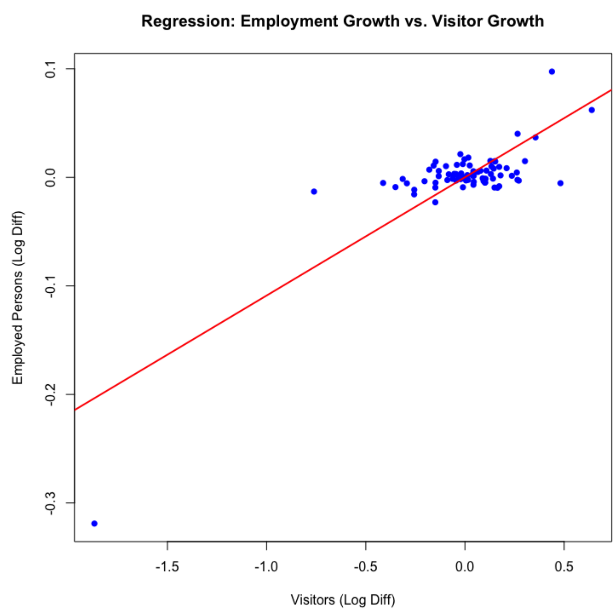

Fig. 2: Employment vs. Visitation Growth Scatter Plot and Fitted Line via Dynamic Regression

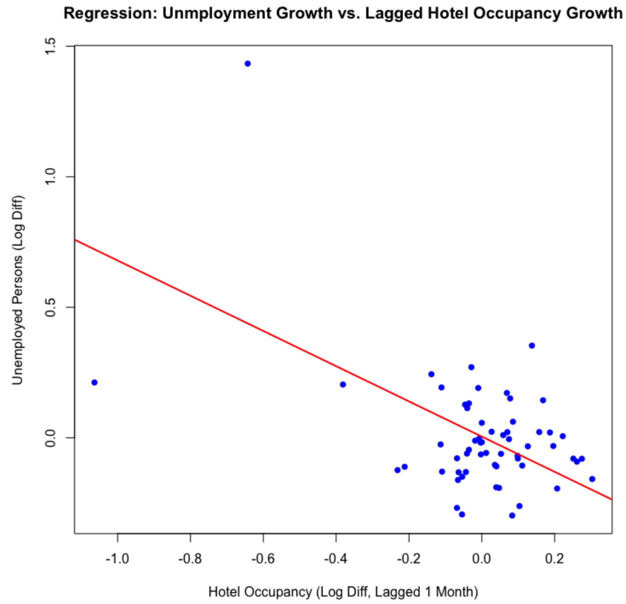

Fig. 3: Unemployment vs. Hotel Occupancy Growth Scatter Plot and Fitted Line via Dynamic Regression

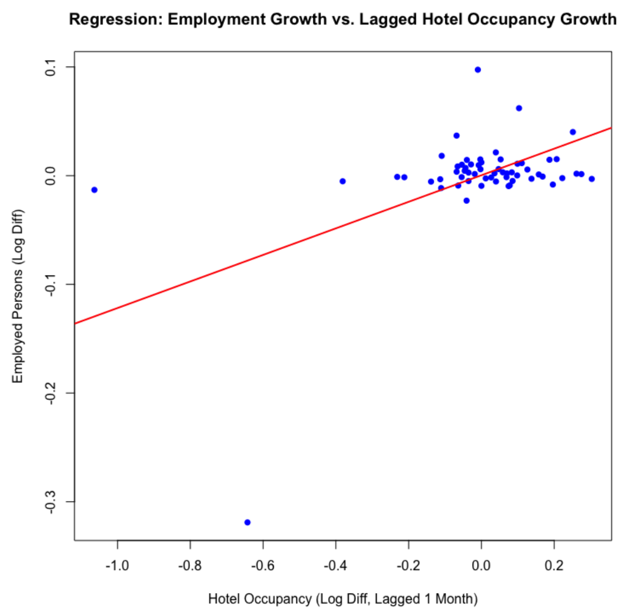

Fig. 4: Employment vs. Hotel Occupancy Growth Scatter Plot and Fitted Line via Dynamic Regression

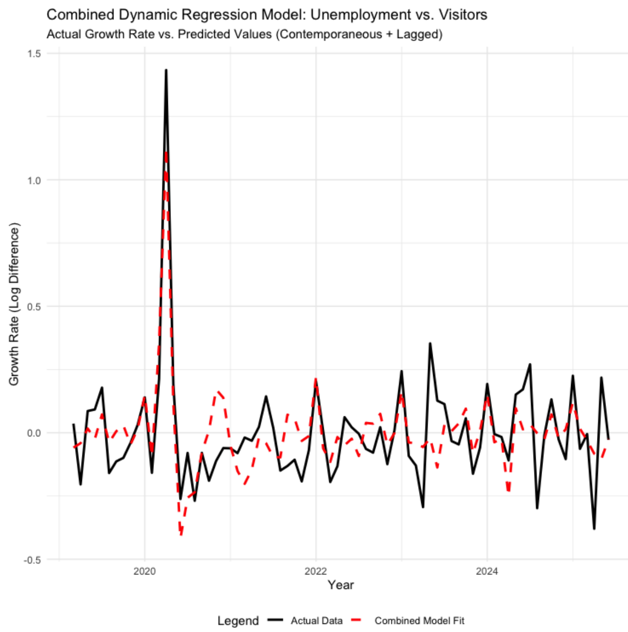

Fig. 5: Fitted values of the dynamic regression with both contemporaneous and lagged downtown visitor growth terms to predict changes in the unemployment growth rate. *Note.* The solid black line represents the actual first difference of the logarithmic adjusted unemployed persons series, while the dashed line represents the fitted values from the dynamic regression.

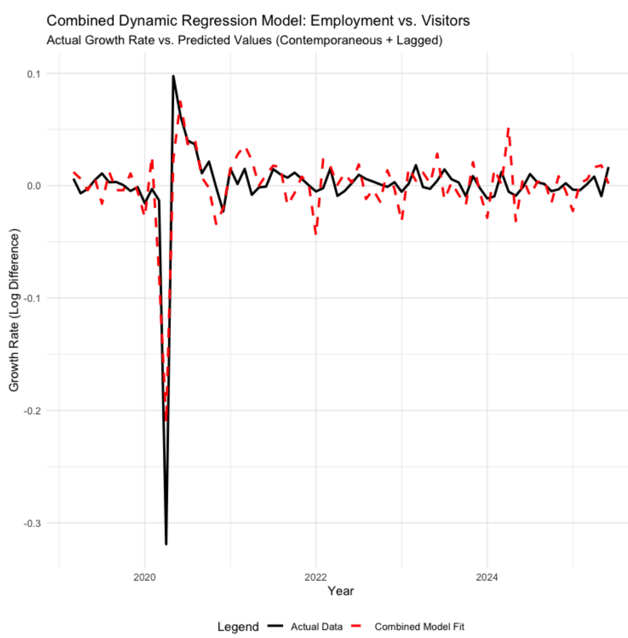

Fig. 6: Fitted values of the dynamic regression with both contemporaneous and lagged downtown visitor growth terms to predict changes in the employment growth rate. *Note.* The solid black line represents the actual first difference of the logarithmic adjusted employed persons series, while the dashed line represents the fitted values from the dynamic regression.

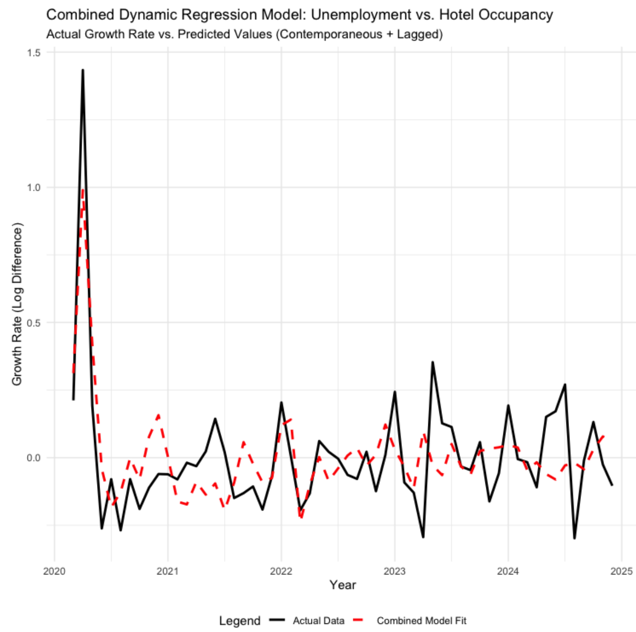

Fig. 7: Fitted values of the dynamic regression with both contemporaneous and lagged hotel occupancy growth terms to predict changes in the unemployment growth rate. *Note.* The solid black line represents the actual first difference of the logarithmic adjusted unemployed persons series, while the dashed line represents the fitted values from the dynamic regression.

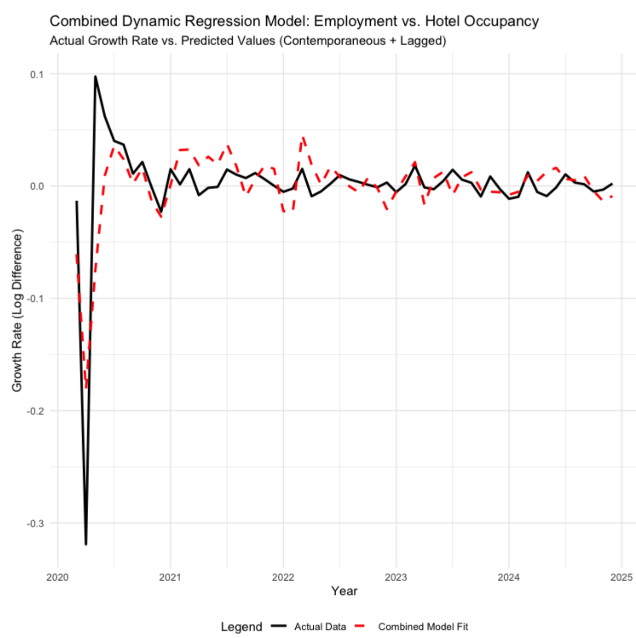

Fig. 8: Fitted values of the dynamic regression with both contemporaneous and lagged hotel occupancy growth terms to predict change in the employment growth rate. *Note.* The solid black line represents the actual first difference of the logarithmic adjusted employed persons series, while the dashed line represents the fitted values from the dynamic regression.

An Analysis of US Pronatalist Fertility Policy at State Level

Yifei Li, Chenbing Zhou, and Tianshu Xia

University of Colorado Denver

Abstract— **In recent decades, the United States has experienced a continuous decrease in fertility, which raises concerns regarding labor force shortages and the sustainability of old-age welfare programs. Several states have implemented pronatalist policies, but their effectiveness has not been fully examined. This paper analyzes the impact of parental leave and the child tax credit on fertility at the state level between 1999 and 2023. We construct a panel dataset including fertility and net fertility as the key dependent variables. Net fertility is defined as the infant-survival-adjusted fertility rate, calculated using the infant mortality rate and the general fertility rate. We apply a two-way fixed effects model controlling for both state and year effects. We find a weak marginal positive effect of parental leave on net fertility, while the child tax credit shows no statistically significant impact. These findings align with previous literature suggesting that work-related leave policies have a more consistent positive impact than direct cash transfers.**

Keywords: fertility, parental leave, child tax credit, pronatalist policy, two-way fixed effects

I. Introduction

In recent years, the fertility rate in the United States has continued to decline. By 2020, the fertility rate had dropped to 1.64, well below the replacement level of 2.1. This downward trend poses serious challenges for future labor supply, economic growth potential, and the long-term sustainability of the social security system. How to design effective public policies that support families and encourage childbirth has become an increasingly important issue for both policymakers and researchers.

Although existing studies on fertility policies have grown in number, most of them focus on nationwide programs at the federal level and offer limited insight into state-level policy variation. In the absence of a unified federal system, such as paid parental leave, many U.S. states have introduced their own fertility-support policies. Among them, two of the most prominent income-support measures are paid parental leave and child tax credits. These policies aim to ease the economic and time burdens faced by families during the early stages of child-rearing, thereby potentially boosting fertility.

This study centers around two main questions: (1) Do state-level income-supportive fertility policies have measurable effects on fertility rates? and (2) How do different types of policies, namely parental leave and tax credits, compare in terms of their effectiveness and impact pathways? Guided by Becker's economic theory of the family, we hypothesize that greater economic support reduces the opportunity costs of having children, which in turn increases the likelihood of childbearing.

Using panel data from 50 states and DC between 1999 and 2023, we employ a two-way fixed effects model to estimate the impacts of both paid parental leave and child tax credits on net fertility. Net fertility is calculated as:

$$\text{Net Fertility}_{it} = (1000 - IMR_{it}) \times \frac{GFR_{it}}{1000} \tag{1}$$

which measures the effective fertility that children are expected to survive into adulthood.

In addition, we conduct subgroup analyses by racial composition and poverty level to investigate heterogeneity in policy effects. Our preliminary findings provide suggestive evidence at the 10% significance level that parental leave has a small positive impact on net fertility, especially in high-poverty states, while the marginal effect of child tax credits remains minimal and statistically insignificant across all model specifications.

This paper contributes to the literature in three ways. First, it fills a gap by offering a side-by-side comparison of two major state-level fertility support policies within a unified empirical framework. Second, it uses net fertility, rather than the standard fertility rate, as the main outcome variable, which better reflects actual population growth by accounting for infant mortality. Third, by leveraging updated data and conducting heterogeneity analyses, we provide targeted and data-driven insights for policy design, especially in identifying high-need populations where policy interventions may generate the greatest marginal returns.

II. Literature Review

A. Theoretical Foundation: Becker's Economic Model of Fertility

Gary Becker's seminal 1973 work was among the first to incorporate fertility decisions into the framework of economic analysis. He treated children as a type of "durable good" and proposed that families face a trade-off between the number and quality of children. As income increases, parents tend to have fewer children while investing more resources in each child's quality, such as education and health. Although children are considered "normal goods," meaning that demand for children increases with income, the substitution effect driven by the rising opportunity cost of time and stronger preferences for child quality can outweigh the income effect, leading higher-income families to have fewer children (Becker & Lewis, 1973). This theoretical framework provides an important foundation for understanding how policies such as parental leave and child tax credits may influence fertility behavior.

B. Evidence on Parental Leave: Mixed Effects and Design Considerations

A large body of empirical research has examined the impact of parental leave policies on fertility. The findings are mixed: some studies show that extended leave boosts fertility, while others find minimal or even negative effects.

In East Germany, after paid maternity leave was extended from 18 to 26 weeks in 1976 (and up to 52 weeks for families with two or more children), birth rates rose sharply in the following years (Lalive & Zweimüller, 2005). In Romania, when paid maternity leave was extended from two months to one year in 1990, the reform increased the probability of eligible mothers having an additional child within seven years by 2.5 percentage points (Hiriscau, 2024). Similarly, a quasi-natural experiment in Austria found that extending paid leave from the child's first to second birthday increased the likelihood of having another child within three years by five percentage points, approximately a 15% increase (Lalive & Zweimüller, 2005). These cases suggest that generous parental leave policies are often associated with higher fertility.

However, not all parental leave policies produce strong fertility effects. Their effectiveness depends heavily on policy design. For example, Norway's maternity leave extensions from four to eight months did not significantly affect fertility rates, possibly because the policy expansion was relatively small and moderate leave coverage already existed beforehand (Farré & González, 2019). In some cases, reforms targeting fathers may even reduce fertility. Spain introduced a two-week paid paternity leave in 2007, and subsequent research found that eligible parents were less likely to have another child within six years compared to ineligible parents. Possible explanations include increased father involvement in childcare raising the perceived costs of additional children and a decline in men's preferences for larger families after direct parenting experience (Farré & González, 2019).

Most of these studies rely on quasi-natural experiments based on policy timing, comparing treatment and control groups through econometric approaches such as difference-in-differences. These methods improve the credibility of causal inference.

C. Evidence on Child Tax Credits: Modest but Varying Fertility Effects

Research on child tax credits and fertility is relatively smaller in scale but has grown steadily in recent years. In general, these financial incentives aim to encourage fertility by reducing the economic costs associated with raising children, although the overall effects are often modest.

During the 1980s, both Canada and the United Kingdom expanded family benefits and experienced increases in fertility rates (Lalive & Zweimüller, 2005). A study covering 22 OECD countries also found a positive and statistically significant relationship between child benefits and fertility, though the estimated effect size remained relatively small (Lalive & Zweimüller, 2005).

More generous and targeted incentives appear to generate stronger fertility responses. For example, Quebec introduced a newborn allowance program during the 1980s that offered up to CAD 8,000 (approximately USD 5,800) per child. The program significantly increased fertility, especially for second and third births, with stronger responses observed among low-income households (Milligan, 2005). Likewise, Spain introduced a one-time EUR 2,500 "baby bonus" in 2007. The policy generated an immediate 2–3% increase in births, suggesting that short-term financial incentives can influence both the timing and likelihood of childbearing. However, fertility rates declined again after the policy ended ("Birth Subsidies and Fertility: Evidence From Spain," 2020).

Overall, these findings suggest that financial incentives can stimulate short-term fertility increases, particularly when policies are large in scale and well targeted. Nevertheless, the broader long-term demographic effects remain limited. Existing literature generally concludes that the effectiveness of such incentives depends heavily on policy generosity, duration, and household behavioral responses.

D. Gaps in U.S. Literature and This Study's Contribution

One major limitation in existing literature is that most U.S.-based studies focus primarily on federal-level programs. Few studies systematically evaluate the combined effects of parental leave and child tax credits on fertility at the state level. Yet in recent years, many U.S. states have independently adopted these types of policies.

To date, limited research has fully utilized variation in the timing and structure of state-level fertility-support policies to estimate their effects. This study seeks to address this gap by constructing a panel dataset covering multiple states and years using updated data through 2023. In addition, we distinguish between fertility and net fertility and compare the relative effects of parental leave and tax incentives within a unified empirical framework. By doing so, this paper contributes new evidence to the growing literature on fertility policy and demographic outcomes.

III. Data

This study constructs a U.S. state-level panel dataset with key variables spanning from 1999 to 2023. The core dependent variables are the general fertility rate and the net fertility rate. The fertility rate is defined as the number of live births per 1,000 women of reproductive age (15–44). The net fertility rate is a derived measure that adjusts the crude fertility rate for infant mortality, calculated as:

$$\text{Net Fertility}_{it} = (1000 - IMR_{it}) \times \frac{GFR_{it}}{1000} \qquad (2)$$

This measure better captures the "effective" number of births expected to survive into adulthood and eventually enter the labor force. Data for both fertility rates and infant mortality rates were obtained from the Centers for Disease Control and Prevention (CDC).

The primary independent variables are fertility-support policies. First, *Parental Status* is a dummy variable indicating whether a paid parental leave policy was implemented in a state for a given year. These data were compiled from the Bipartisan Policy Center. Second, *Credit Eligible* is a continuous variable measuring the cumulative value of child tax credits. It is calculated as the annual credit amount per child multiplied by the eligible age range. These data were collected from the Tax Policy Center.

The analysis also incorporates several demographic and socioeconomic control variables. These include population composition by sex and race, high school graduation rates, marital status distribution, the unemployment rate, Medicaid enrollment rates, and real Gross Domestic Product (GDP). GDP data were obtained from the Bureau of Economic Analysis (BEA), while poverty-related measures were retrieved from the U.S. Census Bureau and incorporated into selected model specifications.

The dataset was assembled by merging information from these sources using state and year identifiers in Stata. Fertility rate, net fertility rate, and cumulative child tax credit values were treated as continuous variables, while parental leave policy was coded as a binary indicator (1 if active, 0 otherwise). In addition, unemployment rate data for 2023 were unavailable at the time of collection, resulting in a final sample of 1,216 state-year observations.

Table 1 presents descriptive statistics for the core variables used in the analysis.

IV. Methodology

Based on panel data from 50 states and DC spanning 1999 to 2023, this study uses net fertility rate as the primary dependent variable. By employing both baseline regression models and the Two-Way Fixed Effects (TWFE) framework, we systematically examine the impacts of income-supportive fertility policies—including paid parental leave and child tax credits—on fertility outcomes. In addition, heterogeneity analyses incorporating poverty levels are conducted, while demographic, economic, and other control variables are included to improve estimation accuracy.

The primary empirical strategy employed in this paper is the Two-Way Fixed Effects (TWFE) regression model, which is a standard and widely used approach for policy evaluation with panel data. This model is particularly suitable for our research objective of estimating the impact of pronatalist policies on fertility outcomes across states and over time. By incorporating both state fixed effects and year fixed effects, the model controls for unobserved time-invariant heterogeneity across states, such as cultural norms and religious beliefs, as well as common temporal shocks affecting all states simultaneously, including national economic cycles and broader macroeconomic conditions. This specification helps mitigate omitted variable bias and improves the credibility of the estimated policy effects.

For comparative purposes, this paper also reports estimates from Ordinary Least Squares (OLS) regressions both with and without fixed effects. These models serve as useful benchmarks for illustrating the confounding influence of unobserved state-level and time-specific factors.

The baseline TWFE specification is presented as follows:

$$\begin{aligned} NF_{it} &= \beta_0 + \beta_1 PL_{it} + \beta_2 CTC_{it} \\ &\quad + \boldsymbol{X}'_{it}\gamma + \alpha_i + \lambda_t + \varepsilon_{it}. \end{aligned} \tag{3}$$

where $NetFertility_{it}$ denotes the net fertility rate in state i during year t, $ParentalLeave_{it}$ indicates whether a paid parental leave policy is active, and $CreditEligible_{it}$ measures the cumulative value of child tax credits. $\boldsymbol{X}_{it}$ represents a vector of demographic and socioeconomic control variables, including racial composition, unemployment rate, educational attainment, Medicaid participation, and GDP. α_i captures state fixed effects, λ_t captures year fixed effects, and ε_{it} denotes the error term.

Despite the advantages of the TWFE framework, we remain mindful of several methodological limitations relevant to our setting. Recent econometric literature has critically examined the standard TWFE estimator in contexts involving staggered policy adoption, emphasizing the possibility of biased estimates when treatment effects are heterogeneous or when the parallel trends assumption is violated. These concerns are particularly relevant for our analysis due to several data limitations, including the relatively low variation in key policy variables, the small number of treated states, and the high multicollinearity among demographic and economic control variables, all of which may reduce estimation precision.

Although our current analysis relies on the TWFE framework as the most appropriate specification given the structure of the available data, we acknowledge that more recent staggered Difference-in-Differences (DID) estimators may offer methodological advantages in future research. These approaches are specifically designed to generate more reliable causal estimates under staggered adoption and heterogeneous treatment effects. However, their implementation generally requires a larger number of treated units and longer policy histories than are currently available in our dataset.

Therefore, we present the TWFE estimates as a foundational analysis of the relationship between fertility-supportive policies and fertility outcomes, while recognizing that future research with expanded data availability and broader policy diffusion would be better positioned to apply more advanced causal inference techniques. In all current specifications, robust standard errors clustered at the state level are reported to account for potential serial correlation within states over time.

V. Results

We first estimate both OLS and fixed effects regressions for the fertility rate and net fertility rate, respectively (see Tables 2 and 3). After comparing the performance of the two

TABLE I: Descriptive Statistics

Variable	N	Mean	Std. Dev.	Min	Max
Fertility Rate	1216	63.690	7.757	44.300	95.000
Net Fertility Rate	1216	62.910	7.424	44.083	94.531
Parental Status	1216	0.049	0.215	0.000	1.000
Credit Eligible	1216	161.975	904.598	0.000	6000.000
Male	1216	48.590	1.111	45.693	51.365
White	1216	79.726	14.016	20.829	98.273
Black	1216	11.398	11.361	0.123	57.605
Asian	1216	3.588	5.704	0.000	41.195
High School Graduation Rate	1216	0.813	0.041	0.686	0.899
Married	1216	0.541	0.040	0.320	0.633
Unemployment Rate	1216	5.304	1.939	2.400	10.900
Medicaid	1216	0.156	0.050	0.054	0.277
GDP	1216	337000	401000	27548.100	2220000
Observations	1216				

Notes: Sample size $N = 1216$. All continuous variables are winsorized at the 1% level.

outcome variables, we select net fertility rate as the primary dependent variable for subsequent analysis.

Table 2 reports the regression results for the fertility rate. In the baseline OLS model without fixed effects (Column 1), the coefficient on parental leave policy is negative and statistically significant at the 1% level. However, after controlling for state and year fixed effects, the coefficient becomes positive and statistically insignificant. The substantial changes in both the magnitude and direction of the coefficient suggest that the estimated relationship between parental leave policy and fertility rate is highly sensitive to omitted state-level and time-specific factors.

In contrast, the net fertility rate exhibits more stable patterns across specifications. As shown in Table 3, the baseline OLS regression does not produce a statistically significant relationship between parental leave policy and net fertility. Moreover, the explanatory power of the fertility rate models remains extremely low, with R^2 values ranging only from 0.02 to 0.05. This is substantially lower than the R^2 values for the net fertility regressions, which range from 0.33 to 0.57. These results suggest that the fertility rate alone may not adequately capture the actual effects of fertility-supportive policies and socioeconomic controls. By accounting for infant mortality, the net fertility rate provides a more reliable measure of effective population growth and is therefore more appropriate for the purposes of this study.

Table 3 presents the main regression results using net fertility rate as the dependent variable. Column 1 reports the baseline OLS regression without controlling for state or year heterogeneity. In this specification, the estimated coefficient on parental leave policy is slightly negative. This counterintuitive result may reflect omitted variable bias arising from unobserved differences across states, such as fertility culture, demographic composition, or long-standing institutional characteristics.

After incorporating state fixed effects and year fixed effects (Columns 2–4), the estimated coefficient on parental leave policy becomes positive. In the full TWFE specification (Column 4), the coefficient reaches statistical significance at the 10% level. Specifically, the implementation of parental leave policy is associated with an approximately 1.67-point increase in the net fertility rate. In contrast, the estimated effects of child tax credits remain economically small and statistically insignificant across all specifications. Nevertheless, the coefficient becomes positive in the TWFE specification, which is broadly consistent with the theoretical expectation that fertility-supportive policies may encourage childbearing behavior.

The overall goodness-of-fit of the model also improves substantially after incorporating both state and year fixed effects. The R^2 increases from 0.419 in the baseline OLS model to 0.574 in the TWFE model. This improvement suggests that the fixed effects specification captures important cross-state differences and common temporal trends omitted in simpler regression models. More broadly, the TWFE framework is well suited to the structure of panel data because it controls for persistent state-specific characteristics, such as fertility culture and institutional differences, as well as year-specific macroeconomic conditions and nationwide demographic trends.

Table 4 reports the results from stepwise TWFE regressions with the gradual inclusion of control variables. Across all specifications, the coefficient on parental leave policy remains consistently positive. After all demographic, economic, and social controls are included, the coefficient remains statistically significant at the 10% level. This pattern suggests that parental leave policies may help reduce the economic and time burdens associated with childbearing and childrearing, thereby contributing to higher net fertility outcomes.

By contrast, the estimated coefficient on child tax credits remains extremely small and statistically insignificant throughout all model specifications. The near-zero marginal effects suggest that child tax credit policies, at least within the context of our sample and measurement approach, do not exert a substantial influence on fertility outcomes.

The estimated effects of control variables are generally consistent with theoretical expectations. Economic stability appears to play an important role in fertility decisions. The coefficient on unemployment rate is negative and statistically significant at the 1% level, indicating that economic downturns and labor market uncertainty are associated with lower net fertility rates. Similarly, GDP exhibits a negative and statistically significant relationship with fertility, consistent with the broader demographic pattern that higher levels of economic development are often associated with lower fertility rates.

The coefficient on Medicaid coverage is positive, suggesting that greater access to medical support may reduce the perceived risks associated with childbirth and increase fertility willingness, although the estimate is not statistically significant. Demographic variables also display heterogeneous relationships with fertility outcomes. The coefficients on the proportions of White and Black populations are positive and generally larger than that of the Asian population share, suggesting meaningful racial differences in fertility patterns across states. The coefficient on high school graduation rate is negative, which may reflect the tendency for more educated populations to prioritize child quality over quantity. In addition, the positive coefficient on marriage rate is broadly consistent with the expectation that marriage provides a more stable environment for childbearing.

To further examine heterogeneity in policy effects, Table 5 divides the sample into low-, middle-, and high-poverty groups according to state poverty levels. States with poverty rates below 10.4 are classified as low poverty, states between 10.5 and 13.5 are classified as middle poverty, and states above 13.6 are classified as high poverty.

Across all poverty subgroups, the estimated coefficient on parental leave policy remains positive. The largest coefficient appears in the high-poverty group, where parental leave policy is associated with an increase of approximately 2.7 points in net fertility. This result suggests that lower-income regions may exhibit stronger fertility responses to parental leave policies. Although statistical power remains limited due to the relatively small sample size within the high-poverty subgroup, the direction of the coefficient remains stable across specifications and does not reverse sign.

By contrast, the estimated coefficients for child tax credits remain close to zero and statistically insignificant across all poverty groups. These findings suggest that the effectiveness of child tax credit policies does not vary substantially by poverty level and that the policy may not generate sufficiently strong incentives to influence fertility behavior in a meaningful way.

VI. Conclusion

Recent policy discussions in the United States have expanded with the federal introduction of "Trump Accounts," a new type of investment account designed for children under the age of eighteen. Under the current proposal, eligible children may receive a pilot deposit of USD $1,000 from the U.S. Treasury at birth. Although the policy remains at an early stage and faces considerable implementation challenges, it reflects growing policy interest in long-term family support and inclusive asset-building strategies. More broadly, these developments highlight the increasing attention policymakers are giving to fertility, child welfare, and intergenerational economic security.

This study investigates the relationship between income-supportive fertility policies—specifically parental leave and child tax credits—and net fertility rates using panel data from 50 states and DC spanning 1999–2023. Several key findings emerge from the analysis.

First, parental leave policies exhibit a weak positive association with net fertility rates in the two-way fixed effects specification, reaching statistical significance at the 10% level. The positive relationship remains relatively stable across subgroup analyses based on poverty levels, with the largest estimated effect appearing in high-poverty states. These findings suggest that parental leave policies may help reduce the economic and time burdens associated with childbirth and childrearing, particularly for lower-income households.

In contrast, child tax credit policies, measured through the *Credit Eligible* variable, display economically small and statistically insignificant coefficients across nearly all model specifications and subgroup analyses. This result suggests that the current structure of child tax credit policies may not sufficiently address the primary constraints influencing fertility decisions. Compared with parental leave policies, direct tax-based financial incentives appear to exert weaker associations with fertility outcomes within the context of this study.

The heterogeneity analysis further highlights important policy implications. Although parental leave policies display positive associations across most subgroups, their estimated effects appear stronger in high-poverty states, suggesting that lower-income households may face greater unmet demand for fertility-related support. By contrast, child tax credit policies exhibit limited variation across poverty groups and do not appear to generate substantial fertility responses in either low- or high-poverty regions.

These findings provide several implications for policymakers seeking to improve fertility-supportive policies. Given the relatively stronger performance of parental leave policies, especially in economically disadvantaged areas, policymakers may consider expanding paid leave duration, increasing subsidy levels, and improving policy accessibility for low-income households and small businesses. For child tax credit programs, the absence of statistically meaningful effects suggests that future reforms may need to move beyond purely financial transfers. Integrating tax credits with practical support mechanisms—such as childcare assistance, healthcare access, or educational subsidies—may better address the multidimensional costs associated with raising children.

Despite the contributions of this study, several limitations should be acknowledged. First, many state-level fertility-supportive policies were implemented relatively recently, limiting the available time horizon for long-run evaluation and making more meaningful lagged-effect analysis difficult. Second, policy adoption has been concentrated disproportionately in economically developed states, resulting in a relatively clustered treatment pattern that may affect estimation precision. Third, several policy and demographic variables exhibit substantial multicollinearity, potentially limiting the explanatory power of certain specifications. Finally, the study period overlaps heavily with the COVID-19 pandemic, which generated significant disruptions to fertility behavior, labor markets, and socioeconomic conditions, thereby complicating efforts to isolate policy-specific effects.

Consequently, while this study provides meaningful evidence regarding the relationship between fertility-supportive policies and fertility outcomes, the findings should be interpreted with appropriate caution given the relatively limited policy history and available variation within the data. Future research would benefit from longer time horizons, broader policy diffusion, and expanded state-level variation. In addition, future studies may benefit from applying more advanced staggered Difference-in-Differences (DID) estimators to improve causal identification under staggered policy adoption settings. Researchers may also extend the analysis to broader outcome measures, including child well-being, family stability, and long-term household welfare, in order to evaluate the multidimensional effects of fertility-supportive policies more comprehensively.

Overall, the findings of this study suggest that parental leave policies exhibit more meaningful associations with net fertility outcomes than child tax credit policies, particularly among economically disadvantaged populations. Future fertility-supportive policy design may therefore benefit from placing greater emphasis on targeted and multidimensional support mechanisms for high-need households. Continued academic research using improved data and more rigorous empirical strategies will be essential for informing future family-support and demographic policy reforms.

REFERENCES

[1] Becker, G. S., and H. G. Lewis. On the interaction between the quantity and quality of children. *Journal of Political Economy* 81.2 (1973). https://www.jstor.org/stable/1840425

[2] BEA Interactive Data Application. 2025. https://apps.bea.gov/itable/

[3] Bipartisan Policy Center. State paid family leave laws across the U.S. Bipartisan Policy Center, 20 Feb. 2025, https://bipartisanpolicy.org/explainer/state-paid-family-leave-laws-across-the-u-s/

[4] Birth subsidies and fertility: Evidence from Spain. Barcelona School of Economics, 5 May 2020, https://focus.bse.eu/birth-subsidies-fertility-evidence-spain/

[5] Farré, L., and L. González. Does paternity leave reduce fertility? *Journal of Public Economics* 172 (2019): 52–66. https://doi.org/10.1016/j.jpubeco.2018.12.002

[6] Fertility rate. Stats of the States, 20 Aug. 2025, https://www.cdc.gov/nchs/state-stats/births/fertility-rate.html

[7] How do state child tax credits work? Tax Policy Center, 2025, https://taxpolicycenter.org/briefing-book/how-do-state-child-tax-credits-work

[8] Hiriscau, A. The effect of paid maternity leave on fertility and mothers' labor force participation. *Journal of Labor Research* 45.3 (2024): 350–384. https://doi.org/10.1007/s12122-024-09361-0

[9] Infant mortality. Stats of the States, 20 Aug. 2025, https://www.cdc.gov/nchs/state-stats/deaths/infant-mortality.html

[10] Lalive, R., and J. Zweimüller. Does parental leave affect fertility and return-to-work? Evidence from a true natural experiment. *SSRN Electronic Journal* (2005). https://doi.org/10.2139/ssrn.731766

[11] Milligan, K. Subsidizing the stork: New evidence on tax incentives and fertility. National Bureau of Economic Research, 2005, https://doi.org/10.3386/w8845

[12] U.S. Census Bureau. Current Population Survey, 1960 to 2025 Annual Social and Economic Supplements (CPS ASEC). https://www2.census.gov/programs-surveys/cps/techdocs/cpsmar25.pdf

Appendix

TABLE II: OLS and Fixed Effects Regression Results for Fertility Rate

VARIABLES	(1) OLS Fertility Rate	(2) Fixed id (state) Fertility Rate	(3) Fixed year Fertility Rate	(4) Fixed id and year Fertility Rate
Parental Status	-2.837***	-1.027	0.976	1.713
	(0.000)	(0.321)	(0.302)	(0.118)
Credit Eligible	-0.000161	-0.000408	-0.000199	-0.000079
	(0.283)	(0.248)	(0.348)	(0.756)
Male	1.850***	0.254	0.515**	0.423**
	(0.000)	(0.298)	(0.013)	(0.042)
White	-0.316***	-0.044	-0.113**	0.009
	(0.000)	(0.769)	(0.031)	(0.940)
Black	-0.188***	-0.007	-0.097	0.024
	(0.001)	(0.967)	(0.165)	(0.877)
Asian	-0.253**	0.040	-0.114	-0.057
	(0.013)	(0.325)	(0.143)	(0.220)
High School Graduate Rate	-8.326***	-6.537***	-1.545	1.125
	(0.000)	(0.000)	(0.405)	(0.583)
Married	4.261***	4.448	-0.051	-1.088
	(0.000)	(0.593)	(0.996)	(0.390)
Unemployment Rate	0.073	0.156*	-0.796***	-0.855***
	(0.395)	(0.064)	(0.000)	(0.000)
Medicaid	-1.974***	-1.640*	-1.722	-1.311
	(0.000)	(0.088)	(0.132)	(0.295)
GDP	-0.050	-0.848	-0.373	-0.967
	(0.202)	(0.710)	(0.301)	(0.521)
Constant	49.334***	109.897***	68.074***	47.604**
	(0.000)	(0.000)	(0.001)	(0.018)
Observations	1,216	1,216	1,216	1,216
R-squared	0.053	0.020	0.034	0.052
Individual Effect	NO	YES	NO	YES
Time Effect	NO	NO	YES	YES
Number of IDs	51	51	51	51

Robust p-values in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

TABLE III: OLS and Fixed Effects Regression Results for Net Fertility Rate

VARIABLES	(1) OLS Net Fertility Rate	(2) Fixed id (state) Net Fertility Rate	(3) Fixed year Net Fertility Rate	(4) Fixed id and year Net Fertility Rate
Parental Status	-0.033	1.093	1.440	1.674*
	(0.453)	(0.945)	(0.743)	(0.075)
Credit Eligible	-0.000032	-0.000256	0.000070	0.000117
	(0.891)	(0.494)	(0.762)	(0.730)
Male	2.533	-1.061	1.453***	-1.019
	(0.138)	(0.446)	(0.001)	(0.485)
White	-0.729**	1.185	-0.622	1.619
	(0.028)	(0.327)	(0.111)	(0.257)
Black	-0.539*	0.847	-0.492	1.217
	(0.073)	(0.320)	(0.182)	(0.242)
Asian	-0.801*	0.479	-0.657	0.249
	(0.093)	(0.354)	(0.134)	(0.475)
High School Graduate Rate	-75.971***	-10.297**	-7.374***	-4.974
	(0.000)	(0.011)	(0.002)	(0.425)
Married	66.592***	1.158	61.921*	1.374
	(0.001)	(0.920)	(0.066)	(0.696)
Unemployment Rate	0.773	0.217*	0.124	-1.347***
	(0.148)	(0.073)	(0.887)	(0.004)
Medicaid	-8.402	2.977	-1.674	2.507
	(0.414)	(0.550)	(0.405)	(0.567)
GDP	-0.113	-0.240***	-0.173***	-0.366***
	(0.258)	(0.004)	(0.002)	(0.001)
Constant	32.221	88.640**	76.654**	-0.387
	(0.660)	(0.036)	(0.044)	(0.993)
Observations	1,216	1,216	1,216	1,216
R-squared	0.419	0.331	0.562	0.574
Individual Effect	NO	YES	NO	YES
Time Effect	NO	NO	YES	YES
Number of IDs	51	51	51	51

Robust p-values in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

TABLE IV: Stepwise Regression Results for Net Fertility Rate

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Parental Status	1.055	1.027	1.008	1.854	3.217	2.715	2.868	2.767	3.261	3.274	1.674*
	(0.525)	(0.542)	(0.543)	(0.431)	(0.292)	(0.273)	(0.280)	(0.278)	(0.184)	(0.190)	(0.075)
Credit Eligible		-0.000108	-0.000130	0.000015	0.000022	0.000010	0.000022	0.000026	0.000023	0.000016	0.000117
		(0.554)	(0.395)	(0.946)	(0.915)	(0.960)	(0.916)	(0.904)	(0.955)	(0.969)	(0.730)
Male			-0.830	-0.979	-1.014	-0.932	-0.934	-0.945	-1.038	-1.021	-1.019
			(0.554)	(0.520)	(0.508)	(0.515)	(0.513)	(0.511)	(0.483)	(0.484)	(0.485)
White				0.783	1.410	1.565	1.587	1.571	1.586	1.645	1.619
				(0.275)	(0.210)	(0.228)	(0.231)	(0.229)	(0.217)	(0.236)	(0.257)
Black					1.087	1.257	1.239	1.252	1.195	1.237	1.217
					(0.187)	(0.214)	(0.212)	(0.217)	(0.202)	(0.221)	(0.242)
Asian						0.272	0.268	0.253	0.227	0.238	0.249
						(0.459)	(0.460)	(0.464)	(0.496)	(0.496)	(0.475)
High School Graduate Rate							-4.364	-4.497	-5.504	-5.313	-4.974
							(0.427)	(0.430)	(0.366)	(0.361)	(0.425)
Married								1.555	0.818	1.404	1.374
								(0.577)	(0.768)	(0.689)	(0.696)
Unemployment Rate									-1.272***	-1.329***	-1.347***
									(0.002)	(0.007)	(0.004)
Medicaid Coverage Rate										2.470	2.507
										(0.579)	(0.567)
GDP											-0.366***
											(0.001)
Constant	64.758***	64.758***	104.918	45.608**	16.860	35.739	3.167	9.157	11.995	-1.224	-0.387
	(0.000)	(0.000)	(0.120)	(0.049)	(0.667)	(0.543)	(0.932)	(0.821)	(0.735)	(0.980)	(0.993)
Observations	1,267	1,267	1,267	1,267	1,267	1,267	1,267	1,267	1,216	1,216	1,216
R-squared	0.320	0.381	0.461	0.492	0.497	0.551	0.554	0.556	0.564	0.567	0.574
Number of IDs	51	51	51	51	51	51	51	51	51	51	51

Robust p-values in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

TABLE V: Heterogeneity Regression Results of Net Fertility Rate by Poverty Level

VARIABLES	(1) Low Poverty Net Fertility Rate	(2) Mid Poverty Net Fertility Rate	(3) High Poverty Net Fertility Rate	(4) Full Sample Net Fertility Rate
Parental Status	1.053	1.071	2.700*	2.700*
	(0.207)	(0.556)	(0.055)	(0.054)
Credit Eligible	0.001070	0.000539	-0.000169	-0.000169
	(0.422)	(0.620)	(0.559)	(0.558)
Male	1.078	-5.210	0.561**	0.561**
	(0.268)	(0.346)	(0.015)	(0.014)
White	2.322	2.381	-0.197	-0.197
	(0.287)	(0.123)	(0.319)	(0.318)
Black	2.884	0.455	-0.259	-0.259
	(0.179)	(0.542)	(0.203)	(0.202)
Asian	0.702	-0.076	-0.153	-0.153
	(0.330)	(0.765)	(0.539)	(0.538)
High School Graduate Rate	50.324	-178.374	15.744	15.744
	(0.366)	(0.296)	(0.381)	(0.380)
Married	-83.876	114.506	-5.768	-5.768
	(0.157)	(0.421)	(0.617)	(0.616)
Unemployment Rate	1.322	-2.243	-0.802***	-0.802***
	(0.507)	(0.121)	(0.009)	(0.008)
Medicaid	10.062	-3.316	-9.431	-9.431
	(0.397)	(0.943)	(0.481)	(0.480)
GDP	0.497	0.100**	-0.132***	-0.132***
	(0.840)	(0.036)	(0.001)	(0.002)
Constant	-219.148	189.655	60.138***	60.138***
	(0.377)	(0.359)	(0.008)	(0.007)
Observations	417	423	376	1,216
R-squared	0.589	0.401	0.509	0.574
Number of IDs	35	45	36	51

Robust p-values in parentheses
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

TABLE VI: Timeline of Fertility Policy Implementation by U.S. State

State	Child Tax Credit Policy Implementation Year	Paid Parental Leave Policy Implementation Year
Arizona	2019	—
Arkansas	—	2023
California	2019	2004
Colorado	2022	2024 (planned)
Connecticut	—	2022
Delaware	—	2026 (planned)
District of Columbia	—	2020
Florida	—	2023
Georgia	—	2022
Hawaii	—	2021
Idaho	2018	—
Kentucky	—	2024 (planned)
Maine	2018	2026 (planned)
Maryland	2020	2026 (planned)
Massachusetts	2023	2021
Minnesota	2023	2026 (planned)
New Hampshire	—	2023
New Jersey	2023	2009
New Mexico	2022	—
New York	2006	2018
Oklahoma	2023	—
Oregon	2023	2023
Rhode Island	—	2014
South Carolina	—	2023
Tennessee	—	2024 (planned)
Texas	—	2023
Utah	2023	—
Vermont	2022	—
Virginia	—	2022
Washington	—	2019

War and Literacy in Liberia: Regional, Cohort, and Ethnic Effects from the Civil Wars (1989 to 2003)

M. Amara Dukuly

University of Minnesota

Abstract— **On Christmas Eve 1989, a Libyan-backed rebel group led by Charles Taylor invaded Liberia, sparking the First Liberian Civil War, which lasted until 1997. A second conflict followed in 1999 and ended in 2003. After nearly a decade and a half of civil war, Liberia is one of the poorest countries in the world and half of the population is illiterate. This paper examines the effects of the First and Second Liberian Civil Wars on literacy by county, cohort, and ethnicity. Using cross-sectional data from the Standard Demographic and Health Surveys (DHS 2007, 2013, 2019) and the Uppsala Conflict Data Program, I construct a weighted treatment variable measuring the total share of conflict deaths per county for individuals who were primary school-aged during the war. In a survey-weighted linear probability model, I interact total death exposure with county, birth cohort, and ethnicity for individuals born between 1964 and 1998, holding survey year, gender, and urban status constant. I find that nearly every county in Liberia showed statistically significant lower levels of literacy. Cohorts born between 1982 and 1993 exhibit significantly lower literacy probabilities relative to the pre-war cohort. Among ethnic groups, the Gio show a significant positive interaction, while the Mandingo, Krahn, and Mano exhibit negative but insignificant interactions. These findings suggest the Civil Wars disrupted literacy through geographic, cohort, and tribally-shaped channels.**

I. Introduction and Historical Context

Liberia is a small country located on the coast of West Africa. The land was inhabited by several indigenous tribes until the American Colonization Society brought former enslaved people to the country and founded what is now known as “Liberia,” which stands for liberty. Following a pattern common to settler-colonial dynamics, the newcomers soon dominated the government and their politics, marginalizing indigenous Liberians. This dynamic between Americo-Liberians and Indigenous Liberians persisted over the next century, with many natives working subservient roles for Americo-Liberians.

Significant economic, political, and wealth disparities persisted until April 15th, 1980, when Sergeant Samuel Doe seized power and assassinated former president William Tolbert in a coup d’état. Over the next decade, Doe drove the once agricultural and export-rich economy to collapse due to corruption and gross mismanagement of funds. Although Doe’s presidency marked the first government led by a native Liberian, he continued the previous dynamic upheld by Americo-Liberians — this time favoring his Krahn and Mandingo counterparts, who would take the majority of government positions despite being minority tribes.

The earliest conflict before the First Liberian Civil War can be dated to 1985, when Thomas Quiwonkpa, a former associate of Doe, attempted to overthrow his regime. Following the failed coup, Doe retaliated against the Gio and Mano tribes in Nimba County (The Advocates for Human Rights, n.d.). On December 24th, 1989, a rebel group led by Charles Taylor invaded Nimba County, backed by the Gio and Mano tribes, marking the beginning of the First Liberian Civil War. Doe was captured and killed in September 1990. The conflict continued until 1997, when Charles Taylor was elected president.

Despite the formal ending of the war, violence remained rampant. Mandingo-led factions invaded Lofa County in 1999, sparking a second civil war. This lasted until 2003, after peace agreements facilitated by Liberian women and surrounding nations. Over 250,000+ civilians were killed during both wars. 80% of the population—2 million people—was displaced (Atkinson, 2006).

II. Literature Review and Motivation

Civil war and armed conflict vary widely in intensity and origin. Low-income countries are highly susceptible to these conflicts, particularly in sub-Saharan Africa. As of early 2026, the majority of countries in the continent are experiencing high or extreme conflict levels.[1] The economics of war are driven by inequality, governance, ethnic and religious division, and proximity to natural resources (Collier & Hoeffler, 2004).

The costs of armed conflict are extensive. Beyond direct casualties and infrastructure destruction, war produces significant declines in economic growth—a “reversal in development” particularly severe in low-income countries with weak institutions (Olaberria & Reinhart, 2022). Women, the poor, and children are often the primary victims. Schools are frequently looted, destroyed, or repurposed as military shelters, forcing students and teachers to flee. For those of primary

. I thank my mother, my sister, and family for sharing their endless knowledge and experience of the war and Liberia. In addition to the many Liberians who shared their vicious stories of war through interviews, documentaries, and other related forms of media. Thank you to Dr. Hannes Malmberg at the University of Minnesota Department of Economics for his assistance in developing the original paper during Fall 2025 in ECON 4331W. Additionally, thank you to the Sadie Collective and Dr. Tristan Reed for his mentorship on this version. Please contact me at maradukuly@gmail.com if you have any questions, concerns, or suggestions.

1. ACLED Conflict Index defines conflict levels by deadliness, danger, geographic diffusion, and armed group fragmentation.

school age during prolonged conflict, interrupted schooling often means permanent illiteracy (Hanneman, 2005; Thompson et al., 2006).

Hanneman (2005) finds that war-induced disruptions to education are primary drivers of low literacy, and that nearly half of the countries with the world's lowest literacy rates had experienced prior or ongoing conflict. Liu (2022) extends this to Liberia specifically, finding that individuals who were school-aged during the civil wars were significantly less likely to attain formal education and showed lower employment rates than their unaffected peers.

Once the fastest-growing economy in Africa, Liberia's real GDP remains at roughly a quarter of its pre-war level. According to the 2022 Liberian Institute of Statistics report, only 58.6% of Liberians aged 15 and older were literate. Building on Liu's (2022) difference-in-differences framework, I apply a survey-weighted linear probability model to estimate the effect of conflict exposure on literacy across county, cohort, and ethnic lines.

I develop three hypotheses: (H1) Liberians with higher conflict exposure will exhibit a significantly lower probability of literacy; (H2) Gio and Mano tribes will exhibit higher literacy probabilities than the Krahn and Mandingo due to wartime political alignments; and (H3) negative effects will be disproportionately concentrated in counties with the highest shares of national conflict deaths.

III. Data

I use cross-sectional data from the Standard Demographic and Health Surveys (DHS) individual and men's recode files, supplemented by IPUMS-DHS harmonized data from 2007, 2013, and 2019. Each respondent is assigned a survey weight, primary sampling unit (PSU), and stratum. I match standardized IPUMS-DHS weights to each respondent via the original DHS case identifier (`caseid`), and apply the same matching procedure to harmonize ethnicity codes across the three survey years.

A significant limitation of cross-sectional DHS data is that it captures residence at the time of survey, not during the war years. Given that an estimated 2 million Liberians were displaced, I restrict my sample to respondents who report having lived in the same county throughout their lives, treating permanent residence as a proxy for wartime location.[2]

The DHS measures literacy through variable `v155`. I recode literacy as a *binary* outcome where only respondents who can read a *whole* sentence are coded as literate (1), and all others as not literate (0). Respondents classified as visually impaired or lacking a required language card are dropped.

To measure conflict exposure, I use the Uppsala Conflict Data Program (UCDP) Geo-referenced Event Dataset for 1989–2003. For each county-year, I calculate the share of total Liberian conflict deaths occurring in that county using UCDP's "best estimate" of total fatalities. Using county-year death shares rather than raw counts normalizes for differences in baseline county population. I restrict the sample to respondents born between 1964 and 1998, capturing pre-war, war-exposed, and early post-war cohorts.

IV. Methodology

To estimate the relationship between conflict exposure and literacy, I employ a survey-weighted linear probability model (LPM), drawing on the difference-in-differences framework of Liu (2022). I extend Liu's framework in three ways: (1) I use annual county-level shares of conflict deaths from the UCDP as the treatment variable; (2) I extend the interaction structure to include county, cohort, and ethnicity; and (3) I examine literacy directly as a proxy for human capital formation rather than as a mediator.

I do not control for educational attainment because schooling is a downstream outcome of conflict exposure, and controlling for it would introduce post-treatment bias.

A. Model Specifications

The survey-weighted LPM takes the following form, where $\mathrm{TSD}_{i,c}$ is the total weighted share of conflict deaths from age 6 onward for individual i in county c, $\mathbf{X}_i$ is a vector of individual controls (gender, urban status), δ_t denotes survey year fixed effects, and ε_i is the error term.

Model 1: County × Conflict

$$\begin{aligned}\mathrm{ltr}_{i,c,t} = \beta(\mathrm{TSD}_{i,c}) + \sum_{k}^{K} \tau_k \left(\mathrm{TSD}_{i,c} \times \mathrm{County}_k\right) \\ + \gamma \mathrm{County}_i + \lambda \mathbf{X}_i + \kappa_b + \delta_t + \varepsilon_i \qquad (1)\end{aligned}$$

where County_k ranges across all 15 Liberian counties with Nimba as the reference. β identifies the conflict–literacy relationship for Nimba specifically, while τ_k estimates the deviation from that baseline for each other county.

Model 2: Cohort × Conflict

$$\begin{aligned}\mathrm{ltr}_{i,c,t} = \beta(\mathrm{TSD}_{i,c}) + \sum_{k}^{K} \tau_k \left(\mathrm{TSD}_{i,c} \times \mathrm{Cohort}_k\right) \\ + \gamma \mathrm{Cohort}_i + \lambda \mathbf{X}_i + \delta_c + \delta_t + \varepsilon_i \qquad (2)\end{aligned}$$

where $\mathrm{Cohort}_k \in \{$1976–1981, 1982–1987, 1988–1993, 1994–1998$\}$, with 1964–1975 as the pre-war reference cohort.

Model 3: Ethnicity × Conflict

$$\begin{aligned}\mathrm{ltr}_{i,c,t} = \beta(\mathrm{TSD}_{i,c}) + \sum_{k}^{K} \tau_k \left(\mathrm{TSD}_{i,c} \times \mathrm{Eth}_k\right) \\ + \gamma \mathrm{Eth}_i + \lambda \mathbf{X}_i + \delta_c + \kappa_b + \delta_t + \varepsilon_i \qquad (3)\end{aligned}$$

where $\mathrm{Eth}_k \in \{$Gio, Mano, Mandingo, Krahn$\}$, relative to all other ethnic groups. κ_b denotes birth cohort fixed effects.

2. See Section VII for further discussion.

V. Descriptive Statistics

The full sample consists of 23,844 respondents with a survey-weighted literacy rate of 42.6%. Men and urban residents are substantially more literate than women and rural residents (Table I).

Across birth cohorts (Table II), those born between 1976 and 1981 show the lowest literacy rate at 35.3%, nearly identical to the pre-war cohort. Among the four focal ethnic groups (Table III), the Gio and Mano exhibit the lowest pre-war literacy rates at 32.6% and 28.6%. The Mandingo stand out for an extreme gender gap: women's literacy in the pre-war cohort is only 2.9% compared to men's at 58.4%. County-level literacy and conflict death shares are shown in Tables IV and V.

TABLE I: Summary Statistics: Literacy by Period

Subsection	N	N (Wtd.)	Ltr %
All Cohorts Combined			
All	23,844	23,485	42.6
Male	8,075	8,002	58.4
Female	15,769	15,483	34.5
Urban	8,872	12,315	59.0
Rural	14,972	11,170	24.6
Pre-War (1964–1975)			
All	6,354	5,936	35.4
War Cohorts (1976–1993)			
All	14,503	14,358	43.9
Post-War (1994–1998)			
All	2,987	3,191	50.6

Ltr% = can read a whole sentence (binary).
Source: DHS Standard Liberia 2007, 2013, 2019.

TABLE II: Literacy by Birth Cohort

Cohort	N	N (Wtd.)	Ltr %
Pre-War (1964–75)	6,354	5,936	35.4
War: 1976–1981	4,427	4,184	35.3
War: 1982–1987	5,196	5,120	45.0
War: 1988–1993	4,880	5,054	49.7
Post-War (1994–98)	2,987	3,191	50.6

Source: DHS Standard Liberia 2007, 2013, 2019.

TABLE III: Literacy (%) by Ethnicity and Birth Cohort

Group	Sub.	Pre-War	76–81	82–87	88–93
Gio	All	32.6	40.8	47.4	35.1
	Male	53.7	63.9	54.9	43.7
	Fem.	18.5	25.9	42.9	30.4
Mano	All	28.6	32.1	39.8	47.8
	Male	54.5	37.8	45.8	59.8
	Fem.	19.2	30.6	38.3	45.3
Mandingo	All	38.4	34.6	43.8	42.3
	Male	58.4	60.2	75.3	65.8
	Fem.	2.9	13.7	14.4	26.3
Krahn	All	45.3	41.8	44.3	55.9
	Male	63.6	63.4	74.2	68.4
	Fem.	35.2	31.2	26.0	48.9
Other	All	35.9	34.8	45.2	52.1
	Male	54.4	51.5	63.3	64.9
	Fem.	24.5	26.1	36.3	45.8

Source: IPUMS DHS Liberia 2007, 2013, 2019.

TABLE IV: Literacy Rate (%) by County and Birth Cohort

County	Pre-War (64–75)	War (76–93)	Post-War (94–98)
Montserrado	59.9	67.4	73.3
Maryland	36.0	44.9	45.3
Margibi	34.4	41.5	37.5
Grand Gedeh	25.9	38.5	35.8
Bomi	18.7	36.5	51.0
Nimba	27.4	35.0	30.0
Grand Kru	34.2	30.7	23.1
Sinoe	33.6	28.2	29.2
Lofa	21.6	26.2	29.0
River Gee	32.8	25.4	24.1
Grand Cape Mount	17.5	24.0	31.9
River Cess	20.6	23.3	35.1
Grand Bassa	16.1	22.2	33.5
Gbarpolu	21.9	22.1	24.5
Bong	21.3	21.9	33.9

Source: DHS Liberia 2007, 2013, 2019.

TABLE V: Conflict Death Share by County, 1989–2003

County	Deaths (UCDP)	Share (%)	Peak Year
Montserrado	5,233	23.1	1996
Lofa	5,216	23.0	1999
Bong	3,043	13.4	1994
Bomi	2,179	9.6	1997
Margibi	1,682	7.4	1998
River Cess	1,528	6.8	1995
Sinoe	1,023	4.5	1994
Nimba	709	3.1	1989
Grand Bassa	449	2.0	1990
River Gee	402	1.8	2003
Grand Cape Mount	343	1.5	1995
Grand Gedeh	316	1.4	1991
Maryland	246	1.1	1995
Grand Kru	148	0.7	1994
Gbarpolu	117	0.5	2002

Source: UCDP GED, 1989–2003.

VI. Results

All models use the pre-war cohort (born 1964–1975) as the reference cohort, Nimba County as the reference county, and *Other* ethnicities as the reference ethnic group. The dependent variable is a binary indicator equal to one if the respondent can read a whole sentence. All associations are correlational; the design does not support causal identification.

A. *County Heterogeneity in Literacy*

The baseline conflict coefficient for Nimba is positive but falls short of conventional significance ($\hat{\beta} = 0.078$, $p = 0.139$), consistent with the Gio and Mano tribes being largely aligned with Taylor's forces and shielded from targeting.

The county interaction terms tell a different story elsewhere. The largest effects are in River Gee ($\hat{\tau} = -0.419$, $p < 0.001$) and Bomi ($\hat{\tau} = -0.309$, $p < 0.001$). Bomi has historically low literacy rates and its peak conflict year of 1997 coincides with the schooling years of the 1982–1987 cohort. Montserrado exhibits $\hat{\tau} = -0.201$, $p < 0.01$, despite its higher pre-war literacy. By contrast, Lofa and

Grand Gedeh show interactions near zero for distinct reasons: Lofa peaked in 1999 after the most school-vulnerable cohorts had largely passed primary school age, while Grand Gedeh accounted for only 1.4% of national deaths.

B. Cohort Heterogeneity in Literacy

The baseline conflict coefficient for the pre-war cohort is positive and significant ($\hat{\beta} = 0.123$, $p < 0.05$), possibly reflecting a selection dynamic: pre-war adults were not school-aged during the conflict.

The 1982–1987 cohort exhibits the largest negative interaction ($\hat{\tau} = -0.310$, $p < 0.001$), consistent with primary school-aged children bearing the greatest burden of educational disruption. The 1988–1993 cohort similarly exhibits a significant negative interaction ($\hat{\tau} = -0.139$, $p < 0.05$). The interaction for the 1976–1981 cohort is near zero and statistically indistinguishable from the pre-war baseline. The 1994–1998 cohort was dropped due to collinearity.

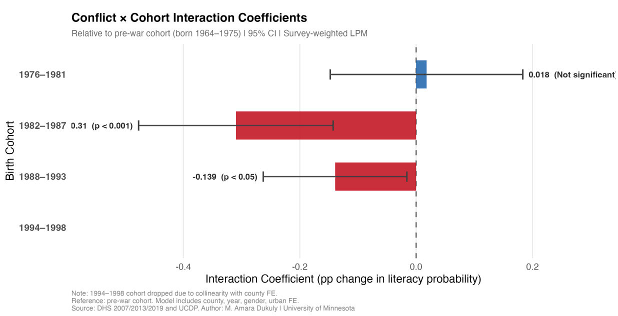

Fig. 1: Conflict × Cohort Interaction Coefficients. Reference: pre-war cohort (1964–1975). 95% CI, survey-weighted LPM.

Cohort main effects are positive and monotonically increasing across the three youngest generations, reflecting net literacy gains over time independent of conflict exposure.

C. Ethnic Heterogeneity in Literacy

The baseline conflict coefficient for the reference group is negative and significant ($\hat{\beta} = -0.054$, $p < 0.01$), consistent with the general hypothesis that conflict reduces literacy.

The Gio exhibit a large positive interaction ($\hat{\tau} = 0.194$, $p < 0.01$), consistent with their historical role as a core constituency of Taylor's NPFL forces and relative protection from direct targeting. The interactions for Mano, Mandingo, and Krahn are all negative and statistically indistinguishable from zero, though the Mandingo exhibited the most negative effect. The lack of significance for these groups likely reflects absorption of county-level variation by county fixed effects.

D. Model Comparison

Across all three specifications, men are substantially more likely to be literate ($\hat{\lambda} = 0.260$, $p < 0.001$), as are urban residents ($\hat{\lambda} = 0.202$, $p < 0.001$). Literacy declined significantly in the 2019 survey wave ($\hat{\delta} = -0.085$, $p < 0.001$). Pseudo-R^2 values are 0.241, 0.240, and 0.242 for Models 1, 2, and 3 respectively. On the basis of AIC, the

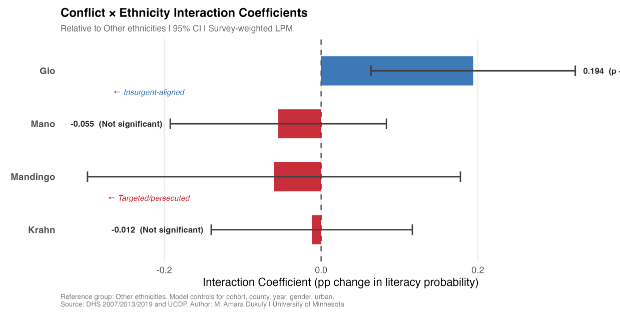

Fig. 2: Conflict × Ethnicity Interaction Coefficients. Reference: Other ethnicities. 95% CI, survey-weighted LPM.

ethnicity model fits marginally better (AIC = 27,699.59) than the cohort (27,704.82) or county (27,715.52) specifications.

Full regression tables are presented in Appendix Tables VI, VII, and VIII below.

VII. Limitations

The most significant limitation is the cross-sectional nature of the DHS data. Because surveys capture county of residence at interview rather than during the war years, and given that an estimated two million Liberians were displaced, there is meaningful potential for county-level exposure misclassification. Restricting to permanent residents addresses this but introduces selection bias, as permanent residents may differ systematically from displaced individuals who likely experienced the highest levels of violence.

The UCDP best estimate records only a fraction of the estimated 250,000 deaths, which may over- or underestimate county-level conflict intensity, particularly for events in rural or poorly documented areas. Third, the design is correlational rather than causal: pre-existing differences in county infrastructure, school quality, and ethnic composition could confound the estimated interactions. Finally, the four focal ethnic groups are minority communities in Liberia, and their small sample sizes reduce statistical power and widen confidence intervals for interaction terms.

VIII. Conclusion

The Liberian Civil Wars spanned fourteen years, displaced nearly two million people, and cost hundreds of thousands of lives. This analysis shows that individuals associated with the lowest literacy rates include women, residents of high-conflict counties such as Lofa, Montserrado, and Bomi, and Liberians born between 1982 and 1987, who were primary school-aged during the peak years of the First Civil War. These results are broadly consistent with the hypothesis that conflict reduced literacy through geographic and cohort channels.

However, results are mixed: some groups exhibited gains during and after the war. The country exhibits higher literacy rates among younger cohorts and counties show notable post-war recovery, suggesting the damage to human capital, while severe and uneven, is not entirely irreversible. The Gio show a positive and significant differential association with conflict exposure not accounted for by county fixed effects, raising questions about within-county variation in schooling access that aggregate analyses would miss.

Future work with finer geographic resolution could help separate the effects of insurgency alignment, displacement, and school destruction more cleanly. Cross-country comparisons with Sierra Leone, which experienced its own civil conflict in this period, could also help assess how much of what is observed in Liberia reflects specific wartime dynamics versus broader patterns common to conflict-affected societies in West Africa.

REFERENCES

[1] ADEA / UNESCO. *Post-Conflict Sierra Leone: Positioning Education Sector Reconstruction for Sustainable Peace*. ADEA Biennial Meeting (2006).

[2] Atkinson, Paul. *RRN Network Paper*, no. 22. Overseas Development Institute (2006). https://www.files.ethz.ch/isn/97730/networkpaper022.pd

[3] Benetech Human Rights Data Analysis Group. *Descriptive Statistics from Statements to the Liberian Truth and Reconciliation Commission* (2013). https://hrdag.org/wp-content/uploads/2013/02/Benetech-TRC-descriptives-final.pdf

[4] Blunch, Niels-Hugo, and Claus C. Pörtner. "Literacy, Skills, and Welfare: Effects of Participation in Adult Literacy Programs." *Economic Development and Cultural Change* 60.1 (2011): 17–66. https://doi.org/10.1086/661219

[5] Center for Global Development. *Liberia Growth Report*. CGDEV. https://www.cgdev.org/sites/default/files/14912_file_Liberia_Growth.pdf

[6] Clément, Jean A. P. "The Economics of Civil War in Sub-Saharan Africa." In *Postconflict Economics in Sub-Saharan Africa: Lessons from the Democratic Republic of the Congo*. International Monetary Fund (2005).

[7] Conteh-Morgan, Earl, and S. Kadivar. "Ethnopolitical Violence in the Liberian Civil War." *Journal of Conflict Studies* 15.1 (1995): 30–44. https://www.erudit.org/en/journals/jcs/1995-v15-n1-jcs_15_1/jcs15_01art02.pdf

[8] DHS Program. "Performing Crosstabulation and Regression Analyses Using DHS Data." Video (2021). https://www.youtube.com/watch?v=MzmS5_jISyY

[9] DHS Program. "Setting the Survey Design in R." Video (2021). https://www.youtube.com/watch?v=pJwd2-m3QBY

[10] Dowd, Caitriona, and Clionadh Raleigh. *The Impact of Civil War on Liberia* (2013). https://rc-services-assets.s3.eu-west-1.amazonaws.com/s3fs-public/CON1222_Accord_23_2.pdf

[11] Hanemann, Ulrike. *Literacy in Conflict Situations*. UNESCO Institute for Education (2005). https://www.researchgate.net/profile/Ulrike-Hanemann/publication/237323833_Literacy_in_conflict_situations/links/5a7af48fa6fdcc772b0960ac/Literacy-in-conflict-situations.pdf

[12] IPUMS. *IPUMS DHS: Harmonized Demographic and Health Surveys Data*. Minnesota Population Center, University of Minnesota (2023). https://doi.org/10.18128/D080.V10.0

[13] Kelly, Jocelyn T. D., Eduardo Colantuoni, Carolyn Robinson, and Michele R. Decker. "Quantifying the Ripple Effects of Civil War: How Armed Conflict Is Associated with More Severe Violence in the Home." *Health and Human Rights* 23.1 (2021): 75–89. https://pubmed.ncbi.nlm.nih.gov/34194203/

[14] Liberia Institute of Statistics and Geo-Information Services (LISGIS). *Liberia Demographic and Health Survey 2007*. DHS Program (2008). https://dhsprogram.com/pubs/pdf/FR201/FR201.pdf

[15] Liberia Institute of Statistics and Geo-Information Services (LISGIS). *Liberia Demographic and Health Survey 2013*. DHS Program (2014). https://dhsprogram.com/pubs/pdf/FR291/FR291.pdf

[16] Liberia Institute of Statistics and Geo-Information Services (LISGIS). *Liberia Demographic and Health Survey 2019–20*. DHS Program (2020). https://dhsprogram.com/pubs/pdf/FR362/FR362.pdf

[17] Liberia Institute of Statistics and Geo-Information Services (LISGIS). *Demographic and Health Survey 2007: DDI Documentation* (2007). https://dhsprogram.com/publications/publication-fr201-dhs-final-reports.cfm

[18] Liberia Institute of Statistics and Geo-Information Services (LISGIS). *Thematic Report on Education and Literacy* (2011). https://lisgis.gov.lr/censusreport/thematic/ThematicReportonEducationandLiteracy.pdf

[19] Liberia Institute of Statistics and Geo-Information Services (LISGIS). *Liberia Demographic and Health Survey 2007: Final Report*. DHS Program (2008). https://dhsprogram.com/pubs/pdf/FR201/FR201.pdf

[20] Liberia Institute of Statistics and Geo-Information Services (LISGIS). *Liberia Demographic and Health Survey 2013: Final Report*. DHS Program (2014). https://dhsprogram.com/pubs/pdf/FR291/FR291.pdf

[21] Liberia Institute of Statistics and Geo-Information Services (LISGIS). *Liberia Demographic and Health Survey 2019–20: Final Report*. DHS Program (2020). https://dhsprogram.com/pubs/pdf/FR362/FR362.pdf

[22] Liu, S. X. "How War-Related Deprivation Affects Political Participation." *Journal of Peace Research* 59.1 (2022): 10–24. https://doi.org/10.1177/00223433211019460

[23] Miguel, Edward, Shanker Satyanath, and Ernest Sergenti. "Economic Shocks and Civil Conflict: An Instrumental Variables Approach." *Journal of Political Economy* 112.4 (2004): 725–753. https://doi.org/10.1086/421174

[24] National Coalition for Literacy. "Key Terms in Adult Literacy." https://nationalcoalitionforliteracy.org/about-adult-literacy/key-terms/

[25] Olaberria, Eduardo, and Carmen M. Reinhart. *The Reversal Problem: Development Going Backwards*. World Bank (2022). https://thedocs.worldbank.org/en/doc/2e421719cc5c8d2e0dc59281dcbdb778-0050012022/original/Reversal-problem-blog-Final.pdf

[26] Radelet, Steven. *Reviving Economic Growth in Liberia*. Center for Global Development (2010). https://www.cgdev.org/sites/default/files/14912_file_Liberia_Growth.pdf

[27] The Advocates for Human Rights. "Background on Liberia and the Conflict." https://www.theadvocatesforhumanrights.org/Res/chapter_4-background_on_liberia_and_the_conflict.pdf

[28] The Center for Justice and Accountability. "Liberia: Country Overview." https://cja.org/where-we-work/liberia/

[29] The DHS Program. *DHS Recode Manual*, Version 5 (2012). https://dhsprogram.com/pubs/pdf/DHSG4/Recode5DHS_23August2012.pdf

[30] The DHS Program. "Analyzing DHS Data." https://dhsprogram.com/data/Guide-to-DHS-Statistics/Analyzing_DHS_Data.htm

[31] World Bank. *Liberia Demographic and Health Survey 2007*. World Bank Microdata Library (2007). https://microdata.worldbank.org/index.php/catalog/1423

[32] World Bank. *Liberia Demographic and Health Survey 2013*. World Bank Microdata Library (2013). https://microdata.worldbank.org/index.php/catalog/2165

[33] World Bank. *Liberia Demographic and Health Survey 2019–2020*. World Bank Microdata Library (2020). https://microdata.worldbank.org/index.php/catalog/3896

[34] Yele, and Specioaza. "Decoding Barriers and Failures in Liberia's Educational Policy Implementation." *World Journal of Advanced Research and Reviews* 23.1 (2024). https://doi.org/10.30574/wjarr.2024.23.1.2009

IX. Appendix

TABLE VI: Model 1 — County × Conflict Interactions

Variable	**Est.**	**Sig.**	**SE**
Intercept	0.074	*	0.035
Conflict (Ref. County: Nimba)			
Total Conflict Share	0.078		0.052
County × Conflict			
× Bomi	-0.309	***	0.073
× Bong	-0.182	**	0.058
× Gbarpolu	-0.041		0.083
× Grand Bassa	-0.193	**	0.067
× Grand Cape Mount	-0.201	*	0.082
× Grand Gedeh	-0.097		0.076
× Grand Kru	-0.058		0.088
× Lofa	0.000		0.124
× Margibi	-0.074		0.058
× Maryland	-0.054		0.078
× Montserrado	-0.201	**	0.066
× River Cess	-0.169	**	0.060
× River Gee	-0.419	***	0.114
× Sinoe	-0.104	.	0.058
County Main Effects (Ref: Nimba)			
Bomi	0.214	***	0.061
Bong	0.011		0.049
Gbarpolu	-0.040		0.053
Grand Bassa	0.021		0.050
Grand Cape Mount	0.082		0.053
Grand Gedeh	0.072		0.059
Grand Kru	0.070		0.057
Lofa	-0.045		0.071
Margibi	0.086	.	0.050
Maryland	0.116	.	0.064
Montserrado	0.331	***	0.050
River Cess	0.068		0.048
River Gee	0.240	**	0.080
Sinoe	0.069		0.043
Birth Cohort (Ref: pre-1976)			
Cohort 1976–1981	0.009		0.014
Cohort 1982–1987	0.091	***	0.015
Cohort 1988–1993	0.113	***	0.012
Cohort 1994–1998	0.090	***	0.017
Survey Year (Ref: 2007)			
Year 2013	0.012		0.015
Year 2019	-0.085	***	0.017
Individual Controls			
Gender (Male = 1)	0.260	***	0.009
Urban	0.202	***	0.019
N=23,844 R^2=0.241 AIC=27715.52			

*** $p<0.001$, ** $p<0.01$, * $p<0.05$, . $p<0.1$
Survey-weighted GLM (Gaussian).

TABLE VII: Model 2 — Cohort × Conflict Interactions

Variable	Est.	Sig.	SE
Intercept	0.047		0.035
Conflict (Ref. Cohort: pre-1976)			
Total Conflict Share	0.123	*	0.058
Cohort × Conflict			
× Cohort 1976–1981	0.018		0.084
× Cohort 1982–1987	-0.310	***	0.085
× Cohort 1988–1993	-0.139	*	0.063
× Cohort 1994–1998	*(omitted—collinear)*		
Cohort Main Effects (Ref: pre-1976)			
Cohort 1976–1981	-0.033		0.051
Cohort 1982–1987	0.275	***	0.057
Cohort 1988–1993	0.179	***	0.033
Cohort 1994–1998	0.139	***	0.023
County FE (Ref: Nimba)			
Bomi	0.070	.	0.039
Bong	-0.063	*	0.030
Gbarpolu	-0.057	.	0.029
Grand Bassa	-0.072	*	0.029
Grand Cape Mount	-0.004		0.030
Grand Gedeh	0.032		0.035
Grand Kru	0.065	*	0.027
Lofa	-0.033		0.031
Margibi	0.054		0.040
Maryland	0.105	**	0.034
Montserrado	0.232	***	0.034
River Cess	0.007		0.029
River Gee	-0.005		0.034
Sinoe	0.030		0.028
Survey Year (Ref: 2007)			
Year 2013	0.011		0.015
Year 2019	-0.085	***	0.017
Individual Controls			
Gender (Male = 1)	0.259	***	0.009
Urban	0.200	***	0.019
N=23,844 R^2=0.240 AIC=27704.82			

*** $p<0.001$, ** $p<0.01$, * $p<0.05$, . $p<0.1$
Survey-weighted GLM (Gaussian).

TABLE VIII: Model 3 — Ethnic Group × Conflict Interactions

Variable	Est.	Sig.	SE
Intercept	0.129	***	0.029
Conflict (Ref: All Other Groups)			
Total Conflict Share	-0.054	**	0.021
Ethnicity × Conflict			
× Gio	0.194	**	0.067
× Mano	-0.055		0.070
× Mandingo	-0.060		0.121
× Krahn	-0.012		0.065
Ethnicity Main Effects (Ref: Other)			
Gio	-0.072		0.046
Mano	0.018		0.051
Mandingo	-0.070		0.075
Krahn	0.053		0.051
Birth Cohort (Ref: pre-1976)			
Cohort 1976–1981	0.009		0.014
Cohort 1982–1987	0.086	***	0.014
Cohort 1988–1993	0.109	***	0.012
Cohort 1994–1998	0.092	***	0.016
County FE (Ref: Nimba)			
Bomi	0.067		0.044
Bong	-0.069	.	0.036
Gbarpolu	-0.036		0.033
Grand Bassa	-0.063	.	0.034
Grand Cape Mount	-0.010		0.035
Grand Gedeh	-0.004		0.047
Grand Kru	0.055	.	0.033
Lofa	-0.002		0.035
Margibi	0.057		0.043
Maryland	0.104	**	0.039
Montserrado	0.241	***	0.037
River Cess	-0.004		0.035
River Gee	0.011		0.038
Sinoe	0.027		0.034
Survey Year (Ref: 2007)			
Year 2013	0.011		0.015
Year 2019	-0.084	***	0.017
Individual Controls			
Gender (Male = 1)	0.262	***	0.009
Urban	0.205	***	0.019
N=23,844 R^2=0.242 AIC=27699.59			

*** $p<0.001$, ** $p<0.01$, * $p<0.05$, . $p<0.1$
Survey-weighted GLM (Gaussian).

Regional Integration and Trade Shocks

Luke Tsuchiya

Tufts University

Abstract— **This paper examines the effects of trade liberalization on different communities and the importance of regional integration for calculating the effect of trade shocks. I test whether there is treatment effect heterogeneity of changes in Chinese import exposure for counties with higher levels of non-local retail sales per capita in the pre-shock period. Import shocks lower median income growth but do so more for counties with lower levels of pre-shock retail sales. This effect and attenuation with retail sales is present in both manufacturing and non-manufacturing industries.**

1. Introduction

There has been extensive research and debate about the effects of China's trade liberalization on U.S. regional economies. For example, Autor et al (2013) find that increased exposure to import competition from China significantly impacted employment in both manufacturing and non-manufacturing industries. Interestingly, the authors find little evidence that people negatively affected by trade exposure relocate to places where job opportunities are growing. Rather, there is evidence that the displaced workers do not migrate and become unemployed or leave the workforce entirely. A related strand of literature examines why the China shock has been so regionally concentrated and persistent. For example, Eriksson et al (2021) find that the effect of import exposure was greater in areas with low education levels, where we would expect to find a lower ability of workers to change industries. This could contribute to the regional concentration of the effect as the shock hit areas least able to adapt. If there was little resulting out-migration towards better job opportunities, this could also increase the duration of the effect as people remain in places with poor employment prospects. Another possible explanation that I put forward in this paper, is that the China shock hit areas that had less non-local demand for local retail goods. Essentially, if the trade shock lowers local demand by reducing employment in manufacturing industries with no resulting increase in employment in non-manufacturing industries, areas with fewer outside sources of demand for local retail will be particularly susceptible to declines in manufacturing employment spilling over to declines in non-manufacturing incomes. This would contribute to the regional concentration of the China shock and potentially its persistence if these declines in non-manufacturing industries similarly result in low sector reallocation or migration.

More specifically, I will consider variation in retail sales that are associated with distance to an interstate highway and test whether there is treatment effect heterogeneity of changes in Chinese import exposure over the period 2000-2007. This is in part motivated by Anderson and Matsa (2011) who document that communities located closer to interstate highways receive more traffic through fast food restaurants than those just slightly further away. The portion of retail sales associated with highway distance is most plausibly the portion that comes from non-local demand, so using highway distance as an instrument for retail sales should allow me to test the importance of specifically non-local demand in determining the effect of the China shock on income growth.

The primary contributions of this paper are to document variations in the effect of trade liberalization on different communities and the importance of regional integration for calculating the effect of trade shocks. The rest of the paper proceeds as follows: section 2 summarizes other relevant literature, section 3 describes the data, section 4 lays out my empirical strategy, section 5 discusses the results, and section 6 concludes.

2. Literature Review

A number of papers document the relative poor performance of highly trade exposed regions after large trade shocks in numerous different countries (e.g. India: Topalova (2010), U.S.: Autor et al (2013) and Kovak (2013), Finland: Costinot et al (2022)) The following literature review is a select summary of a subset of these papers, with particular focus on explanations of possible mechanisms driving the geographic concentration and persistence of the effects of the China shock in particular. It is intended to motivate what follows in the paper.

Autor et al (2013) document three facts about the effect of Chinese import exposure between 1990 and 2007. First, there is little migration response to declines in employment. There is also little increase in employment in non-manufacturing industries. There is however a resulting increase in labor force non-participation and unemployment. This parallels the findings of Foote et al (2015), who show that in the aftermath of mass layoffs after the 2008 recession, movement to labor force non-participation was a significant and growing channel of labor adjustment. In their context, though, they show that migration was still the primary channel. This may reflect the potential difference in willingness, ability, or incentive to migrate for workers affected by trade versus the financial

I would like to thank Professor Jeff Zabel and Adam Storeygard for their invaluable comments and advice

crisis. Second, Autor et al (2013) find declines in wages across all education levels, but the effect is twice as large for non-college educated workers. Third, declines in wages are concentrated among non-manufacturing workers. This would all be consistent with import exposure reducing the demand for non-manufacturing goods from workers who can no longer find work in manufacturing industries.

Autor et al (2025) provides a very detailed follow up to this, decomposing the net change in employment in manufacturing and non-manufacturing into 7 different categories (Aging in and out of the workforce, migrating into or out of a commuting zone, moving into or out of unemployment, and changing sectors). Considering the longer run effect of the China shock from 2000 to 2019, they show that although there were reductions in manufacturing employment and reductions in overall employment prior to 2010, strong growth in non-manufacturing employment after 2010 more than offset the reductions in employment in manufacturing. The demographic groups that move into non-manufacturing industries, however, are distinct from the demographic groups that move out of manufacturing. Those who become employed in non-manufacturing industries are mostly foreign born or native born Hispanic, while the groups moving out of manufacturing are native born white and black workers. Additionally, the majority of the decline in manufacturing employment prior to 2008 is due to reductions in the number of people aging into the workforce and the number of people moving to a commuting zone to work in manufacturing. This would still be consistent with how I've described declines in manufacturing employment leading to declines in non-manufacturing income. Rather than reductions in spending from laid off workers playing the most important role, it could be that in areas with high trade exposure, people who would have aged or moved into manufacturing jobs no longer have this opportunity. Spending on local non-manufacturing goods by the cohort that would have moved into the CZ will disappear completely and spending from the cohort that would have aged into the workforce will decrease if they either never enter the workforce and remain unemployed or they take lower paying jobs in non-manufacturing industries. As this group employed in manufacturing gradually makes up a smaller share of the population, demand for non-manufacturing goods could dry up.

There are a number of papers studying why the effect of the China shock was more geographically concentrated and persistent than expected. Glaeser (2021) provides a good review of this literature that I will briefly summarize here. He notes four main explanations for the geographic concentration and duration of the effect of China shock. First, the shock itself was persistent. That is the China shock was not a one off break in equilibrium but rather a continuous change in the comparative advantage of China to a wider range of manufacturing industries. Suggestive of this, Autor et al (2013) show that China's revealed comparative advantage in manufacturing continues to rise through 2012. Second, low levels of human capital in industries most exposed to rising imports from China may have limitted people's ability to adapt to the shock. For example, Eriksson et al (2021), who as I noted previously, show that the effect of import exposure was greater in areas with low education levels where we would expect to find a lower ability of workers to change industries. In the context of the findings of Autor et al (2025), the important mechanism may not be the inability of workers to change industry after being laid off, but of younger generations who would have worked in manufacturing to find employment outside of manufacturing without a college degree. Third, there might exist significant barriers to adaptation. For example, Kleiner and Kruger (2010) point to the rise of occupational licensing making it difficult to switch industries of find an identical job in a different state due to the need for a new license. Fourth, there may also be other barriers to migration that would reduce people's ability to find a job in a new location. Glaeser and Gyourko (2005) note that because of the durability of housing supply, places with declining productivity, like Detroit, often have very low housing costs. This means that, as pointed out in Horn et al (2024), the return to migrating to areas with higher wages have been declining over time as housing prices have risen in high productivity places and fallen in low productivity ones.

In addition to the explanations cited in Glaeser (2021), I'm aware of two papers that could provide additional explanations for the geographic concentration and or the persistence of the effect of the China shock. Adao et al (2020) show in a spatial general equilibrium model that there are large spillover effects of trade shocks between spatially linked commuting zones. This would mean if the China shock was regionally concentrated, as we will see it was, that counties in regions with higher average changes in import exposure could experience larger total effects from the import shock. This would contribute to the geographic concentration of the effect of the effect of the China shock. In a different area of research, studying the effect of the 2007-2017 coal shock, Krause (2024) shows that counties that had previously experienced declines in educated workers due to an exogenous negative labor demand shift in the 1980s were less resilient to the contemporary coal shock. If true in the context of trade shocks and their effect on migration of educated workers, this could imply that the China shock reduced the resilience of highly affected areas leading to poor economic performance in the face of future shocks. This would increase both the geographic concentration of the shock and its duration.

This paper builds off this literature by considering, in an economic landscape with low returns and high barriers to migration, how non-local demand for retail goods may have impacted the effect of import exposure, and whether the distribution of retail sales is such that counties that were least able to adapt to import competition were also the ones with the highest levels of exposure.

3. Data

There are 6 sources of data used for this project. The analysis in this paper is done at the county level. From the

Fig. 1: Spatial Distribution of Chinese Import Shocks

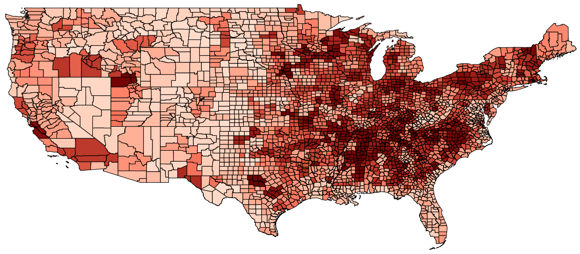

Census Bureau's County Business Patterns (CBP) database, I obtain data on the number of establishments, employees, and annual payroll by NAICS code at the county level between 2000 and 2007. I use the Bureau of Transportation's measure of county land area and I take the measures of trade exposure ($ change in imports from China/worker) at the county level over the period 2000-2007 from Autor et al (2020).

The measure of distance to the nearest highway and urban center comes from Anderson and Matsa (2011) and is provided at the zip code level (i.e. distance to nearest highway exit or urban center from the centroid of a zip code). To match these with county level observations, I take the highway distance of the most populated zip code within a county. In Appendix 1.1, as a robustness check, I repeat the model estimation instead using the minimum highway distance of any zip code within a county. The results are qualitatively similar when including fixed effects, although not statistically significant. Because I'm taking the minimum distance of any zip code within a county, the highway distance measure becomes smaller on average and heavily left skewed. To account for this I restricted the sample to counties with a distance measure less than 6 miles. This seems to correspond with where highway distance is no longer negatively correlated with pre-shock retail sales per capita, but this results in a larger sample than when restricting the original distance measure to less than 16 miles (1,558 v.s. 1308).

County level data on median household income are sourced from the Census Bureau's Small Area Income and Poverty Estimates Program (SAIPE). This data is available for the following years: 1989, 1993, 1995, & 1997-2007. Finally, the measure of pre-shock retail sales per capita come from the census bureau's 1992 economic census of retail trade. To visualize the spatial distribution of import competition, I plot the increase in imports per worker over the 2000-2008 period on a map of U.S. counties in Figure 1.

There are 3,111 counties for which I have import exposure data from Autor et al (2013). To run my empirical specification looking at the effect of trade exposure on median incomes I merge data from the SAIPE and geographic data from Anderson and Matsa (2011). After merging highway and urban distance data, there are 3,107 counties for which I have complete exposure, SAIPE, and geographic data. Due to the larger size of counties in the western U.S. making the

TABLE I: Summary Statistics

	Min	Median	Max	Mean	SD
Panel A. Non-West Census Region Sample					
(Autor et al 2020) Import Exposure Measure ($1000/Worker)	-0.6	0.9	7.2	1.1	0.8
1992 Retail Sales per capita ($1,000)	0.1	5.1	39.6	5.4	2.7
Average Annual Percent Change in Real Median Income 2000-2007	-2.7	0.2	2.7	0.2	0.6
Distance to Interstate (Miles)	0.0	16.0	177.0	22.2	22.8
Distance to Urban Area (Miles)	0	27.9	207.2	36.2	35.2
Land Area (Sq Miles)	2.5	572.3	6671.1	665.7	470.6
Population Density (People/Sq Mile)	0.6	44.6	3140.8	140.9	329.1
1992 Population	465	23,477	2,998,398	64,979	148,180
Observations	2617				
Panel B. Counties Within 16 Miles of an Interstate					
(Autor et al 2020) Import Exposure Measure ($1000/Worker)	-0.0	0.9	5.1	1.1	0.7
1992 Retail Sales per capita ($1,000)	0.1	5.7	30.5	6.0	2.7
Average Annual Percent Change in Real Median Income 2000-2007	-1.8	0.2	2.7	0.2	0.6
Distance to Interstate (Miles)	0.0	4.5	16.0	5.8	4.7
Distance to Urban Area (Miles)	0	15.9	201.5	22.0	26.8
Land Area (Sq Miles)	2.5	535.5	6248.1	601.6	424.5
Population Density (People/Sq Mile)	0.6	81.1	3140.8	236.7	436.9
1992 Population	1161	40,232	2,998,398	106,035	199,162
Observations	1308				

meaning of distance to nearest highway hard to interpret, I drop counties in the west census region. I also drop counties that have 0 recorded retail sales in 1992 and counties with population density above the 99th percentile as population density is heavily right skewed, leaving 2,617 observations at the county level.

Panel A of Table 1 provides summary statistics for this sample. The average annual growth rate in real median income (Base year: 2000) between 2000 and 2007 is 0.2% but ranges from -2.7% to 2.7%. On average, counties are 22.2 miles from an interstate exit and 36.2 miles from an urban center. 1992 retail sales per capita are $5,400 on average but are right skewed so in my main specification I use the log of sales. Additionally, in my main specification, I restrict the sample to counties with a highway distance measure less than 16 miles, leaving 1308 observations. The reason for this is that after a certain point, distance to an interstate is no longer negatively correlated with pre-shock retail sales per capita, meaning for counties beyond this distance, variation in local demand is no longer being driven by outside sources of demand from a highway. Figure 2 shows a binned scatter plot of log retail sales per capita versus highway distance. We can see that the slope seems to change after 16 miles, becoming positive but close to zero. Because the exact choice of a cutoff is arbitrary, in Appendix 1.2 I include coefficient estimates for the three key explanatory variables estimated with different maximum highway distance cutoffs. Results are very similar within a 6 mile window around the 16 mile cutoff.

Counties in this within 16 miles sample are closer to urban centers (22.0 v.s. 36.2 miles), with more people (106,035 v.s. 64,979), and have higher retail sales per capita ($6,000 v.s. $5,400) on average. Summary statistics for this sample can be found in Panel B of Table 1.

To help control for noise in the data, and to ensure comparison of more similar (i.e. closer) counties, I create two types of cluster fixed effects variables. The first employs a spectral clustering algorithm as in Luxbourg (2007) to separate counties into 378 non-overlapping areas. This algorithm performs well with larger cluster sizes but can lead to weird results when the number of clusters is increased. For example it may create long and skinny clusters, or many clusters with only 1 county. In order to test smaller clusters, I also employ

Fig. 2

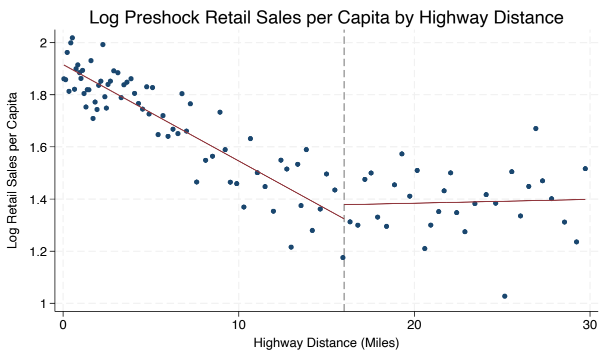

a different clustering algorithm that aims to create smaller but still compact clusters. The algorithm is as follows:

1) Consider the set of counties $\boldsymbol{C} = \{i \mid county\ i \in NonWest\ U.S.\}$
2) Choose an arbitrary county i_1
3) Cluster 1 is the subset $\boldsymbol{C_1} \subset \boldsymbol{C} = \{j \mid county\ j\ borders\ county\ i_1\}$
4) Choose another arbitrary county $i_2 \notin \boldsymbol{C_1}$
5) Cluster 2 is the subset $\boldsymbol{C}_2 \subset \boldsymbol{C} = \{j \mid county\ j \in \boldsymbol{C} \setminus \boldsymbol{C_1}\ \&\ county\ j\ borders\ county\ i_2\}$
6) ...
7) Choose another arbitrary county $i_{n+1} \notin \bigcup_{k=1}^{n} \boldsymbol{C_k}$
8) Cluster n+1 is the subset $\boldsymbol{C}_{n+1} \subset \boldsymbol{C} = \{j \mid county\ j \in \boldsymbol{C} \setminus \bigcup_{k=1}^{n} \boldsymbol{C_k}\ \&\ county\ j\ borders\ county\ i_{n+1}\}$
9) Repeat the previous step until $\boldsymbol{C} \setminus \bigcup_{k=1}^{N} \boldsymbol{C_k} = \emptyset$, that is until there are N clusters and all counties are in a cluster

For the results reported in this paper, I choose the unclustered county with the smallest fips code to determine each new cluster. Comparing, the spectral clusters with the border clusters, the border clusters have a smaller number of counties on average, 4.38 v.s. 7.65, and a smaller maximum number of counties, 10 v.s. 19. Figure 3 gives a visual comparison of these clusterings for both the sample of non-western counties and the sample of counties with highway distance measure less than 16 miles. When included in my econometric model, the border clusters do not change the qualitative results compared to either the model without cluster fixed effects or the model with spectral cluster fixed effects, but they are effective at explaining more of the variation in income growth, with R-squared values of 0.62 v.s. 0.08 or 0.49.

To examine potential mechanisms driving my results, I consider the same effects of trade exposure and pre-shock retail sales for NAICS defined manufacturing, retail, and accommodation and food industries. For each industry, I consider 4 economic outcomes, growth in employees, payroll, establishments, and pay per employee over the 2000-2007 period. Table 2 presents the summary statistics for these variables for the subsample of counties within 16 miles of an interstate highway. The data are also restricted to remove outliers below the 1^{st} and above the 99^{st} percentiles. Dollar

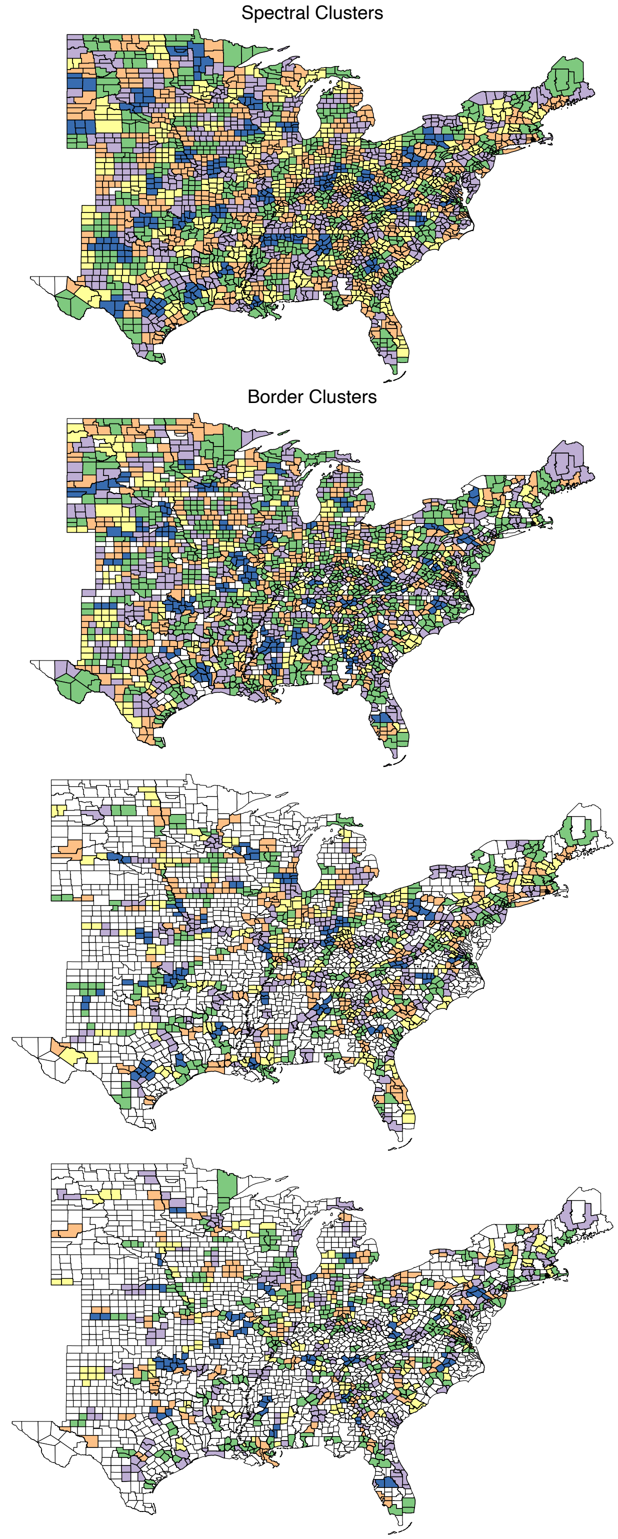

Fig. 3: Maps of clustering results. The bottom maps show only counties that are in the within 16 miles of an interstate sample and have at least 2 counties in their cluster.

TABLE II: Mechanisms Summary Statistics

Manufacturing	Min	Median	Max	Mean	SD
%Δ Employees	-1	-0.341	3.500	-0.290	0.576
%Δ Payroll	-1	0.015	8.880	0.172	1.033
%Δ Establishments	-0.667	0.076	3	0.115	0.214
%Δ Pay	-0.236	0.459	1.827	0.495	0.273
Retail	Min	Median	Max	Mean	SD
%Δ Employees	-1	-0.264	1.739	-0.274	0.318
%Δ Payroll	-1	0.281	4.121	0.253	0.477
%Δ Establishments	-0.353	0.091	0.868	0.110	0.148
%Δ Pay	-0.339	0.465	1.909	0.473	0.173
Accommodation and Food	Min	Median	Max	Mean	SD
%Δ Employees	-1	0.039	4.357	0.080	0.436
%Δ Payroll	-1	0.576	8.753	0.643	0.655
%Δ Establishments	-0.500	0.193	2	0.239	0.275
%Δ Pay	-0.365	0.429	1.656	0.450	0.218

values are in 2000 real terms. The average growth rates for the number of employees in manufacturing and retail are both negative with employment declining by 29.0% and 27.4% in manufacturing and retail respectively between 2000 and 2007. Employment in accommodation and food grew on average by 8.0% over the period. All other values are growing on average over the period.

4. Empirical Specification

The main goal of the empirical model in this paper is to test whether there are heterogeneous treatment effects of the China shock for counties with higher levels of non-local retail sales per capita in the pre-shock period (since non-local sales can't be measured, I proxy for them by instrumenting retail sales with highway distance). In specifying a model of the effect of imports from China, I follow Autor et al (2013) and the literature generally in looking at how the import shock, measured as the sum of changes in imports from China for each manufacturing industry weighted by a county's share of workers in that industry, affects the growth rates of economic outcomes like median income and employment. This is opposed to the effect of either the level of imports on the level of median incomes, the level of imports on the change in incomes, or the change in imports on the level of income. The motivation for regressing Δs on Δs is twofold. First, it requires stronger assumptions to interpret the results from the regression of the level of post-period incomes on either the level of import exposure or the change in import exposure, than it would for a regression using the change in incomes. For the former, we would require that incomes conditional on import exposure are identical in the pre-period, and the counterfactual growth rates over the period in absence of any shock are the same. For the latter, we only need that the counterfactual growth rates over the period are the same. For this reason, I use the change in median income as the dependent variable in my empirical specification. Second, I use the change in imports per worker because the level of imports per worker may not have a causal relationship with the change in incomes. This would be true for example if the level of imports rose in the pre-period and the market had fully adjusted to a new equilibrium, ie there would be no more adjustment of incomes to this level of imports. The change in imports per worker over the period, however, is a new shock that, assuming no anticipation effect, has not already been incorporated into the market equilibrium. This is why I include changes in imports per worker in the empirical specification for this paper.

To study the heterogeneous effect of imports along the dimension of non-local demand, I use the level of retail sales per capita (instrumented using a county's distance to an interstate to capture the part of retail sales that can be considered non-local) in the pre-shock period (1989), instead of the change in retail sales over the period . This is because the mechanism through which I hypothesize non-local demand to create differences in the effect of the China shock relies on the level of non-local demand, not the change in non-local demand. For example, consider a county with many layoffs due to closing manufacturing plants. This will lead to a reduction in demand for local non-manufacturing goods and services. If there is a high level of non-local demand for these goods and services, the reduction in demand will be smaller proportionally resulting in fewer layoffs for non-manufacturing workers. What matters is that there exists significant non-local demand, not that it changes or how it changes. Additionally, by using the level of retail sales I avoid any issue of bad controls in my model.

Structural Model:

$$\begin{aligned}\Delta Y_i = \beta_0 + \beta_1 \cdot Shock_i + \beta_2 \cdot ln(\frac{Preshock\ Retail\ Sales}{Capita})_i \\ + \beta_3 \cdot Shock_i \cdot ln(\frac{Preshock\ Retail\ Sales}{Capita})_i + \vec{\delta_i} + \mu_i\end{aligned} \tag{1}$$

First Stage Regressions:

$$\begin{aligned}Z_i = \alpha_0 + \alpha_1 \cdot Foreign\ Shock_i + \alpha_2 \cdot Highway\ Distance_i \\ + \alpha_3 \cdot Foreign\ Shock_i \cdot Highway\ Distance_i + \vec{\phi_i} + \varepsilon_i\end{aligned} \tag{2}$$

Equation 1 is the main empirical specification for this paper where counties are indexed by i. The trade shock measure (*Shock*), pre-shock retail sales per capita, and their interaction term are instrumented using lagged trade shock measures from other countries (primarily in Europe, as in (Autor et al, 2020)), highway distance, and the interaction between the lagged shock and highway distance measure, as in Equation 2 with Z_i the variable being instrumented. $\vec{\delta_i}$ and $\vec{\phi_i}$ are both vectors of either spectral-cluster or county-border-cluster fixed effects. μ_i and ε_i are error terms. ΔY_i is the average annual percent change in median income over the 2000-2007 period.

The primary coefficient of interest is β_3, which tells us how the effect of import exposure changes with retail trade that is associated with highway distance. My hypothesis is that β_1 should be negative and β_3 should be positive, meaning that the trade shock lowers income growth but this effect is mitigated by higher demand for retail goods and services from highway traffic. The intuition is that jobs in non-manufacturing industries are dependent on local and non-local demand for goods and services. So, when employment in manufacturing falls and, spending decreases because people move to unemployment or out of the labor force entirely rather than finding jobs in new places or industries, workers in non-manufacturing may also lose their jobs because of falling demand for the goods and services they produce. If, however, a county has a higher share of demand for local retail goods and services from non-local sources, for example because of its proximity to a highway, workers in non-manufacturing industries may be less likely to lose their jobs due to spillovers from negative manufacturing shocks.

For each explanatory variable and their interaction, a unique set of assumptions is required for regression results to be considered causal. For all variables, we require that they are significantly associated with the set of instrumental variables (relevance). This of course can be tested by looking at the results from the first stage regressions. The smallest F-statistic for the joint significance of the excluded instruments is 29.44, well above the Stock-Yogo weak instrument cut off of 10. Additionally, the exclusion restriction must hold so the instruments cannot affect median income growth through other channels than changes in the endogenous variables. This seems likely for foreign import exposure, although it's possible that if Chinese exports are substitutes for American exports, then increases in Chinese imports per worker in foreign countries might cause decreases in demand for American exports. To the extent that substitution between American and Chinese goods in foreign countries is due to changes in American productivity or competitiveness, my empirical specification would be misinterpreting this as the effect of increases in Chinese exports. Put simply, it would be conflating changes in demand for American goods with changes in supply of Chinese goods. In the next section, I look at trends in income growth for counties with high and low levels of import exposure and find that they are on almost identical paths pre-2000. This is reassuring as it allows us to rule out changes in demand for U.S. products that do not occur prior to increases in the supply of Chinese exports. However, changes in demand for U.S. production may still be correlated with increases in supply of Chinese exports to foreign countries after 2000, potentially violating the exclusion restriction.

The exclusion restriction assumption for distance to an interstate highway seems less plausible. Counties closer to interstates likely differ along many dimensions other than retail sales, which could affect median income growth. This would also likely violate the independence assumption as these differences could exist both because highway distance has a causal effect on other variables which are important determinants of income growth (exclusion), and because the location of highways is not chosen randomly (independence). For this reason, correlations between income growth and pre-shock retail sales associated with highway distance should not be considered causal. The highway distance instrument, however, is still useful as it restricts the variation in retail sales to variation that is primarily from non-local sources.

Similarly, the exclusion restriction and independence assumptions may be violated for the interaction term. That is the difference in the effect of import exposure between counties with lower and higher levels of pre-shock retail sales may not be caused by this difference in retail sales. If, for instance, re-skilling programs were targeted towards counties with lower income growth and these programs were effective at getting people back into new jobs after being laid off in manufacturing, we would expect this to cause counties with higher retail sales (which also have lower income growth on average) to be more resilient in the face on import competition. Since I have neither evidence for or against the presence of some factor like this, and because there is clearly some pre-shock difference between counties with low levels of retail sales and counties with high levels of retail sales such that counties with high levels of retail sales have lower income growth on average, the results for the relationship between pre-shock retail sales and the size of the effect of import exposure should be considered with caution. That being said, this relationship can still tell us something important about the nature of the China shock if as I will show, the China Shock was felt hardest in areas with low levels of pre-shock retail sales and these counties had the largest reduction in median income growth due to the shock. This would add to an understanding of why the effect of the China Shock was relatively large and geographically concentrated. The next section details these results.

5. Results

A. Primary Regressions

Table 3 provides coefficient estimates for 6 versions of the model, with different combinations of fixed effects and controls, as well as the estimates from the regression of median income growth on only the trade shock for comparison with the results from Autor et al 2013. The results for this specification, listed in column 1 of Table 3, indicate that on average a $1,000 increase in imports per worker is associated with a 0.148 pp decline in the average annual growth rate of real median household income. Compounded over a ten year period, this is equivalent to a 1.49 pp decline in median household income. This is very close to the estimate from Autor et al 2013 of a 1.73 pp decline in median household income over a ten year period due to a $1000 increase in imports per worker. The results are not expected to be identical as I'm using a subsample of counties and considering primarily the period 2000-2007, but it is still encouraging that they are so similar.

Including proximity to an urban center and urban density does not seem to affect coefficient estimates, especially after controlling for border fixed effects. Another valid specification would be to use land area and the log of pre-shock population as controls instead of urban density. I include these results in Appendix 1.3. Most coefficients are similar to those in Table 3, with a notable exception that the interaction coefficient is much smaller after adding controls to the model with spectral cluster fixed effects. There is little

TABLE III: Regression Results

	%ΔMed Inc/Year						
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Shock	-0.148***	-0.404	-0.339	-0.385*	-0.283	-0.569***	-0.527**
	(0.0313)	(0.291)	(0.295)	(0.233)	(0.235)	(0.220)	(0.218)
Log Preshock Retail Sales		-0.588***	-0.524**	-0.636***	-0.436**	-0.872***	-0.772***
		(0.216)	(0.246)	(0.181)	(0.216)	(0.181)	(0.206)
Shock x Sales		0.138	0.0952	0.170	0.0972	0.268**	0.242*
		(0.174)	(0.177)	(0.139)	(0.143)	(0.131)	(0.131)
Constant	0.350***	1.371***	1.330***	1.055***	0.814**	1.953***	1.819***
	(0.0371)	(0.369)	(0.408)	(0.332)	(0.366)	(0.335)	(0.365)
Urban Distance Control	No	No	Yes	No	Yes	No	Yes
Population Density Control	No	No	Yes	No	Yes	No	Yes
Spectral Cluster FE	No	No	No	Yes	Yes	No	No
Border Cluster FE	No	No	No	No	No	Yes	Yes
Shock F-Stat	1360.03	455.15	449.82	254.20	238.74	179.45	157.88
Sales F-Stat		61.77	38.73	43.86	27.25	32.64	23.25
Shock x Sales F-Stat		297.63	296.59	145.33	148.32	98.60	94.12
Observations	1,308	1,308	1,308	1,308	1,308	1,308	1,308
R-squared	0.028	0.079	0.094	0.496	0.523	0.625	0.643

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Reported F-Statistics are from the joint significance test of excluded instruments in the first stage regressions.

difference after including controls in the model with border fixed effects.

Overall, the results across different specifications seem to be consistent with median income growth being negatively associated with import exposure over most of the range of pre-shock retail sales, and this association attenuating with increases in retail sales. Table 3 evaluates these results in terms of economic significance.

Here, I will report on the results from column 6 of Table 3 which I consider the primary specification of my econometric model with the set of county-border fixed effects. At the mean of the log of pre-shock retail sales per capita, a \$1000 increase in imports per worker is associated on average with a 0.115 pp decline in annual average percent median income growth, all else equal. This is equivalent to a 0.143 standard deviation decline in income growth for each standard deviation increase in import exposure. This is consistent with my hypothesis and the results from Autor et at 2013. The coefficient on the interaction term tells us that for every 1 unit increase in the log of pre-shock retail sales per capita (equivalent to a $((e-1)*100\% \approx 172\%$ increase in the level of pre-shock retail sales per capita), the association between median income growth and the import shock increases (becomes more positive) by 0.268 pp. Equivalently, for every $\approx 1\%$ increase in pre-shock retail sales, the association between median income growth and the import shock increases (becomes more positive) by 0.00268 pp. Specifically, this is an increase in the log of pre-shock retail sales that is correlated with distance to an interstate highway.

To consider the economic significance of this result, I evaluate the change in the effect of import exposure due to a non-marginal 100% increase in pre-shock retail sales per capita (equivalent to an increase in the log of pre-shock retail sales per capita by ln(2)) relative to the effect of import exposure at the mean of pre-shock retail sales per capita[1].

[1] $\eta_{Shock*Sales} = \frac{ln(2)*\beta_3}{\beta_1+\mu_{ln(Sales)}*\beta_3}$

TABLE IV: Economic Significance

Variable	Model	Measure	Value
Shock	1	Std Coef	-0.193
Shock	2	Std Coef (μ_{sales})	-0.225
Retail Sales	2	$SD/100\%\Delta_{sales}(\mu_{shock})$	-0.471
Shock * Sales	2	$\%\Delta\beta_1/100\%\Delta_{sales}(\mu_{sales})$	-0.151
Shock	3	Std Coef (μ_{sales})	-0.232
Retail Sales	3	$SD/100\%\Delta_{sales}(\mu_{shock})$	-0.508
Shock * Sales	3	$\%\Delta\beta_1/100\%\Delta_{sales}(\mu_{sales})$	-0.051
Shock	4	Std Coef (μ_{sales})	-0.124
Retail Sales	4	$SD/100\%\Delta_{sales}(\mu_{shock})$	-0.543
Shock * Sales	4	$\%\Delta\beta_1/100\%\Delta_{sales}(\mu_{sales})$	-0.984
Shock	5	Std Coef (μ_{sales})	-0.127
Retail Sales	5	$SD/100\%\Delta_{sales}(\mu_{shock})$	-0.537
Shock * Sales	5	$\%\Delta\beta_1/100\%\Delta_{sales}(\mu_{sales})$	-0.953
Shock	6	Std Coef (μ_{sales})	-0.143
Retail Sales	6	$SD/100\%\Delta_{sales}(\mu_{shock})$	-0.704
Shock * Sales	6	$\%\Delta\beta_1/100\%\Delta_{sales}(\mu_{sales})$	-1.545
Shock	7	Std Coef (μ_{sales})	-0.143
Retail Sales	7	$SD/100\%\Delta_{sales}(\mu_{shock})$	-0.704
Shock * Sales	7	$\%\Delta\beta_1/100\%\Delta_{sales}(\mu_{sales})$	-1.547

Fig. 4

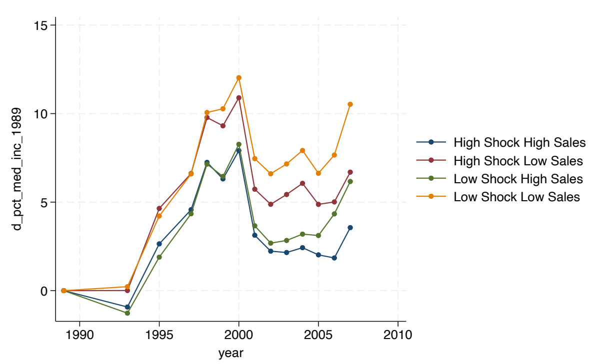

From Table 3 we can see that at the mean of pre-shock retail sales per capita, a 100% increase in sales reduces the effect of import exposure by 154.7%. Equivalently, at the mean of preshock retail sales, a 0.429 unit increase in log preshock sales ($\approx 53.6\%$ increase in the level of sales) would eliminate the negative effect of import exposure on income growth. This is again consistent with my hypothesis.

To see the differential effect of the import exposure shock, I consider counties above and below the median instrumented import shock value and counties above and below the 75th and 25th percentile values for instrumented pre-shock retail sales. In Figure 4, I plot income growth relative to 1989 for the 4 resulting groups of counties. Importantly, this graph provides justification of the exogeneity assumption for the measure of import exposure and it's instrument. We can see this by looking at how counties in the same quartile of pre-shock retail sales have nearly identical income growth trends until after 2000 when the shock is measured. In other words, it's only after 2000 that import exposure between 2000 and

Fig. 5

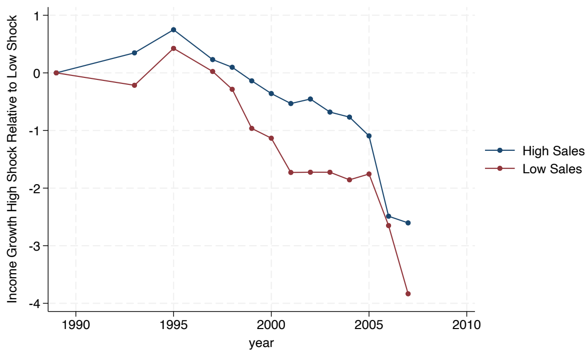

Fig. 6

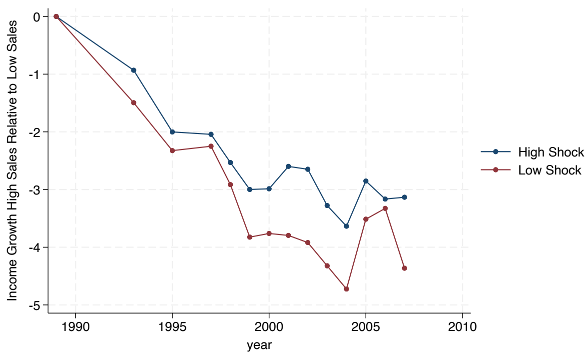

2007 has an effect on county income growth. The same thing can be seen in figure 5 where I'm plotting income growth of counties in the upper quartile of pre-shock retail sales relative to the lower quartile, separated by level of import exposure. Here we can see, though, that there is some evidence of an anticipation effect for counties in the lower quartile of retail sales, although again it's reassuring this only extends at most back to 1998. We might expect that since imports from China have been increasing since 1980, some industries that saw increases in import competition after China's accession to the WTO were already seeing increases in competition before. Since imports in a given industry take time to fully adjust to new trading rules, changes in imports will be auto-correlated over time. The fact that counties in high and low import exposure counties have similar income growth until at least 1998 seems to rule out the possibility that these counties have been exposed to different changes in import competition since trade with China began growing in 1980. Additionally, they should be similar along other dimensions that affect income growth.

The exogeneity of pre-shock retail sales, however, cannot be tested in the same way. The issue is that potential effects of retail sales associated with highway proximity may extend as far back as the construction of the interstates, decades before this study period. We can see in Figure 6 that counties in the upper and lower quartiles of retail sales are on different income growth paths, with counties in the upper quartile

Fig. 7

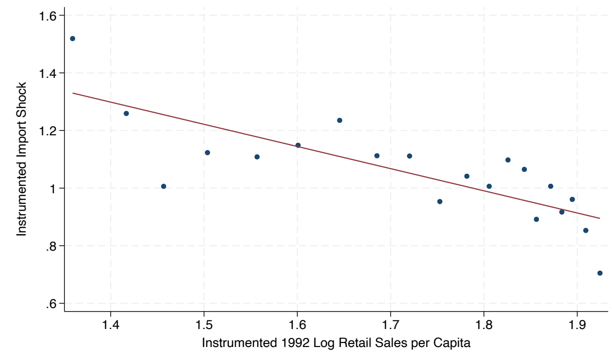

having grown less after 1989 than counties in the lower quartile in every year up to 2007. Additionally, the gap between low retail sales and high retail sales counties grows throughout the whole period, although slightly more for counties that also experience lower trade shocks in the period after 1998. This could be because in high shock counties, having high sales is important for mitigating the effect of the shock reducing the difference between low and high sales counties since low sales counties were growing faster to start with. But, for low shock counties, this sales mechanism isn't as important so there is no "catching up" of high sales relative to low sales counties. This is neither a confirmation nor a rejection of the assumption that counties would have had similar growth with the same level of retail sales, or that the effect of import exposure would have been the same with the same level of retail sales. Figures 5 and 6 are consistent, though, with import exposure having a stronger negative effect in counties with lower pre-shock retail sales.

Additionally, we can see in Figure 7 that counties with lower levels of highway associated retail sales experience higher import shocks on average. So, it's exactly the counties that have stronger negative responses to import competition that are most exposed to increasing trade with China, amplifying the total effect of the shock. Regardless of whether there exists a causal relationship between retail sales associated with proximity to an interstate highway and the effect of import exposure on median wage growth, this could provide additional explanation for why the China shock larger than expected as detailed by Glaeser, (2021). In the next subsection, I explore the effect of import exposure and retail sales on different industries to examine more closely the mechanisms that create the correlations described above.

B. Mechanisms

Given that the effect of import exposure attenuates with retail sales correlated with highway distance, an important follow up question is whether this pattern is similar across different sectors of the economy. Under my primary hypothesis, where non-local retail sales decrease the likelihood of layoffs, reductions in hiring, and wage stagnation caused by declines in manufacturing due to import competition, I would expect some non-manufacturing industries to benefit from outside sources of demand in highly trade exposed counties, but would not expect this to be the case, necessarily, for manufacturing industries. In fact, to the extent that demand for manufacturing production is non-local, the manufacturing sector may act as a placebo for the mitigating effect of pre-shock retail sales. Essentially, if the instrument used for pre-shock retail sales is only capturing exogenous variation in demand for retail goods and services from highway traffic, and demand for manufacturing is entirely non-local, then we should expect null results for the interaction term between import exposure and pre-shock retail sales. To test this, I consider three different industries that I expect will exhibit different correlations with import exposure and pre-shock retail sales: manufacturing, retail, and accommodation and food (NAICS codes 31-33, 44-45, and 72 respectively). For each industry, I look at 4 measures of economic performance: percent changes in number of establishments, total payroll, number of employees, and pay per employee over the 2000-2007 period. Each regression is run using the specification in column 6 of Table 3, with border-cluster fixed effects.

Table 4 reports the results from these regressions. Across all sectors and outcomes, import exposure is negatively correlated with growth, although this relationship is only statistically significant for the change in manufacturing and accommodation and food establishments, and for the change in retail payroll. The coefficients on the interaction term are generally positive with two exceptions and are statistically significant (at the 10% level) for retail payroll and accommodation and food establishments. Although the interaction term is not significant for any of the manufacturing variables, I do not have enough precision in these estimates to reject economically meaningful mitigating effects of pre-shock retail sales on the negative effect of import exposure. Because of this, I can't conclude that the only channel through which pre-shock retail sales mitigate the effect of import exposure is through an effect on non-manufacturing sectors and their supply of non-local demand. One possible explanation is that the effect of the China shock extends further back than the 2000s when China formally joined the WTO. If this was the case and import exposure is positively auto-correlated between periods, then we would be picking up the effect of import exposure and the heterogeneous effect of pre-shock retail sales in previous periods. If for example, manufacturing firms' location and employment decisions are a function of the employment and location decisions of non-manufacturing firms, then a heterogeneous impact on only non-manufacturing firms in the pre-period could create heterogeneity in the treatment effect for manufacturing firms in future periods.

For completeness, I repeat these regressions for all other NAICS 2-digit industry classifications. The results can be found in Appendix section 1.3 Many industries exhibit a negative effect of import exposure that is attenuated by pre-shock retail sales, including NAICS 11 (Agriculture, Forestry, Fishing, and Hunting), NAICS 52 (Finance and Insurance), NAICS 53 (Real Estate, Rental, and Leasing), NAICS 54 (Professional, Scientific, and Technical Services), NAICS 55 (Mangement of Companies and Enterprises), NAICS 71 (Arts, Entertainment, and Recreation), and NAICS 81 (Other Services). In fact, all industries whose economic growth outcomes are negatively correlated with import exposure have effects that shrink with pre-shock retail trade, indicating that the instrument for pre-shock retail trade may be correlated with factors important to the economic performance of a variety of industries affected directly and indirectly by import competition.

6. Discussion

Although my results are consistent with the hypothesis that retail trade from non-local sources played an important role in reducing the negative income effects of the China shock between 2000 and 2007, they do not rule out and in fact may point to other possible causal mechanisms. It is clear that highway distance is not exogenous and is correlated with population, area, density, and distance to an urban area. It could be that any of these factors also help mitigate the effect of the China shock. For instance, more populous counties are likely also more diversified counties. If the set of less affected or unaffected industries in a county is similar to the set of trade exposed industries, which is more probable if a county is diversified, the transition process for workers of moving from one industry to another after a trade shock may be more smooth, not requiring excessive retraining or relocation. This could also mitigate the income losses due to manufacturing decline in trade exposed counties if it decreases the time spent unemployed or out of the labor force for workers who can't find jobs in trade exposed industries. I don't think this would explain away all of my results, as controlling for either urban distance and population density or urban distance, population, and area, only changes estimated coefficients slightly at least comparing within border clusters. The point is, though, I cannot rule out other explanations. That being said, this paper still has something to say about the aftermath of the China shock. Whether due to retail sales or other factors that vary with a county's distance to an interstate highway, there exists significant variation between U.S. counties in the effect of the China shock on income growth and other economic outcomes. Combined with the fact that pre-shock retail sales are negatively correlated with the magnitude of the change in import exposure between 2000 and 2007, it means that counties which were most affected per dollar change in import exposure also experienced the largest change in import exposure. This would have added to the geographic concentration of the effect of the China shock. These findings suggest that in thinking about future economic shocks, it may be similarly important to consider the characteristics of the most exposed populations to fully grasp how such shocks shape outcomes across regions.

7. Conclusion

This paper has examined how regional integration shapes the local consequences of the China shock. Building on literature that documents declines in income growth due to

TABLE V: Sector Decomposition

	Manufacturing				Retail				Accommodation and Food			
VARIABLES	Establishments	Payroll	Employees	Pay	Establishments	Payroll	Employees	Pay	Establishments	Payroll	Employees	Pay
Shock	-0.200*	-0.269	-0.498	-0.0790	-0.0864	-0.466**	-0.0413	-0.0630	-0.315**	-0.576	-0.321	0.135
	(0.105)	(0.697)	(0.417)	(0.148)	(0.0710)	(0.190)	(0.125)	(0.0720)	(0.138)	(0.563)	(0.376)	(0.239)
Log Preshock Retail Sales	-0.181**	-0.406	0.0584	-0.0407	-0.0580	-0.0776	0.280**	0.0197	-0.219**	-0.340	-0.0752	-0.0282
	(0.0847)	(0.523)	(0.327)	(0.135)	(0.0571)	(0.175)	(0.115)	(0.0669)	(0.111)	(0.384)	(0.256)	(0.162)
Shock x Sales	0.0688	0.119	0.320	0.00746	0.0240	0.202*	-0.0293	0.0622	0.133*	0.275	0.118	-0.0758
	(0.0601)	(0.334)	(0.200)	(0.0696)	(0.0406)	(0.114)	(0.0751)	(0.0429)	(0.0789)	(0.305)	(0.203)	(0.127)
Constant	0.554***	5.315***	0.293	0.225	0.332***	0.295	-0.739***	0.420***	0.720***	1.709**	0.493	0.665**
	(0.154)	(1.052)	(0.778)	(0.270)	(0.104)	(0.313)	(0.206)	(0.120)	(0.202)	(0.718)	(0.479)	(0.301)
Observations	1,261	593	592	490	1,261	1,225	1,225	1,179	1,261	1,192	1,192	1,167
R-squared	0.534	0.668	0.622	0.704	0.556	0.592	0.600	0.525	0.515	0.578	0.575	0.484

Standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

increases in import exposure from China, I consider one specific source of treatment effect heterogeneity: the extent to which local economies are linked to outside sources of demand through non-local retail activity. Using pre-shock retail sales per capita, instrumented by distance to interstate highways to capture a plausibly non-local portion of retail sales, I show that counties with lower highway correlated retail sales experienced significantly larger declines in median income growth in response to rising Chinese import exposure between 2000 and 2007.

This is consistent with my hypothesis that reductions in spending due to declines in employment in manufacturing are an important channel through which shocks to manufacturing spread to other industries within a local economy and this spread will be amplified in counties that have fewer sources of non local demand for non-manufacturing production.

However, this pattern of an attenuation in the effect the China shock with the level of pre-shock retail sales is common across many different industries including manufacturing. It points to the fact that there may be other causes of this treatment effect heterogeneity. I therefore emphasize the importance of the existence of treatment effect heterogeneity rather than a causal interpretation of why this treatment effect heterogeneity exists. Combined with the fact that the intensity of the China shock declines on average with pre-shock retail sales, though, I can conclude that the counties most susceptible to declines in median income growth due to changes in import exposure are also the counties that experienced the largest changes in import exposure. This provides further explanation for the question put forward in Glaeser (2021), of why the China shock was so geographically concentrated, leading to important changes in the economic geography of the United States over the last three decades.

From a policy and future research perspective, these findings suggest the importance of understanding the characteristics of the people and places exposed to future economic shocks, in order to consider and prepare for what a response to or recovery from these shocks will look like. For example, changes in the economy due to A.I. will likely look very different from the China shock, but the characteristics of the most exposed populations could play an important role in determining how society adapts to the resulting changes in employment.

References

[Adão, Arkolakis, and Esposito(2019)] Adão, Rodrigo, Costas Arkolakis, and Federico Esposito. 2019. "General Equilibrium Effects in Space: Theory and Measurement." NBER Working Paper No. 25544, National Bureau of Economic Research, Cambridge, MA. February; revised June 2020. http://www.nber.org/papers/w25544.

[Anderson and Matsa(2011)] Anderson, Michael L., and David A. Matsa. 2011. "Are Restaurants Really Supersizing America?" *American Economic Journal: Applied Economics* 3(1): 152–188. https://doi.org/10.1257/app.3.1.152.

[Autor, Dorn, and Hanson(2013)] Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103(6): 2121–2168. https://doi.org/10.1257/aer.103.6.2121.

[Autor et al.(2025)] Autor, David, David Dorn, Gordon H. Hanson, Maggie R. Jones, and Bradley Setzler. 2025. "Places versus People: The Ins and Outs of Labor Market Adjustment to Globalization." NBER Working Paper No. 33424, National Bureau of Economic Research, Cambridge, MA. January; revised June 2025. http://www.nber.org/papers/w33424.

[Eriksson et al.(2019)] Eriksson, Katherine, Katheryn Russ, Jay C. Shambaugh, and Minfei Xu. 2019. "Trade Shocks and the Shifting Landscape of U.S. Manufacturing." NBER Working Paper No. 25646, National Bureau of Economic Research, Cambridge, MA. March; revised January 2020. http://www.nber.org/papers/w25646.

[Foote, Grosz, and Stevens(2015)] Foote, Andrew, Michel Grosz, and Ann Huff Stevens. 2015. "Locate Your Nearest Exit: Mass Layoffs and Local Labor Market Response." NBER Working Paper No. 21618, National Bureau of Economic Research, Cambridge, MA. October. http://www.nber.org/papers/w21618.

[Glaeser and Gyourko(2001)] Glaeser, Edward L., and Joseph Gyourko. 2001. "Urban Decline and Durable Housing." NBER Working Paper No. 8598, National Bureau of Economic Research, Cambridge, MA. http://www.nber.org/papers/w8598.

[Glaeser(2021)] Glaeser, Edward L. 2021. "Comment." *Brookings Papers on Economic Activity*, Fall: 457–460. https://www.brookings.edu/wp-content/uploads/2021/09/15985-BPEA-BPEA-FA21_WEB_Autor-et-al_Glaeser-Comment.pdf

[Horn et al.(2024)] Horn, Keren, Henry Pollakowski, Jeffrey Zabel, and Joseph Chance. 2024. "The Impact of Job Opportunities, Housing Costs, and Mobility Decisions on the Re-employment Experiences of Displaced Workers." MIT Center for Real Estate Research Paper No. 24/05, May. https://dx.doi.org/10.2139/ssrn.4852572

[Kleiner and Krueger(2008)] Kleiner, Morris M., and Alan B. Krueger. 2008. "The Prevalence and Effects of Occupational Licensing." NBER Working Paper No. 14308, National Bureau of Economic Research, Cambridge, MA. September. http://www.nber.org/papers/w14308.

[Kovak(2013)] Kovak, Brian K. 2013. "Regional Effects of Trade Reform: What Is the Correct Measure of Liberalization?" *American Economic Review* 103(5): 1960–1976. August. https://www.jstor.org/stable/42920635.

[Krause(2024)] Krause, Eleanor. 2024. "Job Loss, Selective Migration, and the Accumulation of Disadvantage: Evidence from Appalachia's Coal

Country.” Unpublished manuscript, Harvard University, Cambridge, MA. March 6. https://ssrn.com/abstract=4794041

[Topalova(2010)] Topalova, Petia. 2010. “Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India.” *American Economic Journal: Applied Economics* 2(4): 1–41. October. https://www.aeaweb.org/articles?id=10.1257/app.2.4.1.

[von Luxburg(2007)] von Luxburg, Ulrike. 2007. “A Tutorial on Spectral Clustering.” *Statistics and Computing* 17(4): 395–416. https://doi.org/10.1007/s11222-007-9033-z.

Appendix

A. Alternative Highway Distance Measure

Fig. A1

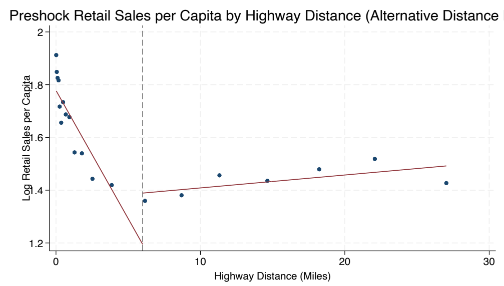

TABLE A1: Regression Results

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				%ΔMed Inc/Year			
Shock	-0.157***	0.183	0.114	-0.322	-0.324	-0.504	-0.498
	(0.0250)	(0.458)	(0.457)	(0.384)	(0.386)	(0.448)	(0.450)
Log Preshock Retail Sales		-0.148	-0.258	-0.648**	-0.689**	-0.814***	-0.859***
		(0.325)	(0.327)	(0.286)	(0.289)	(0.288)	(0.295)
Shock x Sales		-0.223	-0.189	0.128	0.136	0.241	0.252
		(0.282)	(0.281)	(0.235)	(0.237)	(0.283)	(0.285)
Constant	0.376***	0.653	0.886	0.999*	1.072**	1.848***	1.915***
	(0.0305)	(0.541)	(0.544)	(0.510)	(0.514)	(0.483)	(0.493)
Urban Distance Control	No	No	Yes	No	Yes	No	Yes
Spectral Cluster FE	No	No	No	Yes	Yes	No	No
Border Cluster FE	No	No	No	No	No	Yes	Yes
Observations	1,558	1,558	1,558	1,558	1,558	1,558	1,558
R-squared	0.024	0.078	0.078	0.484	0.479	0.636	0.632

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

B. Alternate Max Highway Distances

Fig. A2

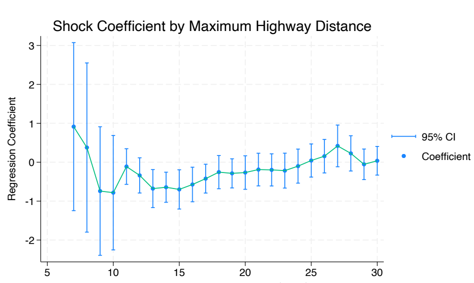

Fig. A3

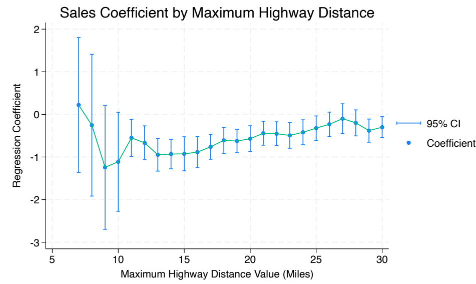

Fig. A4

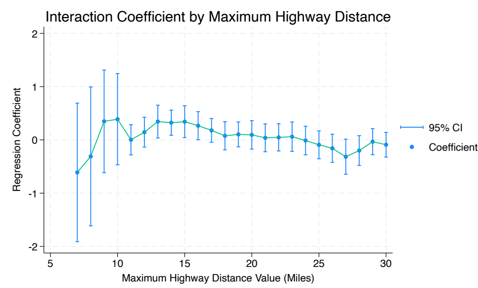

C. Alternate Controls

TABLE A2: Regression Results

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				%ΔMed Inc/Year			
Shock	-0.148***	-0.404	-0.328	-0.385*	-0.187	-0.569***	-0.515**
	(0.0313)	(0.291)	(0.292)	(0.233)	(0.237)	(0.220)	(0.229)
Log Preshock Retail Sales		-0.588***	-0.440	-0.636***	-0.0834	-0.872***	-0.650*
		(0.216)	(0.299)	(0.181)	(0.297)	(0.181)	(0.348)
Shock x Sales		0.138	0.0918	0.170	0.0178	0.268**	0.226
		(0.174)	(0.176)	(0.139)	(0.147)	(0.131)	(0.142)
Constant	0.350***	1.371***	1.865***	1.055***	2.669***	1.953***	2.689***
	(0.0371)	(0.369)	(0.386)	(0.332)	(0.441)	(0.335)	(0.543)
Urban Distance Control	No	No	Yes	No	Yes	No	Yes
Log Population Control	No	No	Yes	No	Yes	No	Yes
County Area Control	No	No	Yes	No	Yes	No	Yes
Spectral Cluster FE	No	No	No	Yes	Yes	No	No
Border Cluster FE	No	No	No	No	No	Yes	Yes
Shock F-Stat	1360.03	455.15	426.24	254.20	236.07	179.45	156.92
Sales F-Stat		61.77	20.58	43.86	14.68	32.64	11.09
Shock x Sales F-Stat		297.63	298.31	145.33	157.48	98.60	104.66
Observations	1,308	1,308	1,308	1,308	1,308	1,308	1,308
R-squared	0.028	0.079	0.109	0.496	0.551	0.625	0.658

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Reported F-Statistics are from the joint significance test of excluded instruments in the first stage regressions.

D. Other Industries

TABLE A3: Agriculture, Forestry, Fishing and Hunting

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	-0.721**	-0.870	-0.163	-0.442
	(0.296)	(0.903)	(0.355)	(0.421)
Log Preshock Retail Sales	-0.399*	1.532	-0.147	-1.602*
	(0.235)	(1.109)	(0.436)	(0.837)
Shock x Sales	0.396**	0.162	0.153	0.261
	(0.165)	(0.509)	(0.200)	(0.251)
Constant	0.615	-2.831	-0.587	2.291*
	(0.426)	(2.030)	(0.799)	(1.205)
Observations	1,058	250	249	130
R-squared	0.458	0.593	0.774	0.701

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A4: Mining

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	0.501	-3.031	3.378	3.161
	(0.326)	(4.609)	(2.452)	(2.462)
Log Preshock Retail Sales	0.520**	-1.342	2.906**	1.469
	(0.263)	(2.637)	(1.403)	(1.200)
Shock x Sales	-0.328*	1.375	-1.115	-1.028
	(0.193)	(1.763)	(0.938)	(0.705)
Constant	-0.509	2.806	-7.408**	-2.990
	(0.459)	(6.442)	(3.427)	(3.197)
Observations	956	170	171	120
R-squared	0.428	0.727	0.744	0.664

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A5: Utilities

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees
Shock	-0.0151	-308.8	28.09
	(0.197)	(1,485)	(173.9)
Log Preshock Retail Sales	0.185	-228.5	22.61
	(0.160)	(1,112)	(130.3)
Shock x Sales	0.00637	145.8	-12.45
	(0.112)	(693.7)	(81.26)
Constant	-0.332	465.6	-48.00
	(0.292)	(2,276)	(266.6)
Observations	1,180	113	113
R-squared	0.553		

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A6: Transportation and Warehousing

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	0.160	0.478	-0.368	0.325
	(0.143)	(0.724)	(0.368)	(0.229)
Log Preshock Retail Sales	0.110	0.179	0.130	0.143
	(0.116)	(0.545)	(0.278)	(0.185)
Shock x Sales	-0.116	-0.143	0.200	-0.0872
	(0.0859)	(0.376)	(0.193)	(0.112)
Constant	-0.205	-0.520	-0.722	0.0871
	(0.210)	(1.082)	(0.552)	(0.390)
Observations	1,321	985	984	808
R-squared	0.533	0.520	0.565	0.603

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A7: Information

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	0.260*	-6.503	1.267	-0.281
	(0.149)	(16.16)	(4.725)	(3.591)
Log Preshock Retail Sales	0.287**	-6.297	-0.0340	-0.288
	(0.133)	(8.024)	(2.178)	(1.855)
Shock x Sales	-0.206**	3.588	-0.465	-0.196
	(0.0907)	(8.139)	(2.364)	(1.656)
Constant	-0.316	12.78	-0.0291	0.926
	(0.263)	(16.49)	(4.479)	(4.021)
Observations	1,205	316	317	242
R-squared	0.476	0.366	0.703	0.715

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A8: Finance and Insurance

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	-0.472***	-1.011**	-0.263	-0.0925
	(0.127)	(0.395)	(0.258)	(0.149)
Log Preshock Retail Sales	-0.0721	-0.783***	-0.147	0.0277
	(0.104)	(0.302)	(0.197)	(0.119)
Shock x Sales	0.207***	0.437*	0.00570	0.0497
	(0.0750)	(0.224)	(0.146)	(0.0815)
Constant	0.603***	2.264***	0.338	0.505**
	(0.192)	(0.585)	(0.382)	(0.236)
Observations	1,335	1,132	1,132	976
R-squared	0.456	0.550	0.525	0.590

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A9: Real Estate Rental and Leasing

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	0.108	-5.554***	-4.727***	3.143**
	(0.158)	(2.046)	(1.118)	(1.577)
Log Preshock Retail Sales	0.312**	-3.848***	-2.200***	1.693**
	(0.129)	(1.154)	(0.627)	(0.760)
Shock x Sales	-0.0875	2.951***	2.456***	-1.693**
	(0.0931)	(1.101)	(0.602)	(0.842)
Constant	-0.137	7.469***	3.918***	-2.237*
	(0.233)	(2.022)	(1.100)	(1.348)
Observations	1,301	957	958	780
R-squared	0.525	0.444	0.310	0.300

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A10: Professional, Scientific, and Technical Services

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	-0.290**	-0.632	-0.979**	-0.503***
	(0.116)	(0.702)	(0.433)	(0.191)
Log Preshock Retail Sales	0.0560	0.376	0.130	-0.194
	(0.0963)	(0.463)	(0.286)	(0.133)
Shock x Sales	0.140**	0.289	0.469**	0.274***
	(0.0684)	(0.373)	(0.230)	(0.106)
Constant	0.312*	0.0164	-0.521	1.050***
	(0.177)	(0.881)	(0.543)	(0.261)
Observations	1,326	1,082	1,082	912
R-squared	0.471	0.505	0.454	0.589

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A11: Management of Companies and Enterprises

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	-2.117***	24.53	21.14	-16.77
	(0.645)	(36.90)	(33.88)	(37.22)
Log Preshock Retail Sales	-1.551***	10.69	10.53	-12.68
	(0.449)	(19.19)	(17.62)	(22.63)
Shock x Sales	1.142***	-11.08	-9.417	7.504
	(0.339)	(15.85)	(14.55)	(16.53)
Constant	2.900***	-22.43	-22.52	26.77
	(0.865)	(41.63)	(38.23)	(47.62)
Observations	1,013	339	339	256
R-squared	0.439	0.197		

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A12: Administrative Support and Waste Management Services

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	-0.143	-1.348	0.0410	-0.0124
	(0.152)	(0.916)	(0.513)	(0.309)
Log Preshock Retail Sales	-0.154	-1.765**	-0.315	0.0611
	(0.123)	(0.703)	(0.389)	(0.232)
Shock x Sales	0.0159	0.439	-0.122	0.0250
	(0.0872)	(0.485)	(0.273)	(0.159)
Constant	0.741***	3.970***	0.505	0.618
	(0.226)	(1.362)	(0.754)	(0.449)
Observations	1,295	893	896	760
R-squared	0.510	0.468	0.548	0.559

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A13: Educational Services

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	0.0214	-24.17	-16.81	4.650
	(0.192)	(21.72)	(14.40)	(5.990)
Log Preshock Retail Sales	0.242	-14.02	-9.170	2.135
	(0.190)	(11.96)	(7.932)	(3.296)
Shock x Sales	-0.0717	11.79	8.340	-2.368
	(0.115)	(11.07)	(7.338)	(3.013)
Constant	0.0894	29.21	18.67	-3.531
	(0.334)	(24.59)	(16.30)	(6.820)
Observations	1,087	517	517	475
R-squared	0.526			0.156

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A14: Healthcare and Social Assistance

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	-0.136	-0.359	-0.251	0.156
	(0.0956)	(0.409)	(0.281)	(0.504)
Log Preshock Retail Sales	0.0785	0.0683	0.216	-0.0723
	(0.0781)	(0.356)	(0.245)	(0.299)
Shock x Sales	0.0183	0.180	0.0667	-0.0837
	(0.0565)	(0.234)	(0.161)	(0.268)
Constant	0.424***	0.843	-0.184	0.577
	(0.145)	(0.667)	(0.458)	(0.549)
Observations	1,335	1,202	1,202	1,119
R-squared	0.514	0.558	0.539	0.520

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A15: Arts, Entertainment, and Recreation

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	0.0342	-5.091**	-2.395*	-1.501**
	(0.196)	(2.010)	(1.312)	(0.740)
Log Preshock Retail Sales	0.0358	-2.015*	-0.917	-0.817**
	(0.143)	(1.107)	(0.724)	(0.412)
Shock x Sales	-0.0180	2.776**	1.249*	0.790*
	(0.111)	(1.113)	(0.726)	(0.404)
Constant	-0.0737	3.264	1.042	2.604***
	(0.256)	(2.040)	(1.334)	(0.761)
Observations	1,281	760	758	631
R-squared	0.527	0.640	0.639	0.598

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

TABLE A16: Other Services

VARIABLES	(1) Establishment	(2) Payroll	(3) Employees	(4) Pay
Shock	-0.166**	-0.650	-0.216	-0.455**
	(0.0744)	(0.469)	(0.259)	(0.211)
Log Preshock Retail Sales	-0.125**	-0.818**	0.0415	-0.315**
	(0.0609)	(0.325)	(0.179)	(0.141)
Shock x Sales	0.0703	0.404	0.0558	0.241**
	(0.0442)	(0.264)	(0.146)	(0.114)
Constant	0.271**	2.025***	-0.280	1.507***
	(0.113)	(0.623)	(0.344)	(0.267)
Observations	1,341	1,259	1,259	1,182
R-squared	0.510	0.382	0.494	0.450

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1