Boston University Undergraduate Economics Journal

Fall 2025 Issue IV

Meet the Team

Ellie Wang
Editor-in-Chief

Ajay Raman
Copy Editor

Nandini Jha
Head Senior Editor

Matthew Morris
Head Layout Editor

Lina Santos
Senior Editor

Dahsom Lim
Senior Editor

Ava Carrel
Senior Editor

Cole Pinto
Junior Editor

Mohammad Aiman Ashad
Junior Editor

Sara Asdhir
Layout Editor

Madi Kim
Layout Editor

Abigail Lam
Layout Editor

BUUEJ Editing Team Fall 2025

Boston University
Undergraduate Economics Journal

Table of Contents

Articles

The Impact of Children on Women's Labor Force Participation: Analyzing the Effect of Family Size on Hours Worked

By Daniel Frank, Koba Jgenti, Lauren Kong

This paper examines how family size influences women's labor force participation by measuring the effect of the number of children (nchild) on weekly hours worked (uhrsworkt). Using IPUMS CPS ASEC microdata from 2010–2019, this paper estimates using both simple and multivariate linear and log-linear models, controlling for age, educational attainment, marital status, race, region, and year. Building on the models, this paper further introduces interaction terms between nchild and race, education, and a post-2015 indicator to uncover heterogeneity. In baseline OLS, results show that each additional child is associated with an average increase of approximately 0.175 hours of work per week; a similar but greater effect persists after controls (∼0.25 hours per child, $p < 0.01$). Interaction analyses reveal that Black mothers experience a larger per-child increase in hours worked (≈ 0.53 hours), while each additional year of schooling attenuates the child effect on hours worked by ∼0.30 hours. Moreover, the child-hour association strengthened by ∼0.0763 hours in the 2015–2019 period. These findings emphasize subtle variations in the effect of children on hours worked and inform policy design to support working mothers across different demographic groups.

I. Introduction

Women's participation in the labor force is a major contributor to gender equality, innovation, and development. There are many factors that influence an individual's decision to commit to the labor force–one such factor is raising children. In 2023, 32.6 million families included children under the age of 18 (US Department of Labor, 2024). This is a significant number of families that include at least one child, which is why it is crucial to understand the support (or possible lack thereof) that mothers get from certain policies, and where new programs could be implemented to further increase this support. From 1990 to 2023, the US saw a trend of women delaying having children, meaning that on average, they were having children later. During this time, there was an increase of 17% in childbearing for women between ages 30 to 34. Meanwhile, during this same period, there was a 71% increase from ages 35 to 39, showing that having children relatively later in life may be preferred as of late (Center for Disease Control and Prevention, 2025). This new trend could be attributed to the difficulties of balancing child needs along with career aspirations, highlighting the importance of accommodating workplace policies.

Historically, there have been several policies aimed at helping workers with children. One such policy was the Family and Medical Leave Act, which was passed on February 5th of 1993. This act allowed employees to take job protected leave for specified family reasons, such as the birth of a child (US Department of Labor, 2024). While programs like this have made small improvements in working conditions, there is still a penalty in the labor force participation and wages for mothers (Bowers et al., 2015). This phenomenon, known as the "motherhood penalty," indicates that mothers may take flexible jobs to better manage childcare responsibilities (ex. taking care of children on sick days).

To shed light on new or already implemented policies that could support women in balancing family and work responsibilities, this study analyzes just how impactful the effects of having children are on women's actual participation in the labor force, measured by the usual number of hours per week spent at all jobs. This paper runs simple and multiple linear regression models and analyzes cross-sectional data from IPUMS CPS from 2010 to 2019. This study proposes the following question: How does the specific number of children that a woman has in the household affect their labor force participation, measured by time spent at all jobs throughout the week? This paper finds evidence that having more children is associated with a slight positive increase in the number of hours worked per week for women. When controlling for factors such as age, education, race, marital status, region, and time period, this effect gets slightly stronger. There are also significant differences between racial groups and across educational groups.

II. Literature Review

A study was conducted using the Norwegian administrative registers, where research was done on the effects on labor supply due to having children. Instrumental variables were used for estimation. Findings showed that having additional children is initially correlated with a significant decrease in labor supply for women. However, as children grow up, this reduction fades, and the relationship turns positive for women without college degrees (Cools, Markussen, & Strøm, 2017). In addition, colleagues found that in the long run, the effects on career were found to be more persistent, leading to women possibly being less likely to get promoted. Just like this study, which controls for education, this study also takes education into account by controlling for years of education completed and splitting its sample into college graduates and non-college graduates. This study is used in this paper to highlight the long-term effects on women labor market outcomes when controlling for education. Moreover, this study's findings on education and its impact on the main relationship of this study align with the results outlined in Section V. However, the results on how and why education affects this relationship are not as in-depth regarding factors such as long-term promotions or wage gaps between men and women.

A different study conducted by J.D. Angrist and William Evans (1998) examined the relationship between fertility rates and work effort. This study utilized an instrumental variable strategy that depended on sibling sex-mix in families with two or more children. The intuition for this study was based on the finding that if a couple's first two children are of the same gender, then they are more likely to have a third child (thus providing an exogenous source of variation for family size). The study tracked children in households from mothers aged 21 to 35, where the oldest child was less than 18 years old at the time of the census. It was concluded that children in the household led to a reduction in female labor supply, and that women with kids were less likely to work in the labor force and worked fewer hours (Angrist & Evans, 1998). Like the Norwegian study described previously, this study also accounted for education by controlling for years of education completed. We use this study to further examine the connection between fertility rates, as well as race, in relation to female labor force participation.

While the results differ from those of Angrist and Evans, they do partially align with those of Cools, Markussen, and Strøm, who also observed a positive effect of children on labor supply in more educated women. This paper's findings similarly suggest that education mediates the relationship between hours worked and fertility; that is, on average, less educated women increase their work hours more than educated women do as their families grow. Additionally, as the sample includes women aged 15-79, it is possible that many of them have children who are relatively older, thus contributing to the non-inverse relationship that was found between hours worked and children. This also partially aligns with Cools' and colleagues' findings, as they observed a fade in the inverse relationship between women's labor supply and children as children grew older.

However, as the results of the main relationship–that is, simply, the effect of children on women's hours worked–do not align with those of the studies above, this paper's findings contribute new evidence to this literature. The positive relationship between the number of children on hours worked–in contrast to the negative relationships found in prior studies–could be due to several outstanding factors. One possible explanation could be related to time period differences and rising prices in childcare. According to a resource that analyzed data from the Bureau of Labor Statistics, prices for day care and preschool as of 2019 were roughly 214% higher than they were in 1990 (Joughin, 2021). Considering that some of the prior studies this paper referenced for the analysis, namely Angrist's and Evans' study, were written before the 21st century, it is possible that these changes over time have contributed to mothers working more hours to meet the financial requirements of childcare. Additionally, social changes–particularly those related to changes in the rate of stay-at-home fathers–have also increased substantially since the turn of the century. A 2023 analysis by Pew Research examined U.S. Census Bureau data from the last 30 years and found that fathers, as of 2023, represent 18% of stay-at-home parents compared to just 11% in 1989. This same analysis found that in 2021, 23% of stay-at-home fathers reported staying home to care for their home and family, compared to only 4% who reported the same in 1989. These shifts suggest changes in family dynamics over time, and with fathers taking on more family responsibilities, mothers gain more flexibility to work more or remain in the workforce.

The third and potentially most significant reason behind the novelty of this paper's findings may be related to differences in how dependent variable, hours worked, was defined. To the best of our knowledge, prior studies, including the ones discussed above, use labor supply–one's willingness and ability to work–as their measurement of labor force participation. Although it is possible that women may report being less willing to work after having children, this does not necessarily mean that actual hours worked in already working women decreases. Since our analysis focuses on hours worked rather than overall labor force participation, it could also be the case that mothers who are already working may increase their overall hours worked due to the financial responsibilities that come with raising children; this is especially possible given the significant increases in childcare prices in recent decades.

III. Econometric Model

Simple Linear Regression Models (Table 4A):

Model 1:

$$uhrsworkt = \beta_1 + \beta_2 nchild + e$$

Model 2:

$$luhrsworkt = \beta_1 + \beta_2 nchild + e$$

To better understand the basic relationship between number of children and hours worked, this study first establishes that a relationship exists through simple linear regression models. In Model 1, 1 represents *uhrsworkt* when nchild = 0, β_2 represents slope, or the change in *uhrsworkt* when *nchild* increases by 1, and e represents error due to omitted factors, which was accounted for in the multiple regression models. If the coefficient for *nchild*, 2, is positive, then it means that as the number of children residing within a woman's household increases, predicted hours worked increases. If 2 takes on a negative value, predicted hours worked in sample decreases as the number of children increases.

Model 2 depicts a logistic-linear functional form of Model 1. This model was used to correct for skewness in hours worked and to reduce variance (which was also done through dropping major outliers). Additionally, interpreting the coefficients of this model allows us to easily examine semi-elasticity (β 1). One limitation of this model is the fact that taking the log of hours worked means that those in the sample with zero hours worked are excluded. However, this is not a substantial problem, since the number of women who work zero hours in the sample is very minute[1]. This model is not necessarily better at describing the relationship, but nonetheless, testing this second functional form allows for a more in-depth analysis.

Multiple Linear Regression Models (Table 4B):

Model 3:

$$\begin{aligned} uhrsworkt = {} & \beta_1 + \beta_2 nchild + \beta_3 age + \beta_4 educclean + \beta_5 black + \beta_6 asian \\ & + \beta_8 midwest + \beta_9 south + \beta_{10} west + \beta_{11} single + \beta_{12} yg2015 + e \end{aligned}$$

Model 4:

$$\begin{aligned} luhrsworkt = {} & \beta_1 + \beta_2 nchild + \beta_3 age + \beta_4 educclean + \beta_5 black + \beta_6 asian + \beta_7 amindian \\ & + \beta_8 midwest + \beta_9 south + \beta_{10} west + \beta_{11} single + \beta_{12} yg2015 + e \end{aligned}$$

For multiple regression analysis, depicted in *Model 3*, examines the effect of number of children on hours worked by women in the last week, controlling for region, educational attainment, marital status, race (White, Black, American Indian, Asian), and year/timeframe (to account for variation in fertility rates across 2010-2019). The total number of independent variables, including all dummy variables, is 12. The inclusion of all of these independent variables helps control for endogeneity, as all of these variables are relevant to the study in that they can have an effect on hours worked[2]. Where one lives, for example, might have a notable impact on labor force participation, as it is possible that an individual residing in a rural community might have more job opportunities than one in a more urban location. The study "Employment Rates Higher Among Rural Mothers than Urban Mothers," conducted by Kristin Smith, found that rural women with young children are

[1] Using the tab command in Stata, we see that 411 women out of the 388,852 in our sample reported working zero hours. This is about 0.11%, which is very minute in comparison to the full sample size.

[2] Supported by levels of significance depicted in Stata regression outputs.

more likely to be employed compared to urban women with young children. This discrepancy was found to be mainly due to jobs in rural communities being more flexible. In addition, this study found that women in urban areas also have a lower employment rate after having children, since there is a higher opportunity cost; their jobs often have longer hours and require higher commitment, which may conflict with the needs of their children and consequently discourage them from returning to work.

Additionally, race, education, and marital status can all have major effects on hours worked. As expected, average hours worked to vary across the four races that was account for, and likewise, expect variations in educational attainment and marital status to have substantial effects on hours worked as well. For example, in the study mentioned in literature review conducted by Cools, Markussen and Strøm, it was found that college educated women experience larger career loss and a significantly larger decrease in work hours from having kids; meanwhile, non-college educated women did have reduced labor supply due to having children, but long-term careers were not as disturbed. Finally, it is important to account for fertility rates across 2010-2019, as it is possible that a fall in these rates could influence labor/social policies, and could also be the result of economic/government policy or a change in social norms. According to the Center for Disease Control and Prevention (CDC), the total fertility rate in the US declined to 1706 births per 1000 women in 2019 which is a record low for the US. In 2019 the general fertility rate also decreased to 58.3 births per 1000 women ages 15-44, and a decline was observed in almost every age group under 35 years of age. Studies have attributed these declines to the Great Recession, in which economic uncertainty, job stability, and financial security were big factors that caused many women to delay having children (Abbamonte, 2025). Furthermore, rising costs of education and childcare, women may be deterred from having children.

β_1, the intercept of *Model 3*, represents reported hours spent at all jobs in the last week by women, specifically when age is zero, the woman has no children, no educational attainment, is white, from the northeast, is not single (either their spouse is present or they have a spouse who is not present), and the year is between 2010 and 2015. Additionally, despite an age value of zero not having a meaningful real-world interpretation, it nonetheless impacts intercept. Because one's age cannot be zero, the intercept of the model may not have practical significance. However, age must be accounted for to control for endogeneity, since it could also have a major impact on hours worked. Someone who is older typically reduces their work hours as they get towards retirement due to health decline. In addition, intercept not having practical significance has little to no effect on this model's ability to estimate the relationship between number of children and hours worked, which is the primary focus of this study.

ε_i, the error term of the model, represents the unexplained variation, as well as unobserved/omitted variables (not included in the model), that could affect hours worked. Based on the OLS assumptions, the expected value of this error term is zero, meaning that the model does not overestimate or underestimate hours worked. Additionally, the expected value of ε_i being zero assumes that variance is constant across all levels of x variables (ie., homoscedasticity).

For *Model 4*, the same variables were used as in Model 3 while incorporating the logistic-linear functional form.

Model 5 (nchild_black interaction):

$$\begin{aligned}\text{uhrsworkt} = \beta_1 &+ \beta_2\text{nchild} + \beta_3\text{age} + \beta_4\text{educclean} + \beta_5\text{black} + \beta_6\text{asian}\\ &+ \beta_7\text{amindian} + \beta_8\text{midwest} + \beta_9\text{south} + \beta_{10}\text{west} + \beta_{11}\text{single}\\ &+ \beta_{12}\text{yg2015} + \beta_{13}(\text{nchild}\cdot\text{black}) + \varepsilon\end{aligned}$$

Model 6 (nchild_educ interaction):

$$\begin{aligned} uhrsworkt = {} & \beta_1 + \beta_2 \text{nchild} + \beta_3 \text{age} + \beta_4 \text{educclean} + \beta_5 \text{black} + \beta_6 \text{asian} \\ & + \beta_7 \text{amindian} + \beta_8 \text{midwest} + \beta_9 \text{south} + \beta_{10} \text{west} + \beta_{11} \text{single} \\ & + \beta_{12} \text{yg2015} + \beta_{13} (\text{nchild} \cdot \text{educclean}) + \varepsilon \end{aligned}$$

Model 7 (nchild_year interaction):

$$\begin{aligned} incwage = {} & \beta_1 + \beta_2 union\_rev + \beta_3 black + \beta_4 asian + \beta_5 female + \beta_6 age + \beta_7 midwest + \beta_8 south \\ & + \beta_9 west + \beta_{10} private\_wkr + \beta_{11} public\_sector + \beta_{12} educ\_rev + \beta_{13} ahrsworkt \\ & + \beta_{14} union\_rev\_public\_sector + e \end{aligned}$$

Models 5, 6, and 7 analyze various interactions between main x variable, *nchild*, and race (*black*), education (*educclean*), and time ($yg2015$). Time is separated into two ranges of years: [2010, 2014], and [2015, 2019].

Through these interaction terms, this paper analyze marginal effects to examine how uhrsworkt varies when race, education, and time change, holding all other independent variables constant. Model 5 focused on the interaction between the number of children and black. The first interaction term in the model is denoted *nchild_black*, and is the product of nchild and black. As stated previously, this paper hypothesized that race will play a role in average hours worked; thus, it is necessary to analyze how the main relationship this paper examines varies across different racial groups. A study conducted by Sandra Florian on racial variation on the effect of motherhood on women's employment found that, when it comes to hours worked according to the number of children a woman has, there are significant discrepancies between ethnicities and racial groups. Specifically, the strongest and longest decline in full-time work after a mother has a child can be seen in the White population. Hispanic women experience a moderate reduction in hours worked, and for Black women, the bearing of a child has the smallest and shortest effect on hours worked compared to other racial groups (Forian, 2018). This difference implies that both race and ethnicity play a major role in women's hours worked, on average, after having a child.

In Model 6, a different interaction term, *nchild_educ*, which is the product of the number of children and level of education was introduced. Higher education is highly correlated with an increase in labor force participation; as of 2025, 72.4% of individuals with a bachelor's degree or higher reported being in the labor force, compared to just 48.1% of those without a high school diploma (U.S. Bureau of Labor Statistics, 2024). Thus, an analysis of how the relationship between number of children and hours worked varies according to different levels of education is paramount to this study.

Third interaction term, $nchild\_yg2015$, which is incorporated in Model 7, is the product of nchild and a dummy variable, $yg2015$, representing the time frame included in this study. The years of study were separated into two categories: 1 if the year is 2015 or greater (maximum of 2019) and 0 if the year is less than 2015 (minimum of 2010). Since fertility rates vary across the years included in this study - being 1.9 from 2010 to 2015, 1.8 from 2016 to 2017, and 1.7 from 2018 to 2019 (World Bank, n.d.) - it is important to examine how the relationship between number of children and hours worked changes in relation to variation in fertility rates across time.

IV. Data and Descriptive Statistics

The data for this study comes from IPUMS CPS, which collects and stores microdata from 1962 until present day from the U.S. Labor Force Survey and the Current Population Survey. To account for major external events that could potentially sway the data, such as the 2008 recession and COVID-19, this paperexamine data collected between the years of 2010 and 2019, inclusive. The data that was used is repeated cross-sectional Annual Social and Economic Supplement (ASEC) data collected through surveys and self-reporting from households. This timeframe was also chosen to account for changes in fertility rates across 2010 to 2019, which are inherently important for analyzing the number of children that the average woman has living in their household. Fertility rates steadily declined during this period, highlighting the trends—delayed childbearing and higher rates of women deciding to not have children–within the United States that followed the 2008 financial crisis (Hamilton et al., 2021). Analyzing these rates is crucial to the analysis, as they play a major role in determining family size and could therefore impact labor force participation.

In defining the main relationship examined, *nchild* was used as primary x (independent) variable and *uhrsworkt* as main y (dependent) variable; *nchild* represents reported children of all ages and forms (biological, adopted, stepchildren, etc.) that reside within the household of everyone in the sample. "*uhrsworkt*" is reported hours spent at all jobs during the week.

Table 1, which can be located on page 1 of the appendix, describes all variables that are used in this study. The nature of each variable - that is, its purpose, as well as whether it is a dummy/continuous variable, is also clarified in this table. In cleaning all variables, observations coded as "NIU" (Not in Universe, code 999, 99, or 9999.99) was dropped. For *uhrsworkt*, code 999 indicated respondents who were not in the labor force. This significantly impacted this research by limiting the analysis to only women who remained employed even after childbirth. Additionally, *uhrsworkt* ¿ 100 was dropped to account for major outliers. Thus, the maximum number of hours worked in this sample is 100. Observations coded as "hours vary" (code 997) were also removed due to their lack of precision, which could potentially obscure trends in hours worked relative to the number of children. By analyzing exact reported weekly hours, this paper aimed to achieve clearer insights into labor force participation trends. Additionally, observations reporting over 100 hours worked weekly were eliminated, identifying these as extreme outliers likely resulting from reporting errors or exceptional circumstances. This helped ensure the robustness and accuracy of this analysis. The education variable (*educ*) originally had inconsistent categorical coding. To address this, it was recoded it into a numerical scale (*educclean*), reflecting completed years of education, thus facilitating a clearer understanding of education's impact on labor force participation.

Several dummy variables were also generated to enhance this regression analysis. *haschild* and *nochild* were created to distinguish respondents with and without children clearly. The variable female explicitly confirmed the inclusion of only female respondents. Regional indicators ($northeast, midwest, south, west$) were introduced to control for geographic variations. Since midwest, south, and west are all included in multiple regression models, the base region is northeast. Dummy variables for race four racial groups - *black*, *amindian* (American Indian), and asian - were created to identify race-based distinctions clearly; other, less specific codes for racial groups (ex. mixed, multiple races) were dropped to keep this results organized across a defined set of groups. As all of these race dummy variables are included in the multiple regression analyses, the base category for race ($black = 0$, $asian = 0$, $amindian = 0$) is White. Employment status was categorized through the creation of fulltime and parttime variables, defined based on reported hours worked; for full-time, the base category is $uhrsworkt < 35$. Lastly, marital status was addressed through the dummy variable single, which is equal to 0 if married and 1 if unmarried, widowed, divorced, or separated. A year dummy variable, $yg2015$ was also generated, to account for variation in fertility rates across the period; the details of this variable can be found in Table 1 as well.

Table 2 contains the descriptive statistics for all of variables. Sample size after cleaning all variables, is 388,852; all of which are women surveyed between the years of 2010 and 2019, inclusive. Mean *uhrsworkt*, or average hours that women in the sample reported spending at all jobs in the last week, is ∼36.69 hours, which falls just above the minimum hours required for a job to be considered full-time. According to the US Bureau of Labor Statistics (2024), women working full time reported an average of about 36.7 hours per week which is consistent with the 36.69 hours observed in the sample. Additionally, the ratio of type of employment of women in the sample can be examined through the fulltime and parttime dummy variables; according to Table 3, 74.5% of women in the sample report working 35 hours or more in the last week. Mean nchild, or average number of children (of all forms: adopted, stepchild, etc.) that women in the sample report having in their household, is 1.005, with 9 as its maximum (since 9 children and 9+ children fall under the same category). The mean level of education in the sample is about 14 years, meaning that the average educational attainment is an associate's degree.

Additionally, most of the women in this sample are from the South, as seen by south dummy variable having the highest frequency compared to the other major regions in the United States (0.339, or 33.9%). Roughly 25% are from the West, 18.5% are from the Northeast, and 22% are from the Midwest (See Table 2). This is interesting, as nearly half (46.7 percent) of people from the southern United States live in rural communities (Davis et al., 2023), while in the Northeast, 84.0 percent of the population resides in urban areas (U.S. Census Bureau, 2021). The fact that the South—the region with the highest representation in the sample—also has the largest rural population in the United States could help explain some of the patterns observed in hours worked.

On average, women in sample have at least one child, as the mean of *haschild* dummy variable (which is either 0 if no child and 1 if there is a child) is 0.541. This study predict that the largest effect of the number of children on hours worked likely takes place between the presence of zero children and one child in the household; thus, the fact that most women in the sample have at least one child residing with them aids in this analysis, as it indicates a nearly perfect 50:50 ratio between women in this sample who have children and women who do not.

Table 2 also contains the standard deviations of all of the variables, as well as their minimum and maximum values. Firstly, and of note, the maximum age of this sample is 79, as individuals who were coded as age 80 and 85 in this sample indicated that the individual was either 80-84 years old, or 85+, respectively. Ages that were unspecified, as such, were dropped from the sample, as age needs to take on exact values to maintain precision in the study. Additionally, as seen in the cleaned education variable (*educclean*), the maximum possible educational attainment is 21 years, which means that an individual completed 21 years of education and received a professional degree or doctorate. Likewise, the minimum is zero years, which means that the individual has either no education, or no more than a kindergarten or preschool education.

Regarding the distribution of races of women in this sample, 79.9% are White, 12.63% Black, 1.27% American Indian, and 6.4% Asian. This can be examined through the means of each race dummy variable. Lastly, and of note, roughly 45% of women in the sample reported being single.

Table 3 displays descriptive statistics based on either the presence or absence of children in the household - that is, Table 3 distinguishes the descriptive statistics between women in the sample who do have at least one child and those who do not. This was generated through the bysort command, as well as creation of a haschild dummy variable (described in Table 1). Specifically, 218,411 women in the sample reported having at least one child in their household, while the other 178,411 reported not having children in their household. The difference in means of hours worked between women with at least one child and women without children is -1.3, meaning that women with at least one child work 1.3 hours more on average compared to those with zero children. The average age of both women with and without children in the sample is roughly the same, with an approximate value of 41.7 years old. Additionally, for all races included in our models, the means are lower for women who do not have children when compared to those who do. This means, for example, that the amount of Black women in samples that do not have children is greater than the amount of Black women who do have children[3]. Other comparisons between regions, racial groups, ages, and more—based on either the presence or absence of at least one child in the household—can be analyzed in Table 3 as well.

V. Estimation Results

In Model 1 of Table 4A, the coefficients of the variables, as well as the levels of significance of each[4], are presented. β_1, or expected hours women with zero children report spending at all jobs in the last week, is

[3] In Table III, we see that the mean for black when haschild = 0 is 0.132; mean for black when haschild = 1 is 0.121. Results described above come from the difference in these two means, which is 0.012.

[4] "***" next to a variable in tables 4 and 5 indicates that a variable is statistically significant at $p < 0.01$; "**" means significant at $p < 0.05$; "*" means significant at $p < 0.1$ (see tables for more information).

$\approx$36.52[5] hours; β_2, or $\approx$0.175, tells us that as the expected number of children that women in the sample have living in their household increases by one, expected hours women report spending at all jobs in the last week increases by 0.175 hours. Based on this, prior to accounting for other variables which can influence hours worked, there is a clear positive relationship between *nchild* and *uhrsworkt*. With this model, any value can be plug in, $0 \leq nchild \leq 9$, to examine how many hours a woman is expected to spend at all jobs per week depending on how many children they have. For example, a woman with two children is expected to spend roughly 36.8674756 hours at all jobs per week[6]. In addition, elasticity at the means of x and y using this model was calculated, as well as Table 2, to determine the percentage change in hours worked after a 1% increase in the number of children a woman has. Elasticity was found to be 0.0048, meaning that, at the mean, a 1% increase in the number of children is associated with a 0.48% increase in hours worked[7].

In Model 2, the intercept is far less than that of the linear model. This is because 1 here represents the log of hours worked when the number of children is zero. Using this model, we find that the estimated hours worked when the number of children is zero is $\sim$ 33.66 [8]. This estimate according to the log-lin is nearly 3 hours less than that of Model 1, suggesting that the standard linear model has some right skewness that the logistic-linear model attempts to correct for. Semi-elasticity through the 2 value of Model 2 was also examined, which is 0.0116, or 1.16%. By interpreting this, we see that an increase in the number of children by 1 is associated with a $\sim$1.16% increase in hours worked.

As seen in Table 4B, β_1, the intercept of *uhrsworkt* is 21.16. This tells us that women's expected hours worked is 21.16 when age is held constant, the woman has no children, no educational attainment, is White, from the Northeast, is not single (either their spouse is present or they have a spouse that is not present), and the year is between 2010 and 2014. This value differs substantially from the intercept of the sample simple regression model, which was 36.517. This difference checks out, as all coefficients in the model are positive (see *Table 4*), and are now accounting for other statistically significant factors that impact hours worked. Holding all other variables constant, as the number of children increases by one (which is roughly the mean number of children for women in the sample), expected hours worked increases by 0.248 hours. Also of note, as age increases by one year, expected hours worked for women in sample increases by 0.097 hours, holding all else constant.

Notably, and also expectedly, education level has a major impact on the number of hours worked; as education level increases by one year, expected hours worked increases by 0.722 hours, controlling for all other x variables[9]. This shows us that aside from the number of children that a woman has, education (in addition to many other variables that are not accounted for as extensively in the study) has a profound effect on labor force participation, when controlling for the other variables in this model.

The races, regions, and years included in this sample are also all statistically significant and have influence on expected hours worked. Specifically, when an individual is Black, or $black = 1$, holding all else constant, expected hours spent at all jobs in the last week increases by 1.228 hours. Additionally, as seen in our model, being American Indian has the largest impact on hours worked out of all the races that are accounted for in the model; expected hours worked increases by 1.310 when an individual is American Indian, holding all else constant. Finally, when all else is controlled for, expected hours worked increases by $\sim$0.188 when examining the years of 2015 and thereafter.

The coefficient of *nchild*, β_2, in the multiple logistic-linear regression displayed in Model 4, tells us that while holding all other x variables constant, each additional child is associated with a 1.363% increase in hours

[5] See Table 4A, Model 1, for exact values.

[6] Setting nchild = 2, 36.51663 + 0.1754228 (2) = 36.8674756

[7] Elasticity, at the means= (Δ_y/Δ_y) * xbar/ybar; 0.1754228 (1.004732/36.69288) = $\sim$0.04

[8] $e^{3.5161} \approx$ 33.66amin

[9] Setting education level at its mean ($\tilde{1}4$) while controlling for the other x variables, we see that expected hours worked is 21.16 + 0.722(14) = $\tilde{3}1.268$ hours.

worked[10]. This value is 0.203% greater than the β_2 of a simple log-linear model, which was 1.16%. Additionally, among other information, Model 4 indicates that when holding other factors constant, Black individuals work approximately 6.1% more hours than White individuals. Also, when controlling for the other x variables, each additional year of education is associated with a 2.4% increase in hours worked[11].

In Models 5, 6, and 7, interaction terms between the number of children and important demographic factors—race, education, and fertility— were included based on research suggesting that the effect of having children on hours worked may vary across these characteristics. One example of past findings that motivated this analysis of these interactions—namely race interaction—came from Angrist's and Evans' study (1998), which found that Black women are often more likely to stay employed compared to White women after having children due to more economic pressure to provide for their children.

Model 5 includes the *nchild_black* interaction term, which allows us to examine how the relationship between number of children and hours worked varies depending on whether a woman is Black or White. Based on the regression analysis, while controlling for other x variables, White women work $\tilde{0}.2036$ more hours per week for each additional child, holding all else constant[12]. For each additional child, Black women work an extra ∼0.3199 hours compared to White women[13], holding all other x variables constant. These results are in line with Angrist's and Evans study, as they suggest that having additional children does not necessarily push Black mothers out of the workforce, but instead may make them more likely to remain employed.

Model 6 includes the *nchild_educ* interaction term, which was used to explore how the relationship between number of children within a household and hours worked changes based on varying levels of education. Through the regression analysis in Model 6, it was found that when controlling for all other x variables, women with no education work approximately 4.37 hours more per week per additional child[14]. Furthermore, this model found that the effect of having children on hours worked decreases as the level of education increases by one year, holding all else constant; This decreases by -0.2994 hours worked per week[15]. Through the effect obtained from Model 5, it shows how each level of education impacts the effect of the number of children on hours worked. For example, the effect of the number of children on hours worked when a woman has completed 7 years of education is $\tilde{2}.280$ more hours worked per week for each additional child, holding all else constant[16]. For a woman who has completed 21 years of education (has a doctorate or professional degree), each additional child is associated with a -1.906 decrease in hours worked per week. This is interesting, as these results indicate that a higher education level for women is associated with a reduction in hours worked when a child/children are present in the household. These results continue to be consistent with Cools', Markussen's, and Strøm's study (2017), which found that women with higher levels of education are more likely to hold demanding and challenging jobs, thereby increasing the likelihood that they will need to take time off work to care for their children. Their study also found that women who have lower education levels are more likely to hold part time jobs or less demanding jobs, allowing them to move around hours more and balance their work needs with raising children. Both results, as well as the findings of Cools, Markussen, and Strøm, underscore the need for more effective workplace policies that aid women with children in balancing work with their personal life (regardless of their job's demand or level of education).

Model 7 includes *nchild_y*2015 interaction term, which was used to explore how the relationship between number of children and hours worked varies according to time period. It found that from 2010 to 2014, the marginal effect of an additional child on hours worked was 0.2107 hours, controlling for the other x variables[17]. From 2015 to 2019 (holding all other x's constant), the effect of an addition child on hours worked was found

---

[10] $\beta_2 * 100 = 0.0136 * 100 = 1.36\%$

[11] See Table 4B for more results

[12] Difference$_{uhrsworkt}$, black = 0 = 0.236 (the marginal effect is 0.2036 + 0.3199*black*)

[13] Addition$_{uhrsworkt,black=0}$ = 1 = 0.2036 + 0.3199 = 0.5235

[14] Difference$_{nchild,educclean=0}$ = 4.3728

[15] Difference$_{nchild,educclean}$ = 4.3728 - 0.2994*educclean* (effect before giving *educclean* a value)

[16] Difference$_{nchild,educclean=7}$ = 7 = 2.280

[17] Addition$_{nchild,yg2015=0}$= 0.2107 + 0.0763(0)

to be 0.2870[18], which is a significant increase compared to the previous time period. This could in part be due to lower fertility rates, which emphasizes a possible trend in women prioritizing their career and stability before having children. According to the CDC, the US fertility rate hit a record low with 1.7 births per woman in 2019 (CDC, 2021), meaning that each additional child likely had less of a marginal impact on a woman's hours worked during that general time period. Based on this, it is possible that women who had children from 2015 to 2019 experienced stronger impacts on their hours worked per additional child—either because of rising childcare costs, fewer family support structures, or more demanding careers; all of these factors are related to changes in fertility rates according to the CDC.

VI. Conclusion and Discussion

This analysis of IPUMS CPS ASEC data from 2010–2019 reveals a small but robust positive relationship between family size and women's hours worked: each additional child is associated with an extra $\approx$0.21 hours of paid work per week in the baseline OLS model and $\approx$0.25 hours per child once controlling for age, education, race, region, marital status, and time period. In semi-elasticity terms, a one-child increase corresponds to a $\approx$1.36% rise in weekly hours worked in the fully specified log-linear model. These estimated effects are statistically significant at p¡0.05, underscoring the reliability of the results.

These findings further add to this literature, as all prior studies to our knowledge have found negative relationships between number of children and workforce participation. As explained in literature review, the novelty of the results could be due to increases in childcare costs, higher rates of stay-at-home fathers since 1990, and most importantly, differences in how to define dependent variables. Because labor force participation can be defined and measured in various ways, it is important to analyze how factors such as number of children can influence not only women's labor supply, but also the type of employment that they pursue (i.e. full-time vs. part-time) and its intensity (e.g. hours worked per week–as used in this study). Further analysis should incorporate these other measures to more accurately examine the impact that children have across all major assessments of labor force participation.

These multiple regression models with interaction terms uncover important heterogeneity across race, education, and time period. Black mothers increase their paid hours by an additional $\approx$0.32 hours per child relative to White mothers, suggesting stronger economic pressures (Model 5). In contrast, and in line with the results of Cools et al. (2017)[19], this paper found that each additional year of schooling attenuates the child-hour effect by $\approx$0.30 hours, with highly educated women (¿16 years) even reducing paid hours as family size grows (Model 6). This paper also find that the child–hours association strengthened modestly after 2015 (by $\approx$ 0.0763 hours per child), coinciding with declining fertility rates and rising childcare costs (Model 7). Although a net increase of 0.25 hours/week per child appears small in isolation, it aggregates to about 13 extra paid hours per year for a two-child family—an amount that could meaningfully bolster household income.

These results have several policy implications. For highly educated women who appear most prone to reducing hours when childrearing demands conflict with high-skill occupations, policies that attempt to reduce the career penalties that stem from motherhood–such as more measures to prevent labor market discrimination based on socioeconomic status and stronger family-leave protections–may help maintain labor force activity when childcare demands are high. For less educated women, policies that aim to reduce work hours when childrearing demands conflict with occupation should be implemented; specifically, expanding access to affordable childcare–potentially by implementing paid family-leave expansions and setting price cap childcare costs based on specific income percentages[20]–could aid less educated mothers in balancing work and family. Additionally,

[18] $\text{Addition}_{nchild,yg2015=1} = 0.2107 + 0.0763$

[19] Cools, Markussen, and Strøm found that the effect (of the birth of a child on mean working hours) constitutes a reduction of 8.6% relative to total hours in college educated women and a 27.6% reduction in women with no college degree.

[20] Sweden, Norway, and Luxemburg have implemented such policies. Specifically, Sweden has a national "maxtaxa" that ensures that parents pay at most 3% of their income for their first child, 2% for their second, and 1% for their third (Statham R., Parkes H., & Nanda S., 2022).

the stronger response among Black mothers underscores the need for equity-focused interventions that recognize disparate economic pressures among demographic groups. Specifically, policies aimed at mitigating unequal access to childcare and lessening racial disparities in paid leave could reduce the disproportionate effect that additional children have on Black mothers[21]. Similar childcare income-based price caps to the ones described in the previous paragraph and in footnote 19 can be implemented for these purposes as well.

Finally, because this paper focus on paid hours, the unpaid caregiving time and shifts in job quality or flexibility that often accompany growing family size likely evolve even more dramatically, suggesting that the estimates capture only part of the motherhood trade-off. Moreover, since *uhrsworkt* is self-reported, these estimates may be subject to recall bias, and the cross-sectional nature of the CPS prevents us from observing individual longitudinal trajectories or fully addressing endogeneity in fertility choices. Additionally, CPS data may be prone to autocorrelation, as data is taken from the same households over time.

Future studies should employ instrumental-variables strategies (e.g., sex-mix instruments) or fixed-effects panel designs and incorporate wage data to measure endogeneity in family choices. Since the decision for a woman to have a certain number of children is often planned, using these strategies can help account for specific variables that can influence these plans (such as health and income). Additionally, implementing instrumental-variable strategies can allow for better, more accurate estimates of the main relationship. Job flexibility metrics and longitudinal career trajectories should also be included, as they can allow for better differentiation between different quality levels of career–i.e. part time (more unstable) work versus full time (quality, more stable) work. Lastly, evaluating how state-level policies—such as childcare subsidies and leave generosity—moderate the effect of the number of children a woman has on their hours worked could offer actionable insights for the creation of more supportive labor-family frameworks.

[21] The Big Back Better legislation is one such program (that relies on programs and policies with substantial track records) that is partially designed to address major gaps in paid leave opportunities and caregiving infrastructure–particularly for workers of color–in the US (Trisi et al., 2021)

References

Angrist, J. D., & Evans, W. N. (1998). Children and their parents' labor supply: Evidence from exogenous variation in family size. *American Economic Review*, *88*(3), 450–477. https://www.jstor.org/stable/116844

Bloom, D. E., Canning, D., Fink, G., & Finlay, J. E. (2007a). Fertility, female labor force participation, and the demographic dividend. *National Bureau of Economic Research*. https://doi.org/10.3386/w13583

Bloom, D. E., Canning, D., Fink, G., & Finlay, J. E. (2007b). *Fertility, female labor force participation, and the demographics dividend* (Working Paper No. 13583). National Bureau of Economic Research. https://www.nber.org/system/files/working_papers/w13583/w13583.pdf

Borrett, A. (2025). Why having a baby remains a career obstacle for women. *Financial Times*. https://www.ft.com/content/254582f8-6253-4cb3-9bd0-f6ae82ac0125

Bureau of Labor Statistics. (2024). Unemployment rate inches up during 2023, labor force participation rises. *U.S. Department of Labor, Monthly Labor Review*. https://www.bls.gov/opub/mlr/2024/article/unemployment-rate-inches-up-during-2023-labor-force-participation-rises.htm

Center for Disease Control and Prevention. (2025). Births: Final data for 2023. https://www.cdc.gov/nchs/data/nvsr/nvsr74/nvsr74-3.pdf

Centers for Disease Control and Prevention. (2018). Vital statistics rapid release: Report no. 8. https://www.cdc.gov/nchs/data/vsrr/vsrr-8-508.pdf

Centers for Disease Control and Prevention. (2020). Vital statistics rapid release: Report no. 20. https://www.cdc.gov/nchs/data/vsrr/vsrr020.pdf

Centers for Disease Control and Prevention. (2021). Births: Final data for 2019. *National Vital Statistics Reports*, *70*(2). https://www.cdc.gov/nchs/data/nvsr/nvsr70/nvsr70-02-508.pdf

Centers for Disease Control and Prevention. (2022). Births: Final data for 2020. https://www.cdc.gov/nchs/data/nvsr/nvsr70/nvsr70-02-508.pdf

Cools, S., Markussen, S., & Strøm, M. (2017). Children and careers: How family size affects parents' labor market outcomes in the long run. *Demography*, *54*(5), 1773–1793. https://doi.org/10.1007/s13524-017-0612-0

Davis, J. C., Cromartie, J., Farrigan, T., Genetin, B., Sanders, A., & Winikoff, J. B. (2023). Rural america at a glance: 2023 edition (eib-261). *U.S. Department of Agriculture, Economic Research Service*. https://doi.org/10.32747/2023.8134362.ers

Dustmann, C., & Fabbri, F. (2017). Children and careers: How family size affects parents' labor market outcomes in the long run. *Demography*, *54*(5), 1773–1794. https://pmc.ncbi.nlm.nih.gov/articles/PMC5624987/

Flood, S., et al. (2024). Integrated public use microdata series, current population survey: Version 12.0.

Florian, S. M. (2018). Racial variation in the effect of motherhood on women's employment: Temporary or enduring effect? *Social Science Research*, *73*, 80–91. https://doi.org/10.1016/j.ssresearch.2018.02.012

Fry, R. (2023). Almost 1 in 5 stay-at-home parents in the u.s. are dads [Retrieved Oct. 19, 2025]. *Pew Research Center*. https://www.pewresearch.org/short-reads/2023/08/03/almost-1-in-5-stay-at-home-parents-in-the-us-are-dads/

Joughin, C. (2021). Economic data underscores the need for significant, sustained investment in child care and early learning [Retrieved Oct. 19, 2025]. *First Five Years Fund*. https://www.ffyf.org/resources/2021/12/latest-economic-data-underscores-the-need-for-significant-sustained-investment-in-child-care-and-early-learning/

Kahn, J. R., Garcia-Manglano, J., & Bianchi, S. M. (2014). The motherhood penalty at midlife: Long-term effects of children on women's careers. *Journal of Marriage and Family*, *76*(1), 56–72. https://doi.org/10.1111/jomf.12086

LendingTree. (2025). How much it costs to raise a child today. https://www.lendingtree.com/debt-consolidation/raising-a-child-study/

National Center for Biotechnology Information. (2021). *Current trends in fertility and childbearing*. https://www.ncbi.nlm.nih.gov/books/NBK588545/

Rector, R. (2021). U.s. fertility is declining due to delayed marriage and childbearing. *The Heritage Foundation*. https://www.heritage.org/marriage-and-family/report/us-fertility-declining-due-delayed-marriage-and-childbearing

Smith, K. (2007). Employment rates higher among rural mothers than urban mothers. https://doi.org/10.34051/p/2020.32

Statham, R., Parkes, H., & Nanda, S. (2022). Towards a childcare guarantee [Retrieved Oct. 21, 2025]. *Institute for Public Policy Research*. https://ippr-org.files.svdcdn.com/production/Downloads/towards-a-childcare-guarantee-sept-22.pdf

Trisi, D., Mazzara, A., Lueck, S., et al. (2021). House build back better legislation advances racial equity [Retrieved Oct 21, 2025]. *Center on Budget and Policy Priorities*. https://www.cbpp.org/research/poverty-and-inequality/house-build-back-better-legislation-advances-racial-equity

U.S. Bureau of Labor Statistics. (2024). Employment characteristics of families - 2023. https://www.bls.gov/news.release/pdf/famee.pdf

U.S. Census Bureau. (2021). Urban and rural classification and urban area criteria [Retrieved March 15, 2025]. https://www.census.gov/programs-surveys/geography/guidance/geo-areas/urban-rural.html

U.S. Department of Labor. (n.d.-a). Employment characteristics of families - 2023 [Retrieved March 20, 2025]. https://www.bls.gov/news.release/pdf/famee.pdf

U.S. Department of Labor. (n.d.-b). Family and medical leave act (fmla) [Retrieved March 20, 2025]. https://www.dol.gov/agencies/whd/fmla

VII. Appendix

Table 1: Description of Variables

Variable Name	Variable Description
nchild	Number of children, regardless of form (biological, adopted, step-child) residing with the surveyed individual.
uhrsworkt	Usual amount of hours spent at all jobs per week reported by the individual, across an unspecified period of time
age	Age
educ	Educational attainment measured at the individual's highest level of school/degree completed
marst	Marital status of each person (includes whether or not the spouse lives in the same household as the surveyed individual)
region	Region of US
northeast	Dummy variable which is 1 if region is northeast and 0 if region is not northeast; specifically: 1 for women from the northeast and 0 for those not.
west	Dummy variable which is 1 if region is west and 0 if the region is not west
south	Dummy variable which is 1 if region is south and 0 if region is not south
midwest	Dummy variable which is 1 if region is midwest and 0 if region is not midwest
female	Dummy variable which is 1 if sex is female and 0 if sex is not.
educclean	Cleaned version of educ variable; number of years of schooling completed
black	Dummy variable which is 1 if race is Black and 0 if race is not black
asian	Dummy variable which is 1 if race is Asian and 0 if otherwise.
amindian	Dummy variable which is 1 if race is American Indian and 0 if not.
fulltime	Dummy variable for uhrsworkt which is 1 if uhrsworkt is ≥ 35 hours per week at all jobs and 0 if uhrsworkt is < 35 hours
parttime	Dummy variable for uhrsworkt which is 1 if uhrsworkt is < 35 but > 0 at all jobs, and 0 if uhrsworkt is ≥ 35
single	Dummy variable which is 1 if marst (marital status) is > 2 (not married) and 0 if otherwise
haschild	Dummy variable which is 1 if nchild (number of children) is > 0 and 0 if nchild is 0
year	year
yg2015	Dummy variable which is 1 if year is > 2015 and 0 if otherwise
luhrsworkt	log of uhrsworkt, for log-lin models
nchild_black	Interaction term for nchild and black
nchild_educ	Interaction term for nchild and educclean
nchild_yg2015	Interaction term for nchild and yg2015

Table 2: Descriptive Statistics of Relevant Variables

VARIABLES	(1) **N**	(2) **mean**	(3) **sd**	(4) **min**	(5) **max**
age	388,852	41.71	13.47	15	79
nchild	388,852	1.005	1.150	0	9
uhrsworkt	388,852	36.69	11.57	0	100
haschild	388,852	0.541	0.498	0	1
educclean	388,852	14.07	2.758	0	21
northeast	388,852	0.185	0.388	0	1
midwest	388,852	0.219	0.414	0	1
south	388,852	0.339	0.473	0	1
west	388,852	0.257	0.437	0	1
black	388,852	0.126	0.332	0	1
asian	388,852	0.0640	0.245	0	1
amindian	388,852	0.0127	0.112	0	1
fulltime	388,852	0.745	0.436	0	1
parttime	388,852	0.254	0.435	0	1
single	388,852	0.447	0.497	0	1
yg2015	388,852	0.484	0.500	0	1
nchild_black	388,852	0.121	0.526	0	9
nchild_educ	388,852	14.12	16.52	0	162
nchild_yg2015	388,852	0.478	0.939	0	9

Table 3: Descriptive Statistics According to haschild

VARIABLES	(1)	(2)	(3)	(4)	***	(5)
	haschild = 0		**haschild = 1**		***	
	N	**mean**	**N**	**mean**	***	**diff**
age	178,441	41.66	210,411	41.75	******	−0.09
uhrsworkt	178,441	35.99	210,411	37.29	******	−1.30
educclean	178,441	14.01	210,411	14.11	******	−0.10
northeast	178,441	0.1830	210,411	0.1870	******	−0.0040
midwest	178,441	0.2130	210,411	0.2240	******	−0.0110
south	178,441	0.3430	210,411	0.3360	***	0.0070
west	178,441	0.2620	210,411	0.2520	***	0.0100
black	178,441	0.1330	210,411	0.1210	***	0.0120
single	178,441	0.6070	210,411	0.3110	***	0.2960
asian	178,441	0.0669	210,411	0.0615	***	0.0054
amindian	178,441	0.0122	210,411	0.0131	******	−0.0009
yg2015	178,441	0.4950	210,411	0.4750	***	0.0200

Table 4A: Simple Regression Models

	(Model 1) Simple Regression Lin Lin	(Model 2) Simple Regression Log Lin
VARIABLES	**uhrsworkt**	**luhrsworkt**
nchild	0.175***	0.0116***
	(0.0161)	(0.000638)
Constant	36.52***	3.516***
	(0.0246)	(0.000974)
Observations	388,852	388,441
R-squared	0.000	0.001

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

Table 4B: Multiple Regression Models

	(Model 3) Multiple Regression Lin-Lin	(Model 4) Multiple Regression Log-Lin
VARIABLES	**uhrsworkt**	**luhrsworkt**
nchild	0.248***	0.0136***
	(0.0166)	(0.000659)
age	0.0968***	0.00355***
	(0.00140)	(5.56e-05)
educclean	0.721***	0.0241***
	(0.00667)	(0.000265)
black	1.228***	0.0610***
	(0.0571)	(0.00227)
asian	0.934***	0.0356***
	(0.0754)	(0.00300)
amindian	1.304***	0.0627***
	(0.163)	(0.00646)
midwest	0.615***	0.0173***
	(0.0573)	(0.00228)
south	1.452***	0.0577***
	(0.0529)	(0.00210)
west	0.487***	0.0166***
	(0.0558)	(0.00222)
single	0.0611	-0.00406**
	(0.0405)	(0.00161)
yg2015	0.187***	0.00668***
	(0.0363)	(0.00144)
Constant	21.16***	2.988***
	(0.129)	(0.00511)
Observations	388,852	388,441
R-squared	0.049	0.040

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

Table 5: Multiple Regressions With Interaction Terms

	(Model 5) *nchildblack* interaction	(Model 6) *nchildeducclean* interaction	(Model 7) *nchildyg*2015 interaction
VARIABLES	**uhrsworkt**	**uhrsworkt**	**uhrsworkt**
nchild	0.204***	4.373***	0.211***
	(0.0178)	(0.0765)	(0.0225)
age	0.0967***	0.0945***	0.0967***
	(0.00140)	(0.00139)	(0.00140)
educclean	0.722***	1.046***	0.721***
	(0.00667)	(0.00888)	(0.00667)
black	0.922***	1.213***	1.229***
	(0.0723)	(0.0569)	(0.0571)
asian	0.929***	0.834***	0.934***
	(0.0754)	(0.0751)	(0.0754)
amindian	1.310***	1.275***	1.302***
	(0.163)	(0.162)	(0.163)
midwest	0.617***	0.639***	0.614***
	(0.0573)	(0.0571)	(0.0573)
south	1.453***	1.405***	1.451***
	(0.0529)	(0.0527)	(0.0529)
west	0.489***	0.407***	0.487***
	(0.0558)	(0.0556)	(0.0558)
single	0.0481	-0.0497	0.0610
	(0.0405)	(0.0404)	(0.0405)
nchild_black	0.320***		
	(0.0465)		
yg2015	0.188***	0.203***	0.111**
	(0.0363)	(0.0362)	(0.0481)
nchild_educ		-0.299***	
		(0.00542)	
nchild_yg2015			0.0763**
			(0.0315)
Constant	21.20***	16.85***	21.20***
	(0.129)	(0.150)	(0.130)
Observations	388,852	388,852	388,852
R-squared	0.049	0.056	0.049

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

The Citizen Wage Gap: Income Differences Between Native-Born and Naturalized Americans

By Keane Adam, Daniel Demtschenko, Matthew Yeo

Income inequality across citizenship groups remains an important topic in discussions of immigrant integration in the United States. This paper examines whether naturalized U.S. citizens and birthright citizens differ in their annual income levels, using microdata from the Current Population Survey (CPS) from 2004 to 2011. The goal of this study is to document and quantify the association between citizenship type and income while controlling for demographic and economic characteristics. Using log-linear multiple regression models, this paper accounts for education, hours worked, years spent in the United States, disability status, marital status, metropolitan residence, and others. Interaction terms are included to explore whether the association between naturalization and income varies by hours worked or by living in a metropolitan area. Results show that naturalized citizens who live in a metropolitan area and work 40 hours per week are associated with earning 21.29% higher income than native-born Americans. Naturalized citizens who don't live in a metropolitan area earn 28.5% greater income, given a 40-hour work week. These findings contribute to the literature by shifting the focus from comparisons between naturalized and non-citizen immigrants to comparisons between two groups of full citizens. Understanding these income differences provides new context for discussions on immigrant integration, the economic characteristics of naturalized citizens, and how labor market outcomes vary across citizenship types in the United States.

I. Introduction

The idea of the American dream that promises prosperity and success in exchange for hard work and dedication has defined US history for decades, attracting millions of immigrants every year. Although many visa categories permit people to enter the US, only two options allow immigrants to remain permanently in the country: obtaining a Green Card (Permanent Resident Card) or securing citizenship status through naturalization. Of these two options, naturalization, at least on paper, offers the same benefits as a citizen who was born in the US. However, whether this equality is reflected in income data is the motivation of our research. This topic is especially relevant at the time of publication, as recent policies enacted by the Trump administration around immigration and birthright citizenship have reignited national debates over who deserves access to the rights and privileges of full citizenship.

To analyze this underlying question, this paper examines the income disparities between naturalized citizens and those with birthright citizenship. Using Annual Social Economic Supplement (ASEC) data collected from the Current Population Survey (CPS) from 2004 to 2011, this paper finds whether naturalized citizens face differences in income relative to citizens born in the United States. Understanding this disparity is crucial because it explores whether naturalized citizens, despite achieving full legal status, experience differences in income that may reflect systemic barriers to economic mobility. Economic theory also provides several reasons to expect potential income differences between the two groups. Human capital theory suggests that naturalized citizens may differ in education, work experience, or skill accumulation, all of which directly influence earnings. In addition, selection theory argues that immigrants who choose to naturalize are often a non-random and highly motivated group, which may lead them to have different labor market outcomes than birthright citizens. If naturalized citizens earn consistently less, it could indicate deeply rooted social issues such as discrimination or unequal labor market access. By contrast, if the opposite is true, it could signal that the naturalization system selects individuals who are already high-income earners.

Overall, the empirical models in this paper consistently find that naturalized citizens are associated with earning higher average incomes than birthright citizens in the CPS sample. In the multiple regression that includes demographic, geographic, and human capital controls, the naturalization coefficient remains positive and statistically significant. These associations suggest that individuals who naturalize may differ from birthright citizens in characteristics such as educational background, labor market attachment, and time spent accumulating human capital in the United States. As the barriers to naturalization rise, the individuals who successfully naturalize may also represent a more motivated and highly selected group, which could relate to the higher incomes observed in the data. These patterns, however, reflect correlation rather than causal effects. This paper contributes to the existing literature by shifting the focus away from the well-studied comparison between naturalized and non-citizen immigrants and instead examining income differences between two types of U.S. citizens. While previous studies have concentrated on how naturalization affects earnings within the immigrant population, far fewer have analyzed how naturalized citizens compare directly to birthright citizens once key economic and demographic characteristics are controlled for. By using CPS data from 2004 to 2011 and applying log-linear multiple regression models with interaction terms, this paper provides new descriptive evidence on whether the earnings of naturalized citizens and birthright citizens diverge.

II. Literature Review

Research on citizenship and earnings has largely focused on differences within the immigrant population, especially between naturalized and non-naturalized immigrants. A central contribution in this area is Bratsberg et al. (2002). Their paper investigated whether wage growth after naturalization is driven by the change in legal status itself or by immigrants' human capital investment prior to becoming citizens. Using panel data from the National Longitudinal Survey of Youth and a fixed effects regression, they were able to track the same individuals before and after naturalization. This method allowed them to show that wage acceleration occurs only after citizenship is obtained, which challenged the earlier belief that immigrants increased their wages mainly through pre-naturalization investments. While this finding is important, the focus of the paper is limited to comparing naturalized immigrants to immigrants who have not naturalized. It does not compare naturalized immigrants to American citizens by birth, the focus of this paper.

A related study by Hall, Greenman, and Farkas (2010) looked at wage disparities among Mexican immigrants by legal status. Their analysis separated undocumented immigrants, documented immigrants, native-born Mexican Americans, and non-Latino citizens using data from the Survey of Income and Program Participation. They found a 17% wage gap between documented and undocumented Mexican men, and a 9% gap for women. They also discovered lower returns to human capital and slower wage growth for undocumented workers in their research. As expected, these findings support the idea that legal status shapes labor market access and wage growth. Similar to Bratsberg et al., however, the comparison remains within immigrant groups and does not address income differences between naturalized and birthright citizens. These outcomes are expected, since naturalization generally opens access to better jobs, especially in white-collar and public-sector employment. What remains less clear is whether naturalized immigrants experience higher or lower wages than birthright citizens, who theoretically face the same legal rights and economic opportunities. This is the central question that serves as motivation for this paper, but existing literature has not directly answered it.

One descriptive source that begins to touch on this comparison is an Economic Policy Institute briefing paper by Shierholz (2010). Using CPS data, Shierholz reported that in 2007, adult citizen immigrants had a median family income of US$57,823, which was slightly higher than the US$56,000 dollars earned by native-born adult Americans. Although the paper does not use econometric models or control factors that influence income, it highlights a potentially meaningful pattern that motivates further study. Overall, the existing research has clearly established that naturalized immigrants tend to earn more than non-citizen immigrants. What has not been studied nearly as extensively is the income relationship between the two types of citizens in the United States. Our paper fills this gap by comparing naturalized citizens to birthright citizens using CPS data from 2004 to 2011. By applying simple log-linear and multiple log-linear regression models with interaction terms, our study builds the foundations of prior work while extending the literature to a new and relatively unexplored comparison.

III. Data and Descriptive Statistics

The data for this study comes from Integrated Public Use Microdata Series - Current Population Survey (IPUMS CPS), from which INCTOT served as the dependent variable, and CITIZEN as the independent variable. CPS data is collected from select households who are surveyed for 4 months, then out for 8 months, then back in for another 4 months, specifically targeting civilian, non-institutionalized populations of the US. The chosen data range is from 2004 to 2011. This data range was chosen for a sharp spike in naturalization rates which warranted further investigation. In 2008, a relative peak in naturalizations was observed, which only lasted one year before dropping again in 2009 (Batalova, 2021). This observation raises interest to dig deeper into this time period.

The final cleaned sample consists of 606,206 observations, with 561,942 native-born and 44,264 naturalized citizens[1]. Therefore, the data identified 7.3% of those people to be naturalized citizens of the US and 92.7% to be birthright citizens. The dependent variable is the annual pre-tax personal income, and the key independent variable is a binary indicator for naturalized citizenship (naturalized). Additional controls include demographic, educational, and geographical variables. Those who are U.S.-born citizens have a mean annual income of $36,538.93 per year, while naturalized citizens make $41,397.06 per year. For U.S.-born citizens, observations typically vary by $55,381.59 from the mean annual income, while observations for naturalized citizens typically vary by $64938.88 from their mean annual income. Both types of citizenship share the same minimum annual

[1] All variables were cleaned to remove top codes, missing values, and non-relevant observations. Dummy variables were generated for key demographic and regional controls and standardized continuous variables where applicable. Detailed definitions appear in Appendix Table 1.1.

income, \$0 per year, while their maximums vary. Native born citizens have a maximum annual income of \$2,950,301 per year, while naturalized citizens have a maximum annual income of \$2,864,712 per year. While naturalized citizens have a higher mean annual income, they do not have a higher maximum annual income than native-born citizens of the United States.

Table 2 shows the descriptive statistics for all the variables. While the mean income is \$46,700.73, there is a large standard deviation of \$53,257.10 and a maximum income of around \$1.7 million, which suggests that the distribution is highly right-skewed. This skewed income distribution justifies the use of a log-linear simple regression model to reduce heteroskedasticity. The table also shows only 7.3% of all citizens in the sample obtained their citizenship by naturalization. The gender variable is split roughly evenly, 50.3% of all people in the sample being male, 61.2% married, and 53.8% with children. The mean level of years of education is around 13.8 years, slightly above a high school diploma. Disability is quite low, around 2% of all people in the sample have a disability, which impacts work, even though the CDC reports a much higher number for general disability (CDC Newsroom, 2016). In terms of region, 78.5% of respondents live in metropolitan areas, and the most respondents, 30.5%, reside in the southern region. The timing of the survey plays a critical role in this research, and according to the table, 38.1% of the respondents were surveyed before 2006, 25.5% were surveyed during the peak period (2006-2008), and 36.1% after 2008.

Table 3 reports mean values of key demographics and economic variables for naturalized and native-born US citizens. Several meaningful differences emerge across groups, highlighting the importance of controlling for characteristics like these in the regression analysis. Naturalized citizens make \$2,360 more on average than native-born citizens. This suggests a positive raw income premium for naturalized citizens. Interestingly, naturalized citizens receive slightly less average education than native-born, 13.62 years and 13.85 years, respectively. As expected, naturalized citizens have considerably fewer average years in the United States than native-born citizens, 23.16 and 40.90 years respectively. For naturalized citizens, marriage rates and the likelihood of having children are higher by over 10%, which could indicate that people tend to migrate to the US already as a family. Naturalized citizens are overrepresented in the West (38.0%) and underrepresented in the South (26.2%), compared to native-born (22.6% West, 30.8% South). This reflects known migration patterns, with immigrants more likely to settle in coastal and urban areas. This is also supported by the fact that 94.2% of naturalized people live in metropolitan areas, compared to only 77.3% of native-born; a difference of 17 percentage points. Overall, naturalized citizens earn more despite having slightly less education and fewer years in the US, which could suggest strong labor market assimilation

IV. Econometric Model

$$\widehat{linctotrev} = \beta_1 + \beta_2 naturalized + e$$

A log-linear functional form was determined to be the best method to explore the relationship between citizenship status and income, as the independent variable is a dummy variable and this functional form addresses the heavily right-skewed data. Because of the log transformation, the coefficients of variables represent percent differences in income. This simple log-linear model only predicts log income *linctot_rev* from *naturalized* with no controlling variables. As a result, substantial variation is unexplained. β_1 represents the expected value of log income for native-born citizens, while β_2 represents the percent difference in income between native-born and naturalized citizens, that is, the naturalization premium.

$$2.\ \widehat{linctot\_rev} = \beta_1 + \beta_2\, naturalized + X_i'\beta + e$$

By adding 16 controlling variables, this multiple log-linear regression model accounts for demographic and human-capital relationships on income to more precisely isolate the association between naturalization status and log income, holding these characteristics constant. Thus, the model controls for other variables that influence income.

β_1 and β_2 hold the same meaning as compared to the simple log-linear model, but only by holding all other variables constant. X is the vector of the controlling variables and $X_i'\beta$ captures the effects of race, marriage status, years of education, disability status, region, number of kids, whether someone lives in a metropolitan area, hours worked, years in the US, before and during a peak in naturalizations, and age on the log of income *linctôt_rev*. The specific controlled variables are outlined in Table 1.1.

The omitted variables selected are based on human capital theory, which suggests that factors such as education, experience, and demographic characteristics determine an individual's earnings potential. The model controls for many of these key determinants which help in reducing the omitted variable bias.

$$linct\hat{o}t\_rev = \beta_1 + \beta_2\, naturalized + \beta_3\, naturalized\_uhrsworkt + \beta_4\, naturalized\_metrop + X_i'\beta + e$$

To arrive at the final model, interaction terms are introduced to the multiple log-linear model. This model controls for demographic and human-capital variables, and the analysis aims to capture differences in the effects of naturalization across contexts. The use of interaction terms offers valuable insight on the impact of naturalized citizenship by other characteristics. Accounting for the marginal effects of relevant variables and maximizing control for endogeneity by including omitted variables, the multiple regression is optimized to prevent omitted variable bias and provide greater insights

In this log-linear multiple regression model 2 is the baseline naturalization premium when metrop and *uhrsworkt_rev* are zero, holding all other variables constant. For interaction terms, this paper included two key interaction effects: *naturalized_uhrsworkt* and *naturalized_metrop*. *naturalized_uhrsworkt*, measured by 3, allows the model to capture how the return to an additional hour of work differs between naturalized and native-born citizens. Likewise, *naturalized_metrop*, measured by 4, captures how the association between metropolitan residence and income differs for naturalized citizens compared to native-born citizens. Taken together, these coefficients allow the premium to vary across labor-market conditions.

The interaction variable *naturalized_uhrsworkt* is the product of *uhrsworkt_rev* and *naturalized*. To justify this variable's inclusion in the model, the following possibility is considered: at a constant number of hours worked per week, a naturalized citizen may have a lower income than a native-born citizen because the naturalized citizen could have less economic and social integration than the native-born citizen. This aligns with the descriptive statistics, which show that naturalized and native-born citizens differ in mean hours worked, years in the United States, years of education, and other factors related to labor-market integration. With higher average hours worked and fewer average years in education and in the United States, naturalized citizens may experience lower returns to an additional hour of work. Therefore, naturalized citizens who work a high number of hours per week may have a lesser marginal effect on their income compared to native-born individuals who work the same number of hours per week.

The interaction variable *naturalized_metrop* is the product of naturalized and metrop. This variable's inclusion in the model is justified by the following rationale: for a naturalized citizen, living in a metropolitan area may imply a higher income compared to a naturalized citizen who does not live in a metropolitan area, as metropolitan areas typically have higher costs of living and thus are compensated with higher wages and incomes. Furthermore, this term addresses the disproportionate geographic clustering of naturalized citizens in metropolitan areas found in the descriptive statistics. As a result of this clustering, in terms of ratios, naturalized citizens participate much more heavily in urban labor markets and wage structures compared to the native-born citizens. Thus, naturalized citizens who work in a metropolitan area may exhibit a greater marginal effect on their income compared to naturalized citizens who live outside of a metropolitan area.

Table 1.1: Description of Final Variables

Variable Name	**Variable Description**
inctot_rev (income)	Reports each respondent's total pre-tax personal income from all sources for the previous calendar year, given their income is greater than or equal to 0.
linctôt_rev (log income)	This variable is the natural log of inctot_rev.
yrs_us (Years in the US)	Continuous variable that represents the number of years between when a respondent "came to the U.S. to stay" and the year of the CPS survey.
naturalized	Dummy variable that describes the type of citizen of mainland US a person is. 1 = a naturalized citizen, 0 = a birthright citizen.
married	Dummy variable that states whether a person is married or not married. 1 = a married individual, 0 = an unmarried individual.
age	Gives each person's age at their last birthday. Cleaned to remove top codes in given years.
male	Dummy variable that describes the gender of each person. 1 = male, 0 = female.
race	Gives each person's race, cleaned to only include white (100), black (200), and Asian (651) for simplicity purposes.
white	Dummy variable that describes the race of each person. 1 = white, 0 = not white.
black	Dummy variable that describes the race of each person. 1 = black, 0 = not black.
asian	Dummy variable that describes the race of each person. 1 = asian, 0 = not asian.
midwest	Dummy variable that describes the region of each person. 1 = from Midwest, 0 = not from Midwest.
south	Dummy variable that describes the region of each person. 1 = from South, 0 = not from South.
west	Dummy variable that describes the region of each person. 1 = from West, 0 = not from West.
educ_rev	Continuous variable that measures highest year of school or degree completed.
disabled	Dummy variable that represents whether someone's disability impacts their ability or type of work. 1 = disabled, 0 = not disabled.
Metrop (Metropolitan Area)	Dummy variable that describes whether a person lives in a metropolitan area. 1 = lives in a metropolitan area, 0 = does not live in a metropolitan area.
kids	Dummy variable that describes whether the person has kids or not. 1 = has kids, 0 = does not have kids.
uhrs_workt_rev (hours worked)	Reports the usual number of hours per week the respondent reports being at all jobs.
before_peak (before peak)	Dummy variable that describes whether the survey information is from before the peak (2006) or not. 1 = before peak, 0 = not during peak.
during_peak (during peak)	Dummy variable that describes whether the survey information is from during the peak (2006-2008) or not. 1 = during peak, 0 = after peak.
naturalized_uhrsworkt	Interaction term between naturalized and uhrs_workt_rev that allows the effect of hours worked differ for naturalized and native-born citizens.
naturalized_metrop	Interaction term between naturalized and metrop that allows the effect of living in a metropolitan area differ for naturalized and native-born citizens.

V. Sample Regression Model and Estimation Results

$$\widehat{linctot\_rev} = 10.19 naturalized$$

Equation 4 is the sample regression model for the simple log-linear regression with the coefficients obtained form Table 4. According to the model, naturalized citizens earn approximately 10.6% higher income than native-born citizens on average. The intercept of 10.19 represents the expected log income of native-born citizens. The R-squared value is approximately zero which is typical for a regression based on one dummy variable and is a reflection that log income has a substantial amount unexplained variation in this model based on unconditional association.

$$\widehat{linctot\_rev} = 5.377 + 0.218 naturalized + X_i'\beta$$

Equation 5 with values from Table 4 is the sample regression model for the multiple log-linear regression without interaction variables. 5.377 represents the expected log income for the baseline category, that is: individuals who are native-born, black, female, unmarried, non-disabled, northeastern-residing, non-metropolitan, and were surveyed after the peak in naturalizations. Because continuous variables are set to zero in the baseline, the intercept does not carry substantive economic value. Compared to Equation 4, this model has a higher naturalization premium at 0.218, reflecting that naturalized citizens are associated with a 21.8% higher income than native-born citizens. Thus, introducing demographic and human capital variables results in a larger conditional association between naturalization and log income. The R-squared value of this model is approximately 0.279, with the increase from the previous model indicating that this model explains more variation in log income than the model given by Equation 4.

$$\widehat{linctot\_rev} = 5.359 + 0.619 naturalized - 0.00835 naturalized\_uhrsworkt - 0.0721 naturalized\_metrop + X_i'$$

Equation 6 is the final sample regression model for the multiple log-linear regression with the interaction variables. The values of this model were obtained in Table 4. β_1, 5.359, is the baseline expected log income when all variables are set to zero, though it is not economically meaningful nor is it the focus of the paper. A naturalized β_2 coefficient of 0.619 indicates that naturalized citizens are associated with 61.9% higher income than native-born citizens in the baseline case explicitly stated above, which corresponds to individuals living outside a metropolitan area and holding all other variables constant.

The β_3 coefficient, -0.00835, represents the marginal effect on the percentage return to income from an additional hour of work for naturalized citizens relative to native-born citizens. Hence, the estimated percentage return to each additional hour of work is 0.835 percentage points lower than that of native-born citizens. While naturalized citizens have a higher baseline income, the incremental gain from an additional hour of work is flatter, potentially suggesting a declining naturalization premium in hours worked.[2] Correspondingly, the coefficient for 4 -0.0721 indicates that the metropolitan premium is about 7.21% smaller for naturalized citizens compared to native-born citizens.[3] Because β_4 is negative, it suggests the possibility that naturalization premium is smaller in metropolitan areas.

$$\frac{\partial \widehat{linctot}_{rev}}{\partial naturalized} = 0.619 - 0.00835\, uhrsworkt_{rev} - 0.0721\, metrop$$

[2]One speculative explanation not tested in the paper is that if hourly returns reflect productivity, native-born citizens may have greater marginal productivity than birthright citizens perhaps due to an institution which favors them.

[3]Another speculative explanation not tested in this paper: the negative interaction coefficient may reflect an increased concentration of blue-collar or lower-wage occupations among naturalized citizens, lowering their relative income premium in percentage terms compared to native-born metropolitan citizens.

Equation 7 is the partial derivative of the multiple log-linear regression with interaction terms with respect to naturalized. β_2, 0.619, is the baseline naturalization premium which is adjusted by β_3 $uhrsworkt - rev$ and β_4 *metrop* respectively. From this derivative, the total naturalization premium can be found at any given number of hours worked per week and metropolitan residence status.

Although the interaction term of being a metropolitan naturalized citizen is negative, the total effect, $\beta_2+\beta_4$, is still positive at approximately 54.69%, representing the total estimated percent increase in income for naturalized metropolitan citizens compared to non-metropolitan native-born individuals when hours worked are set to zero.

Likewise, the interaction coefficient of naturalized and weekly hours worked is negative, but the total effect $\beta_2 + \beta_3$ is still positive for reasonable work weeks. Fixing $uhrsworkt_rev$ at 40, the standard fulltime work week, and setting *metro* equal to zero, the naturalization premium is equal to 0.285. This means that at 40 hours per week, a non-metropolitan naturalized citizen is associated with a 28.5% higher income compared to non-metropolitan native-born citizens, holding all else constant. Likewise, a metropolitan naturalized citizen at the same level of hours is associated with a 21.29% higher income relative to a metropolitan native-born citizen, holding all else constant. These values are reasonably close to the coefficient obtained in the previous model for the naturalization premium without interaction terms.

All coefficients in the final model are statistically significant at the 1% level ($p < 0.01$). The model's adjusted R-squared of 0.280 indicates moderate explanatory power for a cross-sectional dataset. Furthermore, this value of R-squared is the highest across all the models, indicating that this model explains the most variation in log income. The confidence intervals provide additional insight into the precision of the estimated coefficients in the final model. For the naturalized coefficient, the 95% confidence interval ranges from 0.543 to 0.695, indicating that the baseline naturalization premium remains positive for all plausible values. The 95% confidence interval for the *naturalized_uhrsworkt* coefficient ranges from -0.00956 to -0.00715 and does not include zero. This suggests that the lower marginal return to an additional hour worked for naturalized citizens is a statistically precise estimate. Finally, the 95% confidence interval for the coefficient of *naturalized_metrop* ranges from -0.127 to -0.0168. This strictly negative range demonstrates a statistically significant reduction in naturalization premium in metropolitan areas. Together, these intervals show the key coefficients of interest are precisely estimated and robust to sampling variation.

VI. Conclusion

While this analysis provides meaningful descriptive insight into income differences across citizenship types, it is important to acknowledge that the CPS ASEC data are cross-sectional and cannot capture unobserved characteristics that may influence both naturalization and earnings. Factors such as language ability, pre-migration experience, motivation, or selection into naturalization may play an important role, and these cannot be fully addressed with the available data. As a result, some of the patterns observed may reflect underlying selection effects rather than causal relationships. The interpretations offered here should therefore be understood as possible explanations rather than definitive mechanisms.

Looking at individuals who work 40 hours per week, naturalized citizens earn about 21.29% more than native-born Americans when they live in metropolitan areas, and roughly 28.5% more than native-born Americans when they live outside those areas, holding all else constant. These patterns highlight that the income premium associated with naturalization appears consistently across different geographic contexts, although the reasons behind these differences cannot be determined from the available data and should be interpreted as possible explanations rather than definitive causes. Moreover, the interaction terms offered speculative insights outside of the model that require further research. Naturalized citizens start at a higher baseline, but they earn less on each additional hour worked. This suggests that although legal status may provide a higher income, it may still struggle to unlock higher marginal productivity, potentially reflecting differences in access to jobs, licensing, or employer-sponsored training. This seems to suggest the presence of institutional frictions for naturalized and native-born citizens.

This paper also identified temporal cohort effects. Relative to after 2008, immigrants naturalizing before 2004 earn approximately 11.8% less, while those who naturalized during 2004–2008 earned about 4.37% less, despite the broader economy not recovering until 2011. This could reflect how barriers for naturalizations are slowly lifting and naturalized citizens are finding more success as the economic climate becomes more accommodative in the long run. The broader set of results also highlights persistent structural disparities across demographic groups. The estimated gender gap of 27.6%, racial differences in returns to income, and the income penalty associated with disability all align with long-standing evidence of unequal labor market outcomes in the United States. These findings show that even within the category of citizens, economic opportunity is shaped by a complex interaction of human capital characteristics, demographic factors, and institutional conditions.

The paper's conclusions can help guide naturalization policies, and US citizenship seekers. Politicians looking to reform the naturalization process may refer to the wage gap found in this paper to highlight that perhaps the US could benefit from being more open to immigration and naturalization. Naturalized citizens fare better than birthright citizens despite their weaker marginal effects, so an increase in naturalizations could reflect an improvement in the US economy. Furthermore, measures to remedy what may be causing weaker marginal effects could allow naturalized citizens to contribute much more strongly to the economy. A possible hypothesis for this behavior is that naturalized citizens earn more because of their achievement of legal citizenship. As the barriers to citizenship rise, those who have achieved naturalization may continue to invest that same rising effort and determination in their work. Additionally, naturalized citizens may earn more if they are being selectively recruited by domestic firms that may support their subsequent naturalization. Future work should be done to explain the income discrepancy between the two forms of citizens. Research should also be conducted to find if institutional barriers are responsible for the discrepancy in marginal effects. Doing so will offer even deeper insights for policy makers to target certain shortcomings in policy to ensure equitable earnings opportunities for both naturalized and birthright citizens.

References

Bratsberg, B., Ragan Jr, J. F., & Nasir, Z. M. (2002). The effect of naturalization on wage growth: A panel study of young male immigrants. *Journal of Labor Economics*, *20*(3), 568–597. https://doi.org/10.1086/339616

CDC. (2016, January). CDC Newsroom. https://www.cdc.gov/media/releases/2024/s0716-Adult-disability.html

Flood, S., King, M., Rodgers, R., Ruggles, S., & Warren, J. R. (2023). Integrated Public Use Microdata Series, Current Population Survey: Version 11.0 [dataset]. https://doi.org/10.18128/D030.V11.0

Hall, M., Greenman, E., & Farkas, G. (2010). Legal status and wage disparities for mexican immigrants. *Social Forces*, *89*(2), 491–513. https://doi.org/10.1353/sof.2010.0082

Leclerc, C. (2022). Immigrants' earnings and neighbourhood economic wealth: The conditioning role of citizenship. *Journal of Ethnic and Migration Studies*, *48*(7), 1591–1609. https://doi.org/10.1080/1369183X.2021.1971958

Marrow, B., & Batalova, J. (2024, July). Naturalized citizens in the United States. https://www.migrationpolicy.org/article/naturalization-trends-united-states

Shierholz, H. (2010, February). *The effects of citizenship on family income and poverty* (EPI Briefing Paper No. 256). Economic Policy Institute. https://files.epi.org/page/-/bp256/bp256.pdf

Steinhardt, M. F. (2008, November). *Does citizenship matter? The economic impact of naturalizations in Germany* (Development Studies Working Paper No. 266). Centro Studi Luca d'Agliano. https://doi.org/10.2139/ssrn.1295995

U.S. Department of Homeland Security. (2023). *2022 Yearbook of Immigration Statistics* (tech. rep.). U.S. Department of Homeland Security, Office of Homeland Security Statistics. Washington, D.C.

VII. Appendix

Table 1.2: Summary Statistics for Total Annual Income by Naturalization Status

Naturalization Status	Observations	Mean Income ($)	Std. Deviation	Minimum	Maximum
Native-born (0)	561,942	46,528.22	52,950.46	0	1,700,287
Naturalized (1)	44,264	48,890.73	56,961.98	0	1,199,998

Table 2: Descriptive Statistics

Variable	Observations	Mean	Std. Deviation	Minimum	Maximum
inctot_rev	606,206	46,700.73	53,257.10	0	1,700,287
naturalized	606,206	0.0730	0.2602	0	1
black	606,206	0.1090	0.3116	0	1
white	606,206	0.8527	0.3544	0	1
male	606,206	0.5031	0.5000	0	1
married	606,206	0.6118	0.4873	0	1
educ_rev	606,206	13.8337	2.6021	0	21
disabled	606,206	0.0196	0.1387	0	1
midwest	606,206	0.2517	0.4339	0	1
south	606,206	0.3049	0.4604	0	1
west	606,206	0.2370	0.4253	0	1
kids	606,206	0.5377	0.4986	0	1
metrop	606,206	0.7855	0.4105	0	1
uhrsworkt_rev	606,206	39.3614	12.3685	0	198
yrs_us	605,954	39.6126	13.7758	0	85
before_peak	606,206	0.3808	0.4856	0	1
during_peak	606,206	0.2545	0.4356	0	1
age_rev	606,206	41.1776	13.1029	15	85
yrimmig	606,206	145.0766	516.5697	0	2011
naturalize_t	606,206	2.9518	10.9229	0	160
naturalize_p	606,206	0.0688	0.2531	0	1

Table 3: Mean Summary Statistics according to Naturalized

Variable	**Native-born (0)**	**Naturalized (1)**	**Total Sample**
inctot_rev	46,528.22	48,890.73	46,700.73
black	0.1095	0.1026	0.1090
white	0.8751	0.5679	0.8527
male	0.5021	0.5158	0.5031
married	0.6034	0.7186	0.6118
educ_rev	13.8502	13.6201	13.8334
disabled	0.0201	0.0128	0.0196
midwest	0.2625	0.1149	0.2517
south	0.3083	0.2614	0.3049
west	0.2258	0.3798	0.2370
kids	0.5299	0.6366	0.5377
metrop	0.7731	0.9424	0.7855
uhrswor_kt_rev	39.2776	40.4260	39.3614
yrs_us	40.9034	23.1565	39.6126
before_peak	0.3841	0.3389	0.3808
during_peak	0.2547	0.2521	0.2545
age_rev	40.9065	44.6199	41.1776
yrimmig	0.2081	1984.216	145.0766
naturalize_t	0	40.4260	2.9518
naturalize_p	0	0.9424	0.0688

Table 4: Regression Outputs

VARIABLES	(1) Simple Log-Linear	(2) Multiple Log-Linear Model	(3) Multiple Log-Linear Model with Interactions
naturalized	0.106***	0.218***	0.619***
	(0.00791)	(0.0142)	(0.0389)
white		0.0753***	0.0752***
		(0.00586)	(0.00586)
asian		0.0504***	0.0512***
		(0.0115)	(0.0115)
male		0.276***	0.276***
		(0.00362)	(0.00362)
married		0.138***	0.138***
		(0.00427)	(0.00427)
educ_rev		0.138***	0.138***
		(0.000704)	(0.000704)
disabled		-0.153***	-0.152***
		(0.0127)	(0.0127)
midwest		-0.0514***	-0.0516***
		(0.00522)	(0.00522)
south		-0.0926***	-0.0931***
		(0.00506)	(0.00505)
west		-0.0344***	-0.0352***
		(0.00530)	(0.00530)
kids		0.236***	0.235***
		(0.00390)	(0.00390)
metrop		0.182***	0.184***
		(0.00435)	(0.00440)
uhrsworkt_rev		0.0345***	0.0350***
		(0.000151)	(0.000155)
yrs_us		0.0141***	0.0139***
		(0.000576)	(0.000576)
before_peak		-0.118***	-0.118***
		(0.00406)	(0.00406)
during_peak		-0.0439***	-0.0437***
		(0.00452)	(0.00451)
age_rev		0.0129***	0.0131***
		(0.000574)	(0.000574)
naturalized_uhrsworkt			-0.00835***
			(0.000614)
naturalized_metrop			-0.0721**
			(0.0282)
Constant	10.19***	5.377***	5.359***
	(0.00214)	(0.0135)	(0.0135)
Observations	606,206	605,954	605,954
R-squared	0.000	0.279	0.280

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

Exploring the Effects of Access to Electricity on GDP Per Capita

By Ali el Sayed

This paper examines the impact of electricity access on GDP per capita using panel data from 217 countries between 1998 and 2023. A fixed-effects log-linear regression finds that a one percentage point increase in access to electricity is associated with a 1.2% rise in GDP per capita. Model testing confirms the relationship is quadratic, with both linear and squared terms significant at the 1% level. The final model, which controls for rural population, institutional stability, educational attainment and foreign direct investment, explains 77% of the variation in GDP per capita. These findings highlight the economic importance of access to electricity and the role of stable institutions in supporting development.

I. Introduction

Electricity access plays a crucial role in shaping a nation's economic growth. It improves productivity, enables the delivery of education and health services, and increases industrial and technological development. However, most developing nations still suffer from infrastructure shortages that impede inclusive growth. Thus, this article examines the impact of access to electricity on GDP per capita using a global panel data set of 217 countries from 1998 to 2023. The main objective is to assess whether higher electricity access is associated with higher GDP per capita, while controlling for various socioeconomic and institutional factors. Previous research has confirmed a strong link between access to electricity and growth, although causality has been hard to establish. Stern, Burke, and Bruns (2019) observe a strong positive correlation between electricity consumption and per capita GDP, particularly in low-income countries, but note the limitations of causality. David I. Stern is an environmental and energy economist at the Australian National University (ANU), Paul J. Burke is also an economist at ANU, and Stephan B. Bruns is an economist whose work overlaps with energy economics and econometric methods, meta-analysis and empirical research.

Others, like Andersen and Dalgaard (2013), note that unreliable electricity, especially in the case of countries like those in Sub-Saharan Africa, can negatively affect growth and argue that both access to and reliability of electricity are central to this. Micro-level data, such as Lee, Miguel, and Wolfram (2016), show that electrification of rural Kenyan homes is associated with income gains, while Kanagawa and Nakata (2008) correlate electricity supply with increased literacy rates, one of the main catalysts of sustainable development. Therefore, the evidence points to the economic significance of electrification, but more work should be done so that its macroeconomic effects on a broad cross-regional global population are well understood. Using fixed-effects panel regression, we regress the link between electricity availability (in terms of percentage population with access) and GDP per capita, controlling for rural population ratio to total population, political stability, foreign direct investment, and share of the labor force with intermediate education (graduated with a high school diploma). We observe a robust and statistically significant link between access to electricity and GDP per capita. These results affirm the economic usefulness of investment in infrastructure, particularly in access to energy, as a vehicle towards development.

Building on this background, this paper investigates whether higher access to electricity is associated with higher GDP per capita across 217 countries from 1998 to 2023. By applying fixed-effects panel regression and controlling for key institutional, demographic, and educational factors, I aim to clarify how electrification contributes to income growth and whether its impact strengthens as access expands.

II. Literature Review

Prior studies indicate a strong relationship between electricity access and economic growth, though with weak evidence of proving causality. Stern, Burke, and Bruns (2019) confirm that electricity use, and per capita GDP are positively related, especially in less developed countries, but few have shown causality. Andersen and Dalgaard (2013) also confirm that frequent power blackouts significantly lower economic growth rates in regions like Sub-Saharan Africa and suggest that reliability is as crucial as access to electricity. Moreover, Calderón, Moral-Benito, and Servén (2015) conclude that investment in electricity infrastructure positively and significantly contributes to GDP growth for countries, but it is challenging to identify electricity's impacts. Lee, Miguel, and Wolfram (2016) micro-level evidence from households in small regions in Kenya also shows the economic benefit of electrification, which includes higher incomes to Kenya's rural households following electrification. In general, while there is strong evidence of a positive relationship between access to electricity and economic development, the study concluded that further rigorous causal analysis is needed.

Jamil and Ahmad (2010) analyze the relationship between electricity consumption, electricity prices, and real GDP of Pakistan with respect to annual data within 1960-2008. They find unidirectional causality from economic growth to electricity consumption, which means that an increase in GDP causes an increase in electricity demand rather than vice versa. The research studies multiple sectors and most exhibit the same pattern as in commercial and domestic markets, while the agricultural market determines bi-direction causality. They concluded that policies aiming to conserve electricity may not adversely limit economic growth, although investment is used to increase the capacity of electricity generation to fulfil increasing demands. Therefore, the findings show that electric planning must be accompanied by general economic development strategies (Jamil & Ahmad, 2010).

Kanagawa and Nakata (2008) examine the link between socio-economic development in rural regions of developing countries and access to electricity. Through their research, they find that there is a significant and positive relationship between household electrification and improved literacy levels, which are principal determinants of long-run economic growth. Using a multiple regression model, they show rises in electrification rates to be statistically associated with higher literacy achievement, and electrifying towns fully could dramatically increase educational levels. Their finding harmonizes with wider evidence demonstrating that electricity access indirectly contributes to GDP growth by enhancing productivity and human capital. Even though the study focuses on education rather than GDP per capita directly, the authors demonstrate that electrification raises key indicators of development strongly related to higher levels of national income. This supports my research question with evidence at a micro level, where extending the availability of electricity can generate economic growth through human capital development (Kanagawa & Nakata, 2008).

Overall, previous research shows that electricity access is linked to economic development, but it also leaves important questions open. Large-scale studies (like Stern, Burke, & Bruns, 2019; Calderón, Moral-Benito, & Servén, 2015) find positive relationships but often struggle to prove causality or cover differences across countries. Other work looks at reliability (Andersen & Dalgaard, 2013) or focuses on smaller, local settings (Lee, Miguel, & Wolfram, 2016; Kanagawa & Nakata, 2008), while case studies such as Jamil & Ahmad (2010) suggest that the direction of the relationship can change depending on the context. This paper builds on these insights by using a large global dataset (217 countries, 1998–2023) and applying fixed-effects and log-linear models to reduce bias and capture both short- and long-term effects. This paper also tests for non-linear patterns, something earlier work rarely explored. In this way, this study adds to the literature by providing a broad, global picture of how electricity access connects to GDP per capita, while addressing gaps such as the lack of cross-country coverage, the absence of institutional and demographic controls, the limited attention to how the strength of the relationship changes as access increases, and the failure to account for how political stability or violence may shape the economic benefits of electrification.

III. Data Description

This research paper employs panel data collected at the country level between 1998 to 2023 from 217 countries. The data is sourced from the World Bank's World Development Indicators. The dependent variable is GDPpercap, while the key independent variable is *AccessElec*. To account for other determinants of AccessElec, the model also includes *RuralPop*, *LaborFInter*, PSA, and FDI. Summary statistics show a wide range between countries and years, with GDPpercap ranging from $96.32 to 256{,}581$, and AccessElec ranging from 0.8% to 100%. There are 5,431 observations for GDPpercap and 5,258 for AccessElec in the data, with all variables cleaned and normalized for comparability. Some countries report full electrification at 100%, while others remain near zero, showing stark global inequality. Likewise, the wide spread in the standard deviation of $z2.262x1010 in FDI highlights how unevenly financial resources for development are distributed across nations. This is a solid foundation for examining how increasing electricity access is linked to income levels after accounting for demographic, institutional, educational, and financial factors.

Table 1: Variable Descriptions

Variable	Description
Time	Year
CountryName	Country Name (duplicate in original list)
GDPpercap	GDP per capita using constant 2015 value of the dollar
AccessElec	Access to electricity (% of population)
LaborFInter	Labor force with intermediate education (% of total working population with high school diploma)
PSA	Political Stability and Absence of Violence/Terrorism: percentile rank
FDI	Foreign direct investment, constant 2015 value of the dollar
RuralPop	Rural population (% of total population)

Table 2: Summary Statistics

VARIABLES	N	Mean	SD	Min	Max
Time	5,642	2,010	7.501	1,998	2,023
GDPpercap	5,431	15,349	24,129	96.32	256,581
AccessElec	5,258	81.05	29.25	0.800	100
RuralPop	5,590	41.39	24.27	0	92.17
LaborFInter	2,160	65.61	9.781	21.67	100
PSA	4,839	49.25	29.01	0	100
FDI	4,475	4.235×10^{8}	2.262×10^{10}	3.454×10^{11}	2.183×10^{11}

IV. Empirical Model

We first estimate a linear regression model to examine *GDPpercap* as a function of *AccessElec* using country-level panel data with time and country fixed effects. In this fixed effects model, holding time *(t)* and country *(i)* constant, *GDPpercap* represents the GDP per capita in country *(i)* at time *(t)*, and is the dependent continuous variable of the equation. The main independent variable, *AccessElec*, measures the percentage of the population with access to electricity in country *(i)* and year *(t)*. This regression model is estimated to be linear and the coefficient on AccessElec is expected to be positive as more electricity access is correlated with improved productivity and economic growth.

$$\log(GDPpercap_{it}) = \beta_0 + \beta_1 AccessElec_{it} + \beta_2 RuralPop_{it} + \beta_3 LaborFInter_{it} + \beta_4 PSA_{it} + \beta_5 FDI_{it} + \alpha_i + \delta_t + \epsilon_{it}$$

This panel data analysis spans the years 1998 to 2023 across 217 countries. Country fixed effects (α) control for unobservable time-invariant factors such as geography and baseline governance, while time fixed effects (δ) account for global events or shocks. The error term (ϵ) captures all other unexplained variations such as differences in infrastructure quality, short-term economic shocks, measurement error in national statistics, or policy changes that are not included in the model. The model is estimated using fixed effects regression, and robust standard errors are clustered at the country level to account for heteroskedasticity and omitted variable bias. Heteroskedasticity means the error term does not have constant variance, which can make standard errors unreliable. Omitted variable bias occurs when an important factor is left out of the model, causing biased estimates of the included variables. The final model is designed to isolate the within-country variation in *AccessElec* and its effect on *GDPpercap* over time.

Figure 1 shows a positive relationship between *log(GDPpercap)* and *AccessElec*, with the fitted curve indicating a quadratic trend. To test this, *log(GDPpercap)* was regressed on *AccessElec* and *AccessElec*^{2}, both of which were significant at the 1% level. This confirms the relationship is better modeled as quadratic, suggesting stronger effects of electrification on *GDPpercap* at higher levels of access.

Figure 1: Personal Longevity Expectation by Financial Wellbeing Category

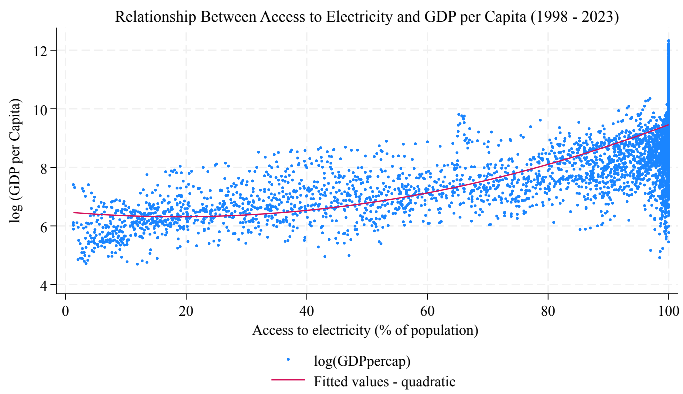

Figure 2 displays *AccessElec* over time from 1998 to 2023, with data points colored by year. This plot was used to assess whether electricity access has increased consistently across countries. While some clustering appears at higher access levels in later years, the spread remains wide across all years, indicating that a general upward trend over time cannot be clearly assumed.

Figure 2: Differing Associations of Financial Wellbeing on Personal Longevity Expectation

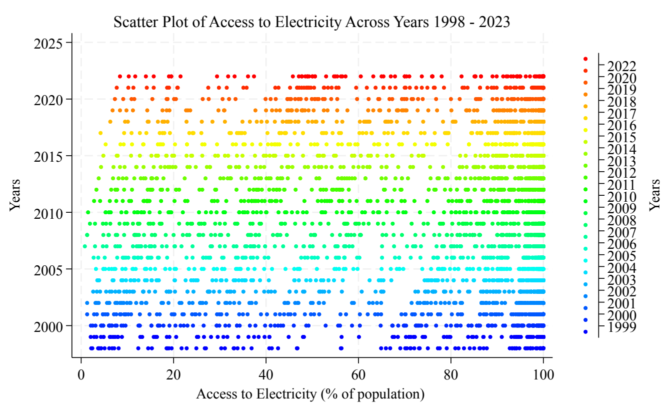

Figure 3 tracks *AccessElec* from 1998 to 2023 for six countries. The United States serves as a developed country benchmark, while Brazil and India represent developing nations in South America and Asia respectively. Three Sub-Saharan African countries: Nigeria, Angola, and Uganda are included due to their historically low access to electricity rates as seen in the studies from research from Lee, Miguel, and Wolfram (2016). The figure highlights persistent gaps, with Sub-Saharan countries showing slower and more uneven progress compared to the other selected countries.

Figure 3: Scatter Plot of Selected Countries Access to Electricity (1998-2023)

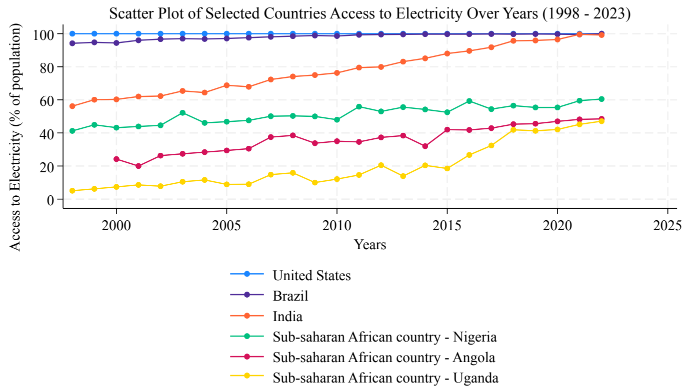

V. Results

To determine the appropriate model specification, I followed the 10-step diagnostic process to test whether the dependent variable, *GDPpercap*, or the independent variable, *AccessElec*, should be logged or both. By comparing adjusted R-squared values and evaluating residual patterns, I found that logging the dependent variable improved model fit and interpretability. Therefore, I used *log(GDPpercap)* as the outcome variable in the main regression specifications of model 2 to 5.

Models 3 to 5 presents a series of five regressions estimating the relationship between *AccessElec* and *GDPpercap*, progressively adding controls and fixed effects. Model (1) uses a linear regression of *GDPpercap* on *AccessElec*, showing a strong and significant positive relationship. However, this model has a very low adjusted R-squared, indicating limited causation. In model (2), *GDPpercap* is logged, resulting in a substantially higher adjusted Rsquared (0.243), and *AccessElec* remains highly significant at the 1% level, with a coefficient of 0.0292. This suggests that a 1 percentage point increase in electricity access is associated with approximately a 2.9% increase in GDP per capita.

In model (3), certain controls *(RuralPop, LaborFInter, PSA and FDI)* are introduced. The coefficient on *AccessElec* remains positive and statistically significant at 1%, though slightly reduced (0.0186). *RuralPop* has a significant negative effect, implying that more rural populations are associated with lower income levels. PSA significant at the 5% level. *LaborFInter, FDI* are not significant in this model. In model (4), I introduce country fixed effects. The coefficient on *AccessElec* remains positive (0.0152) and significant at the 1% level, while *RuralPop* becomes more negative and significant, and *LaborFInter* and *FDI* remain insignificant.

Finally, model (5) includes both country and year fixed effects. This specification yields the highest adjusted R-squared (0.769), indicating the strongest model fit. The coefficient on *AccessElec* is 0.0119 and remains significant at the 1% level, reinforcing the main hypothesis that electricity access is positively associated with GDP per capita. In this fully specified model, *PSA* becomes significant and positive, while *RuralPop* and LaborFInter lose significance, suggesting that once trends are accounted for, institutional stability plays a more prominent role. Additionally, *FDI* is excluded from the regressions due to insufficient variation.

Across all models, *AccessElec* consistently shows a positive and significant association with *log(GDPpercap)*, supporting the idea that expanding access to electricity contributes to economic development. The results also highlight the importance of controlling for both cross-country differences and temporal shocks, as model fit improves substantially with the inclusion of fixed effects.

Table 3: Regression Results

VARIABLES	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**
Outcome Variable	**GDPpercap**	**log(GDPpercap)**	**log(GDPpercap)**	**log(GDPpercap)**	**log(GDPpercap)**
AccessElec	72.48***	0.0292***	0.0186***	0.0152***	0.0119***
	(8.547)	(0.00207)	(0.00204)	(0.00237)	(0.00187)
RuralPop			-0.0422***	-0.0683***	0.00562
			(0.00312)	(0.00679)	(0.00887)
LaborFInter			-0.00299	-0.00444	0.00201
			(0.00275)	(0.00294)	(0.00185)
PSA			0.00375**	0.00132	0.00579***
			(0.00148)	(0.00160)	(0.00111)
FDI			-0	-0	-0
			(0)	(0)	(0)
Constant	9,526***	6.201***	8.846***	10.33***	6.534***
	(1,423)	(1.80)	(0.332)	(0.398)	(0.450)
Observations	5,103	5,103	1,907	1,907	1,907
R-squared				0.356	0.773
Adjusted R-squared	0.00249	0.243	0.354	0.354	0.769
Country FE	NO	NO	NO	YES	YES
Year FE	NO	NO	NO	NO	YES
Number of Countries	212	212	173	173	173

VI. Conclusion

This paper examines the impact of electricity access on GDP per capita across 217 countries from 1998 to 2023, addressing the research question of whether electrification supports economic growth. The hypothesized equation holds to be partially true, as *AccessElec* is consistently positive and significant in explaining *log(GDPpercap)*. Using a fixed effects log-linear model, the analysis shows that a one percentage point increase in electricity access leads to approximately a 1.2% increase in GDP per capita in model (5) in Table 3. Further testing showed that the relationship is not purely linear: including both *AccessElec* and $AccessElec^2$ confirms a quadratic relationship, with both terms significant at the 1% level. This suggests the economic benefits of electrification grow stronger at higher access to electricity levels.

These findings highlight the developmental value of electricity access, particularly in stable environments. Among control variables, *PSA* becomes significant in the full model (coefficient 0.00579), pointing to the importance of institutional quality and safety. *RuralPop* is significant in earlier models but drops out with fixed effects. *LaborFInter* and *FDI* remain mostly insignificant but were kept in the model to avoid omitted variable bias, as they may still account for underlying factors influencing GDP per capita. The final model achieves a high adjusted R-squared (0.769), showing strong fit and reinforcing the role of electrification as a driver of economic development. However, the study has limitations. Fixed effects help control for unobserved country and time factors, but reverse causality remains a concern because wealthier countries are better positioned to invest in electrification. Several controls, including *FDI* and *LaborFInter*, were insignificant or had to be dropped due to missing data, which limits the model's ability to capture all relevant drivers of GDP per capita.

Country-level averages may also hide important within-country differences, especially in countries with large urban–rural divides.

Future work should use causal methods such as instrumental variables to better isolate the effect of electrification. Subnational data and interaction terms could clarify how electricity access interacts with education, governance, and geographic inequality. Including longer timeframes and lagged variables would help capture long-run effects, since infrastructure benefits often emerge slowly. Finally, future research should distinguish between access and reliability because frequent outages can limit the economic gains even where official access is high. This paper contributes to the literature by showing that electricity access has a strong and statistically significant association with GDP per capita, and that its benefits become larger as access expands. While challenges of causality remain, the results underline the critical importance of expanding and stabilizing electrification as part of broader development policy.

References

Stern, D. I., Burke, P. J., & Bruns, S. B. (2019). The Impact of Electricity on Economic Development: A Macroeconomic Perspective. EEG State-of-Knowledge Paper Series.

Andersen, T. B., & Dalgaard, C.-J. (2013). Power outages and economic growth in Africa. Energy Economics, 38, 19–23.

Calderón, C., Moral-Benito, E., & Servén, L. (2015). Is infrastructure capital productive? A dynamic heterogeneous approach. Journal of Applied Econometrics, 30(2), 177–198.

Lee, K., Miguel, E., & Wolfram, C. (2016). Experimental evidence on the demand for and costs of rural electrification. Journal of Political Economy, 124(6), 1768–1812.

Jamil, F., & Ahmad, E. (2010). The relationship between electricity consumption, electricity prices and GDP in Pakistan. Energy Policy, 38(10), 6016–6025.

Kanagawa, M., & Nakata, T. (2008). Assessment of access to electricity and the socioeconomic impacts in rural areas of developing countries. Energy Policy, 36(6), 2016–2029.

Financial Wellbeing and Its Association with Longevity Optimism

By Anthony Benavides

This paper assesses whether financial wellbeing defined as of one's overall financial stability and freedom, as measured by Financial Wellbeing Score (CFPB), is positively associated with longevity optimism. Longevity optimism is defined as a dispositional optimism regarding an expectation of increased longevity, measured by an individual's percentage expectation of living beyond the age of 75. Using an ordinary least squares regression with an interaction term for health insurance possession and controls for various health and demographic related variables to isolate the main relationship, a model was constructed that accounted for 21.3% of the variance in individual longevity expectations and found a positive linear correlation between the dependent and independent variables with a main coefficient of 0.219 significant at the p<0.01 level, suggesting that an individual's financial security is a significant predictor of their longevity outlook.

I. Introduction

This paper aims to discover whether one's financial wellbeing is correlated with their longevity optimism. Optimism, when defined as "a psychological attribute characterized as the general expectation that good things will happen," has been shown to be correlated with actual longevity (Lee et al, 2019). Opening the door for further research into how financial standing can affect one's optimism and thus expected longevity, although it is important to note that this is a broader form of optimism than is analyzed in this paper. The correlation between financial wellbeing (as measured through income) and life expectancy has been studied and found to be statistically significant, having a P value $< .001$ (Chetty et al, 2016).This leads to the question whether one's perception of their own financial wellbeing follows a similar pattern. Furthermore, controlling for health and demographic factors isolates this relationship to determine if it holds true independent of a higher prevalence of illness at lower incomes. In this paper, longevity optimism will be measured by one's personal expectation of living beyond the age of 75 (lifeexpect), while financial wellbeing will be measured by their Financial Wellbeing Score (fwbscore) (as outlined in the Data Description section).

In the following sections, a hypothesized ordinary least squares population regression–an empirical model is created–that expects a positive relationship between financial wellbeing and one's personal expectation of living beyond the age of 75 based on current research and known correlations between important variables. This relationship was depicted through graphs and regression testing, culminating in a final model that explains 21.3% of the variation in one's personal expectation of living beyond the age of 75. Ultimately, a positive linear relationship dependent on the possession of health insurance was found as outlined in the empirical model, with a one unit increase in one's Financial Wellbeing Score being associated with a 0.080 and 0.219 percentage point increase in their personal expectation of living beyond the age of 75 for the insured and uninsured, respectively, with a β_1 that was significant at the one percent level. This finding suggests the importance of considering financial wellbeing beyond its application in economics, encouraging further research into how one's financial position and possession of health insurance contribute to perceptions of longevity, and how those perceptions contribute to actual longevity. In addition, this importance beyond economics could apply to other non-longevity related factors such as mental health, long-term decision makings such as retirement savings, and investments in one's health. This leads to various opportunities for policy research and expansion, such as the secondary benefits of health insurance expansion.

II. Literature Review

Recent research suggesting a correlation between optimism and physical health provides an encouraging motive for studying the correlation between financial wellbeing and optimism. In fact, in the paper *Optimism is associated with exceptional longevity in 2 epidemiologic cohorts of men and women*, Lee et al (2019) found an "association of higher optimism levels at baseline with increased longevity" at the 1% significance level. This research, split into a male and female cohort and defining "increased longevity" as living beyond the age of 85, found an average 1.5x higher rate of longevity among women and 1.7x higher rate of longevity among men for those with the highest versus lowest optimism levels when controlling for demographics and health conditions.

Chetty et al's 2016 paper, *The Association Between Income and Life Expectancy in the United States*, is an example of an established connection between financial wellbeing and actual health outcomes. Here, the researchers used over 1.4 billion human-hours of data collected from 1999-2014 using deidentified tax and Social Security mortality records that included data on income and mortality rates at different ages. Household income was compared with life expectancy, reaching the conclusion that "higher income was associated with greater longevity." It was also found that differences in life expectancy correlated to behaviors regarding health and characteristics of local areas, raising the theory that differences in medical care are related to health and longevity. This shows that while income and life expectancy are correlated, there are many other factors at play. This raises a question: could the suspicion or knowledge of this discrepancy could be internalized, and would this affect optimism?

Ryu, S., and Fan, L's 2023 paper *The Relationship Between Financial Worries and Psychological Distress Among U.S. Adults* is included because stress was expected to be included as a control variable, and this paper measured its correlation with financial wellbeing. This paper found a statistically significant positive relationship between financial worries and psychological distress, implying that this correlation does exist. This is important since stress and optimism are similar as mental health indicators, meaning stress could have a significant confounding effect and could cause omitted variable bias if not accounted for.

Finally, Arrow, K.J.'s classic paper *Uncertainty and the welfare economics of medical care* (1963), a foundational piece in healthcare economics, is included because it discusses the willingness for individuals to pool their illness risk, demonstrating the "safety net" that health insurance can play and offering reasons for why it exists. This paper focuses on uncertainty as the key driver of the healthcare industry, differing it from a normal competitive market. It is important to understand the role this plays in one's decision making and outlook regarding healthcare, especially in relation to health insurance and the effects this pooling of risk have on actions.

III. Data Description

In this paper, all the data was obtained from a 2017 Consumer Financial Protection Bureau survey focused on gathering Financial Wellbeing Scores at the individual level fielded on the GfK KnowledgePanel, a recruited, nationally representative sample. This dataset is cross-sectional with a sample size (n) of 6,394 people, collected to be representative of all 50 states and the District of Columbia with the aim to obtain the "first-ever picture of the current state of financial well being of American adults overall." (CFPB)

The Financial Wellbeing Score is a standardized measure from 1-100 meant to quantify "the extent to which someone's financial situation and the financial capability that they have developed provide them with security and freedom of choice" (CFPB). This score provides a more robust measure of financial stability than any single raw variable could, since it is formulated using multiple surveyed sub-questions aiming to get a full picture of one's financial situation. The questions measured relate to consumer definitions of what financial wellbeing means to them, including control over finances, the ability to absorb shocks, and the feeling of financial freedom. This extensive array of factors are weighed in one variable, FWBscore. The weighted questions in this score vary from "I am securing my financial future" to agreement or disagreement with "Giving a gift for a wedding, birthday or other occasion would put a strain on my finances for the month," with all responses quantified in a numerical range (CFPB, 2017). The breadth of questions included speak to the wide variety of perceptions and financial situations represented in a single variable.

Within the data, there are multiple different categories of variables. The categories relevant to this paper include IRT score, survey items, and panel data. IRT scores are variables created with Item Response Theory methods that are used to measure unobservable traits by their observable outcomes, done by fitting a model linking underlying traits to the probability of one's response to each item. Survey items come directly from the survey, where respondents are usually asked to rate their agreement to a statement on a provided scale. Panel data does not indicate a panel dataset but instead represents data drawn from the GfK KnowledgePanel members sampled in this survey, including federal poverty level and certain demographic variables such as ethnicity and household size as a part of standard procedure.

The nature of the data used meant that, while there were a few outliers as defined by points further from the third quartile (Q3) than 1.5 times the range between the first and third quartiles (1.5IQR), none were implausible, so none were removed. In the end, the only data removed were those with negative values, usually serving as a placeholder for a survey non-response. Three created variables will be discussed: *"fwb_healthinsured," "fwbscore_sq,"* and *"fwb2_healthinsured."* The first is an interaction term between Financial Wellbeing Score (*fwbscore*) and possession of health insurance (*healthinsured,* a dummy variable), the second is a squared version of *fwbscore*, and the third is an interaction term between *fwbscore_sq* and *healthinsured*. These three will be discussed in depth in the Empirical Model and Results sections of this paper.

Table 1: Variables Descriptions

Variable Name	Variable Description	Variable Type
lifeexpect	One's personal expectation of living past the age of 75 (%)	Survey
fwbscore	Financial Wellbeing Score (14–95)	IRT
fwbscore_sq	Squared Financial Wellbeing Score	Created
healthinsured	Possession of health insurance (0, 1)	Survey
fwb_healthinsured	fwbscore and healthinsured interaction term	Created
fwb2_healthinsured	fwbscore sq and healthinsured interaction term	Created
health	Personal rating of health (1–5)	Survey
distress	Personal rating of stress (1–5)	Survey
hincome	Household income (1–9, < \$20,000 to > \$150,000)	Panel
gender	Gender (1, 2; M, F)	Panel
generation	Generation (1–4, pre-boomer to millennial)	Panel
ppeduc	Education (1–5, < high school to graduate degree)	Panel
gen_optimism	Optimism about future, 1–7	Survey

Table 2: Summary Statistics

Variable	Mean	Std. Dev.	Min	Max
lifeexpect	73.36	29.13	0	100
fwbscore	56.08	14.07	14	95
healthinsured	0.71	0.45	0	1
health	3.45	0.92	1	5
distress	3.15	1.09	1	5
hincome	5.51	2.67	1	9
gender	1.48	0.50	1	2
generation	2.55	1.05	1	4
ppeduc	3.16	1.18	1	5
gen_optimism	5.42	1.43	1	7

Observations: 6,394

IV. Empirical Model

The following ordinary least squares hypothetical population model was constructed to estimate the effects of one's financial wellbeing score on their personal expectation of living beyond the age of 75:

Model 1

$$lifeexpect_i = \beta_0 + \beta_1 fwbscore_i + \beta_2 prodhave\_3_i + \beta_3 fwbscore{*}prodhave\_3_i + \beta_4 health_i + \beta_5 ppgender_i + \beta_6 generation_i + \beta_7 ppincimp_i + u_i$$

In this model, *lifeexpect* serves as the main dependent variable. It represents the surveyed response to the question "How likely do you believe it is that you will live beyond age 75?" for individuals i. The primary explanatory variable is *fwbscore*, or an individual's financial wellbeing score. The second explanatory variable is *healthinsured*, a dummy variable representing possession of health insurance (=1). The third explanatory variable is an interaction term between the first two, to measure whether the effect of financial wellbeing on personal life expectancy is dependent on the possession of health insurance. The fourth through seventh explanatory variables are controls for health and demographics. *health* is a surveyed rating of one's own health, *gender* is one's gender (M or F), generation is the generation of birth, and *hincome* is household income.

There is predicted to be a positive linear relationship between *fwbscore* and *lifeexpect*, or a positive β_1. This is due to the paper *The Association Between Income and Life Expectancy in the United States* measuring the correlation between life expectancy and income, a pattern which may alter the expectations of those surveyed in this direction (Chetty et al, 2016). The relationship between *healthinsured* (possession of health insurance) and both *fwbscore* and *lifexpect* is also predicted to be positive. First, it is predicted to have a positive relationship with *lifeexpect* because health insurance leads to easier access for preventative and emergency lifesaving care, and that is likely to be understood by the survey participants. Second, it is predicted to have a positive relationship with *fwbscore* because of the construction of *fwbscore*, which weighs the "capacity to absorb a financial shock" such as a health emergency in its creation (CFPB, 2015). The interaction term is predicted to be significant, because the ability to rely on health insurance could mean a weaker impact of financial wellbeing on personal life expectancy due to possession of a "safety net" of sorts; therefore, β_3 is expected to be negative. Since the healthcare industry is heavily driven on uncertainty (Arrow 1963), this safety net emerges from a willingness to pool risk and alleviate the burden of a medical emergency. Health is predicted to have a positive effect on *lifeexpect* and *fwbscore*. It is predicted to have this relationship with *lifeexpect* because a better self-perception of health likely ties to a better self-perception of longevity, and it is predicted to have this relationship with *fwbscore* for similar reasons as *healthinsured* (possession of health insurance) along with the positive correlation between income and life expectancy found in Chetty et al's paper (2016) that studied how household income affected longevity and found a correlation between the two. *gender* and *generation* are included to control for demographics. Finally, *hincome* (household income) is predicted to have a positive relationship with both *lifeexpect* and *fwbscore*. This is because of the Chetty et al (2016) study finding life expectancy and income to be correlated, and because income is considered in the construction of financial wellbeing score.

Figure 1: Personal Longevity Expectation by Financial Wellbeing Category

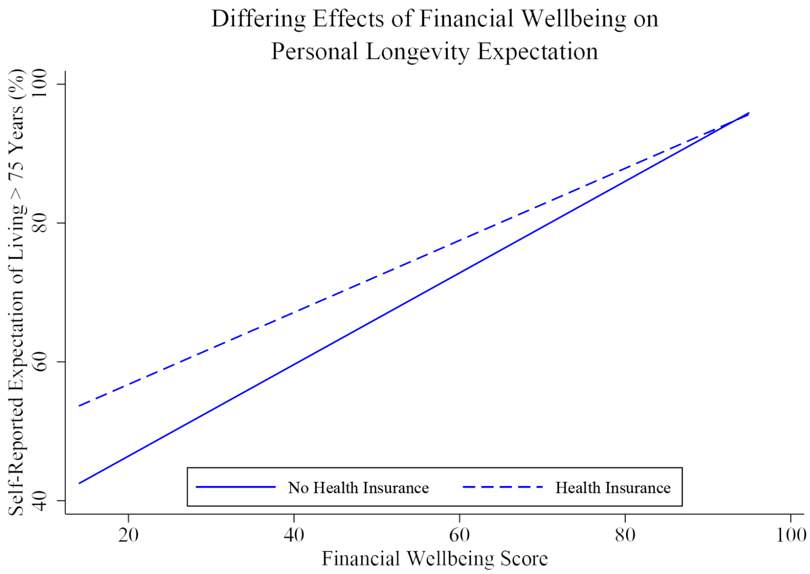

Figure 2: Differing Associations of Financial Wellbeing on Personal Longevity Expectation

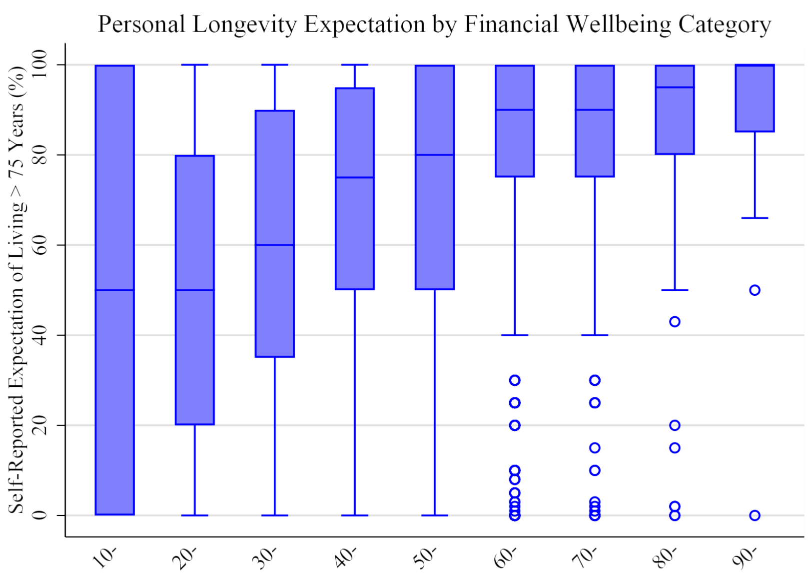

V. Results

Table 3: Regression Results

lifeexpect	(1) Basic Linear	(2) Interaction I (Emperical)	(3) Quadratic	(4) Quadratic Interaction	(5) Interaction II (Final)
fwbscore	0.602***	0.329***	0.556***	0.562**	0.219***
	(0.0278)	(0.0544)	(0.153)	(0.261)	(0.0545)
fwbscore_sq			-0.00280**	-0.00225	
			(0.00126)	(0.00233)	
healthinsured		8.966***	2.761***	7.171	9.605***
		(3.467)	(0.886)	(9.082)	(3.397)
fwb_healthinsured		-0.117*		-0.0739	-0.139**
		(0.0606)		(0.322)	(0.0597)
fwb2_healthinsured				-0.000180	
				(0.00280)	
health		9.795***	9.775***	9.765***	8.402***
		(0.455)	(0.455)	(0.456)	(0.469)
gender		4.282***	4.284***	4.284***	4.910***
		(0.702)	(0.702)	(0.702)	(0.710)
generation		-4.601***	-4.635***	-4.636***	-4.941***
		(0.372)	(0.372)	(0.372)	(0.382)
hincome		0.657***	0.645***	0.645***	
		(0.159)	(0.159)	(0.159)	
distress					-1.158***
					(0.384)
ppeduc					2.972***
					(0.324)
gen_optimism					2.451***
					(0.322)
Constant	40.05***	22.03***	18.43***	16.66**	18.00***
	(1.686)	(3.502)	(4.740)	(7.235)	(3.908)
Observations	5,632	5,609	5,609	5,609	5,560
Adjusted R-squared	0.083	0.193	0.193	0.193	0.213

Robust standard errors in parentheses
**** p<0.01, ** p<0.05, * p<0.1*

The first model, Basic Linear, is a basic linear regression that shows a positive relationship between *fwbscore* (financial wellbeing score) and *lifeexpect* (one's personal expectation of living past the age of 75) that is significant at the 1% level. This signifies that for every 1-unit increase in financial wellbeing score, one's personal expectation of living past the age of 75 increases on average by 0.602 percentage points. The adjusted r squared, a measure variation accounted for by included variables that only increases when new variables are significant, is 0.083. This is the lowest of any model tested.

The second model, Interaction I (Empirical), adds an interaction term between financial wellbeing score and possession of health insurance, as well as several control variables such as *health*, *gender*, *generation* (for age) and household income (*hincome*). It is based on the expected population regression from the *Empirical Model* section. While the control variables have an overall positive and overstating effect on β_1, it remains significant at the 1% level. All the control variables are statistically significant at the 1% level. Furthermore, the significance of the interaction term implies that β_1 is the expected effect of *fwbscore* on *lifeexpect* only when *healthinsured* (possession of health insurance) = 0, or when the individual does not have health insurance. In fact, the expected change of life expectancy associated with a one unit change in *fwbscore* drops to (approximately) 0.212 when the individual possesses health insurance, since the coefficient of the interaction term is negative. The adjusted r squared of 0.193 suggested a large increase in the amount of variation explained by the model when compared to the first. Since adjusted r squared accounts for insignificant predictors and only increases when new predictors are significant, the benefit of including the interaction term and the various control variables is demonstrated.

The third model, Quadratic, was created to test for a possible quadratic relationship where the main independent variable of *fwbscore* was squared. This was done to see if a non-linear curve fit the data better, which could imply diminishing returns or a turning point where increased financial wellbeing leads to a less positive or negative relationship with longevity optimism. With *fwbscore_sq* being significant at the 5% level, this relationship is found, with a turning point at *fwbscore* = 98.58, an x-value that is 3 points higher than the highest *fwbscore* collected. This high turning point implies that there was no point at which increased financial wellbeing was associated with decreased longevity optimism, but it did imply the possibility of diminishing returns in this relationship. For now, the interaction term was removed, to be re-added later in the correct manner associated with quadratic models. At this point, a linear log model was tested; however, the adjusted r squared was nearly identical to the interaction term and quadratic models (both 0.193), so it was not included. This is because the interpretation of a logged index variable is less straightforward than its non-logged counterpart.

Model four, Quadratic Interaction, keeps the quadratic variable while reintroducing interaction terms to test a combination of the previous two models in a statistically appropriate manner where there is a second interaction term including the squared version of *fwbscore*. The inclusion of interaction terms for both the linear and quadratic versions of *fwbscore* leads to a drop in significance of *fwbscore*'s coefficient (β_1)to the 5% level and a complete loss of significance for both interaction terms, *healthinsured* (possession of health insurance), and *fwbscore_sq*. Because of this unusual result, joint f-tests were run on the two *fwbscore* variables and their interaction term counterparts to evaluate their significance when measured together rather than individually. This method determines the collective significance of grouped independent variables on a dependent variable, therefore deciding whether these groupings were worth including in the model. First, a test was applied on all four variables, finding them to be jointly significant. This indicated a high probability that at least one coefficient was not equal to zero, meaning at least one of the grouped variables had a statistically significant association with longevity optimism. Next, a test was applied to *fwbscore_sq* and its associated interaction term. Here, they were found to be jointly insignificant. This justified their removal from the final model, which kept the original interaction term. Therefore, this model was not selected as the final, further justified by the unchanged adjusted R squared.

For the fifth model, Interaction III (Final), several more control variables were added, removing them as necessary when they did not cause omitted variable bias in order to create a model that suffered from less omitted variable bias overall. These include controlling for ethnicity (removed), education (included), stress (could impact optimism and is possibly influenced by financial wellbeing) and optimism about the future. In the end, the most significant of these was *gen_optimism*, or the optimism variable measured by the question "I am optimistic about my future," which led *fwbscore*'s coefficient (β_1) to move by an entire standard error. This drastic change implies that there was significant omitted variable bias when general optimism was unaccounted for in previous models. Furthermore, several tests were run for omitted variable bias on previously included control variables, removing them if there was not found to be a >0.5 standard error change in β_1. Because of this, *hincome* s(household income) was removed. This was initially surprising, as it had caused significant omitted variable bias previously. The drop in its association with β_1 (and significance) occurred when education was first controlled, suggesting multicollinearity between the two that could possibly be explained by the frequently high cost of post-high school education that was measured by the education variable. Education's expected effect was more significant on β_1 and *hincome* (household income) dropped to insignificance, so education was kept in the final model while *hincome* (household income) was removed. With the lower β_1 in the final model, the interaction term has a more dramatic effect. Now, when an individual possesses health insurance, the expected change of *lifeexpect* associated with a one unit change in *fwbscore* drops to 0.08. However, when they do not possess health insurance, a one unit increase in *fwbscore* is associated with a much larger 0.219 percentage point increase in *lifeexpect*. This suggests that the possession of health insurance leads to a sense of a "safety net" in respondents, since even when actual health is held constant and respondents have similar Financial Wellbeing Scores, those without health insurance have a lower longevity expectation that is more strongly driven by financial wellbeing as opposed to those with health insurance (Model 5, Figure 2). This model had the highest adjusted r squared at 0.213 and being that it was a significant improvement over the initial linear (0.083) and quadratic (0.193) models, it was chosen as the final model, with a β_1 that was significant at the one percent level.

VI. Conclusion

In the end, the results of the regressions are consistent with the hypothesis outlined in the model section. When controlling for various health and demographic variables, there is a statistically significant linear correlation of one's financial wellbeing on their personal expectations of living past the age of 75 at the one percent level, dependent on the possession of health insurance. When an individual does not have health insurance, a one-unit increase in their Financial Wellbeing Score is associated with a 0.219 percentage point increase in their expectation of living past the age of 75; when they have health insurance, the expected increase drops to 0.080 percentage points.

These results suggest that those with poor financial wellbeing tend to expect a shorter lifespan. Since health-related factors were controlled for, the more pessimistic view on longevity expected to be held by those who are less financially secure can be interpreted independently of these factors. This pattern makes sense because of the strain that poor finances can have on perceived control of one's own life, which could decrease expectations of longevity. The importance of this relationship is highlighted by findings such as those by Lee et al (2019) that find an association between optimism and actual longevity. These findings encourage further research into the mental health impacts of financial wellbeing, as this could uncover more links between the two that could deepen the understanding of the issues they surface and how to aid those currently suffering.

The significance of these findings lies in their framing of financial wellbeing in greater scope of health (particularly mental) and wellbeing rather than just in economics. This leads to notable areas for policy exploration, such as in the possibility of secondary benefits coming from health insurance expansion in the form of an increase in longevity expectation resulting in alleviated psychological tolls.

The primary limitation in this paper stems from the inability to have a similar objectivity in the measure of health as is present in the measure of financial wellbeing. Health is likely the most crucial factor to be controlled for when trying to measure the association between financial wellbeing and longevity optimism, and the subjective nature of its incorporation in the CFPB survey obscures its true effect. One person's personal health rating of 1 could be another's 3, and there is no way to know how personal perspectives altered these patterns. This is reflected in other subjective control variables, such as *gen_optimism* (optimism) and *distress* (stress). Another limitation is the cross-sectional nature of the data. The ability to run a fixed effects model using panel data could substantially increase the power of the research, looking at changes within individuals rather than across them. Finally, the data used were from 2017, 8 years in the past at the time of writing. Since this time, changes in the composition of income, perceptions of health altered by the Covid-19 pandemic, and medical discoveries altering life expectancy could play a possible role in the modern understanding of longevity optimism.

The fact that the model found that an increase of 1 in one's financial wellbeing score was associated with an increase of 0.219 in one's percentage expectation of living beyond the age of 75 at the $p<0.01$ level provides a compelling basis for further research. In the future, a panel survey with more objective measures for key variables could provide a more interesting and applicable look at the relationships found in this paper, seeing if they apply to the modern environment.

Furthermore, an extension of research into the exact ways in which optimism, financial wellbeing, and health play into longevity could provide valuable insights that could positively impact future policy decisions. For example, a study in the style of Lee et al's *Optimism is associated with exceptional longevity in 2 epidemiologic cohorts of men and women* focusing specifically on longevity optimism defined as a "dispositional optimism regarding an expectation of increased longevity" and how it affects actual longevity could provide a more secure link between this paper's research and the welfare and longevity effects that result from financial stability. Finally, because of the stronger positive correlation between financial wellbeing and longevity optimism found for those not in possession of health insurance when compared to those in possession of health insurance (an expected 0.219 increase in lifeexpect as opposed to 0.080 increase when fwbscore increases by 1, respectively), further research into the physical and mental health benefits of health insurance possessions could be used to make policy arguments in favor of the expansion of healthcare access if relationships were to be found.

References

Arrow, K. J. (2004). Uncertainty and the welfare economics of medical care. 1963. *Bulletin of the World Health Organization*, 82(2), 141–149.

Chetty, R., Stepner, M., Abraham, S., Lin, S., Scuderi, B., Turner, N., Bergeron, A., & Cutler, D. (2016). The association between income and life expectancy in the United States, 2001–2014. *JAMA*, 315(16), 1750–1766. https://doi.org/10.1001/jama.2016.4226

Consumer Financial Protection Bureau. (2015). *Financial well-being: The goal of financial education*. Washington, D.C.: Consumer Financial Protection Bureau. https://www.consumerfinance.gov/data-research/research-reports/financial-well-being/

Consumer Financial Protection Bureau. (2017). *Financial well-being in America: Technical report*. https://www.consumerfinance.gov/data-research/research-reports/financial-well-being-technical-report/

Consumer Financial Protection Bureau. (2017). *Financial well-being survey data*. Washington, D.C.: Consumer Financial Protection Bureau. https://www.consumerfinance.gov/data-research/financial-well-being-survey-data/

Lee, L. O., James, P., Zevon, E. S., Kim, E. S., Trudel-Fitzgerald, C., Spiro, A., Grodstein, F., & Kubzansky, L. D. (2019). Optimism is associated with exceptional longevity in 2 epidemiologic cohorts of men and women. *Proceedings of the National Academy of Sciences of the United States of America*, 116(37), 18357–18362. https://doi.org/10.1073/pnas.1900712116

Ryu, S., & Fan, L. (2023). The relationship between financial worries and psychological distress among U.S. adults. *Journal of Family and Economic Issues*, 44(1), 16–33. https://doi.org/10.1007/s10834-022-09820-9

Union Membership and Income Behavior: Estimating the Wage Premium Effect

By Otavio Agostini and Emiliano Vielma

This paper explores whether the wage benefits of union membership are distributed evenly across different demographic subgroups and occupational groups in the U.S. It employs microdata from IPUMS CPS, spanning from 2021 to 2024. The study examines factors like race, gender, age, job sector, and region, and how they impact the relationship between union membership and income levels. To address potential omitted variable bias, the analysis includes 11 regression models, ranging from linear–linear to log–linear, controlling for a wide range of demographic and occupational variables. Initial results from simple regression models show that union members have a wage premium of $4,897$ *relative to non-union workers. After controlling for additional variables, the wage differential associated with union membership turns slightly negative* (-721.82)*, suggesting that the initial wage premium reflected demographic factors rather than union membership alone. The introduction of interaction terms presents interesting patterns. The results show that the union wage premium for Black individuals is* $7,559.66$ *larger than that of White individuals, suggesting that unionization potentially plays a role in reducing the racial wage gap. The paper's general findings imply that union wage benefits are not evenly distributed across demographic and occupational subgroups.*

I. Introduction

In recent years, several union strikes have taken place, which have been publicly documented and shared online, as was the case at Boston University with the Graduate Student strike one year ago. Undergraduate students directly observed this across campus, and felt the consequences through the absences of teaching assistants and instructors. Examples such as this have sparked a growing public interest in labor rights and workplace equity, drawing attention to wage disparities and labor market outcomes across union affiliation, gender, race, employment sector, and region. According to the U.S. Bureau of Labor Statistics (2023), unionized workers tend to earn higher median weekly wages than their non-unionized counterparts, a difference commonly referred to as the union wage premium. From an economic perspective, this occurs because unions allow workers to bargain collectively, giving them greater negotiating leverage that raises wages above the market rate. In doing so, unions can help narrow the wage gap for groups with less individual bargaining power, such as minority workers. However, this premium is not experienced uniformly across demographic groups. Research from the Economic Policy Institute (2022) highlights that union wage premiums are larger at the bottom of the income distribution – typically for African Americans, women, and people with lower levels of education.

Eleven regression models were employed to analyze the relationship between union membership and income, including simple and multiple linear-linear and log-linear models. The simple regression models provided a baseline estimate of the relationship between *union_rev* and *incwage*, while the multiple regression models introduced control variables like race, sex, age, region, employment sector, education, and weekly hours worked to reduce omitted variable bias. Additionally, multiple regression models, including interaction terms like *union_rev_black*, *union_rev_female*, *union_rev_south*, *union_rev_private_wkr*, and *union_rev_public_sector*, were utilized. These were used to examine whether union wage premiums varied across different demographic and occupational subgroups.

The findings highlight a wage penalty for Black individuals compared to their White counterparts, even after controlling for other variables. However, when testing the interaction term between *union_rev_black*, Black individuals appear to have a higher union wage premium compared to Whites. This suggests that unionization may help reduce the racial wage gap between these two groups. Additionally, the regression models showed that female individuals earned significantly less than males, indicating a persistent gender wage gap. However, the interaction term between union membership and being female was not statistically significant, suggesting that unionization does not offset the wage penalty for women within this sample. This contrasts with past research (Budd and Na, 1999), which finds that union wage premiums are stronger at the lower end of the income distribution, where women are generally more concentrated. This raises questions about the consistency of union benefits across sex-based demographic groups.

The structure of this paper is as follows: Section II reviews existing literature on union wage effects and income disparities across demographic and occupational subgroups; Section III outlines the data sources and descriptive statistics; Section IV describes the study's econometric models; Section V presents and interprets the results; Section VI discusses implications, links to existing literature, limitations, and potential extensions of the study.

II. Literature Review

Card (2001) examines how unionisation affects wage structures using longitudinal data, accounting for misclassification errors and unobserved productivity differences. He addresses selection bias in past work by developing a regression model that separates this bias from what he terms the "true union effect." Card finds that lower-skilled workers receive a larger union wage premium than higher-skilled workers, thereby compressing the wage distribution across demographic groups. Consistent with this compression mechanism, our estimates similarly indicate that Black workers receive a larger union wage premium than white workers, narrowing the racial wage gap.

Building on early work examining individual-level union effects, Reilly (1996) introduces the concept of union density, defined as the share of unionised workers in an establishment. The study finds that at lower levels of union density, individuals experience a significant union wage premium regardless of whether they are union members, consistent with the bargaining-power effect. Although our study focuses on the individual effect of union membership rather than density, we likewise observe a union wage premium emerging from our regression models, reinforcing the broader empirical regularity of positive union effects on wages.

Expanding this literature on demographic heterogeneity, Budd and Na (1999) studied whether union members earn higher wages than their non-unionized counterparts, using data from the CPS from 1983 to 1993. They specifically analyzed full-time private sector employees covered by collective bargaining agreements to control for omitted variable bias. They found that a 12% to 14% union wage premium exists under the same employment agreements. Similarly to our paper, Budd and Na introduced different interaction terms to test the union wage premium across different demographics. They found that the premium was higher for men than for women and found no significant difference when testing for race or marital status. These results differ from our own, as we found higher union wage premiums for demographic subgroups concentrated in lower income distribution levels, generally women or Black individuals.

Further contributing to this strand of research on subgroup variation, Blanchflower and Bryson (2003) analyze how the effects of unionization vary across occupational and demographic groups in the United States. Using CPS microdata from 1973 to 2002, they find that the wage premium associated with union membership declined substantially over the period. Their results show that younger, less-educated, and blue-collar workers tend to experience larger union wage premiums than highly educated or white-collar workers. Moreover, their regression analysis reveals that public-sector employees receive higher union wage premiums than those in the private sector, much like the results we found in our study. The approach taken by Blanchflower and Bryson is similar to ours in that it examines differences in union wage effects across demographic groups. However, while their analysis focuses on how the union wage premium varies based on the interaction between variables such as union membership, occupation, and education, our study applies a similar framework but shifts the focus to union membership's interaction with race and gender, in order to highlight disparities across these groups. Additionally, the CPS data used in our study is more recent micro data from 2021 to 2025 as opposed to Blanchflower and Bryson's CPS data spanning from 1973 to 2002.

III. Data and Descriptive Statistics

The data used in this study were sourced from IPUMS CPS ASEC in a yearly format, covering the period from 2021 to 2024. The primary explanatory variable, *union_rev*, was reorganized into two categories from the original three. The third category, "Covered by union but not a member," was included with union members for this study. In the U.S., it is common for individuals to be covered by a union contract—which determines wages—without being formal members of the union, typically because they belong to a bargaining unit. For example, an Economic Policy Institute report found that in 2023, 11.2% of people were represented by a union, while only 10% were members (Shierholz et al.). To ensure accuracy, individuals with missing or invalid responses regarding union membership were dropped. The variable was therefore converted into a dummy, where 1 indicates union coverage or membership and 0 otherwise. For the dependent variable, *incwage*, which measures total pre-tax wage and salary income, all observations with missing values were dropped to maintain consistency in income analyses. Additionally, the sample includes only civilian labour-force participants aged 15 or older.

Covariates also underwent refinement. For the variable *ahrsworkt*, which records total hours worked last week, all missing observations were removed, and only responses within the range of 1–99 hours were retained. The variable for the region, *region_rev*, was re-categorized from its original ten subcategories into the four standard U.S. regions (Northeast, West, Midwest, and South). Respondents with missing region information or unspecified states were dropped. Similarly, race was modified to include only the three main racial groups

(Black, White, and Asian). Dummy variables were created for Black and Asian, with White as the reference category.

After removing missing variables, a dummy variable for private workers, *private_wkr*, was created, where 1 denotes employment in the private sector and 0 denotes employment in the government (the base category). This variable was derived from the worker variable *CLASSWKR*, which identifies whether a respondent was self-employed, works in private industry, in the public sector, in the armed forces, or without pay in a family business or farm (Flood et al., 2025). The last variable modified was *industry*. Because the union membership rate was 32.2% in the public sector compared with 5.9% in the private sector, the public sector was of particular interest. Accordingly, the dummy variable *public_sector* was created from the variable *IND*, which reports the industry type in which the respondent performs their primary occupation. Industries classified as part of the public sector (e.g., executive offices, legislative bodies, and public finance activities) were grouped into this category, where 1 indicates public-sector employment and 0 indicates all other industries. Missing observations and respondents who did not specify their industry were dropped.

Table 2 presents summary statistics for all variables. The typical respondent in this sample is 42.56 years old, works 38.45 hours per week, and has 14.24 years of schooling, implying that the majority have at least some college experience. The average annual pre-tax wage income is $64,829, though the range is wide, with a maximum exceeding $2.1 million. The log-transformed income variable (*lincwage*) averages 10.70, indicating a more normalized distribution than raw income. Within the sample, 11.1% of respondents report union membership. The gender distribution is nearly even, with women making up 49% of all respondents. Approximately 10% identify as Black and 6.6% as Asian; the remaining 83.4% identify as White. Regionally, 36.6% reside in the South, 27.4% in the West, and 19.9% in the Midwest, with the remainder in the Northeast. Employment-wise, 83.5% of the sample works in the private sector, while 4.9% work in the public sector.

Table 3 provides a breakdown of sample characteristics by union membership status. Union members are generally older than non-members (44.77 vs. 42.28) and work slightly more weekly hours (40.67 vs. 38.17). They also report higher average annual income ($69,181 vs. $64,285) and a higher average log wage (10.92 vs. 10.67). Union members have slightly more education (14.71 vs. 14.18 years). Regionally, union membership is more concentrated in the West (35.0% vs. 26.5% among non-members) and less concentrated in the South (22.7% vs. 38.4%). Sectoral differences are especially pronounced: 14.9% of union members work in the public sector, compared with only 3.67% of non-members. Non-union respondents are far more likely to work in the private sector (87.7%) than union members (50%). These figures underscore the distinct demographic and labor-market profiles of union and non-union workers—differences crucial for interpreting the estimated union wage premium.

IV. Econometric Modeling

Multiple Regression Models

$$\begin{aligned} incwage = {} & \beta_1 + \beta_2 union\_rev + \beta_3 black + \beta_4 asian + \beta_5 female + \beta_6 age + \beta_7 midwest + \beta_8 south \\ & + \beta_9 west + \beta_{10} private\_wkr + \beta_{11} public\_sector + \beta_{12} educ\_rev + \beta_{13} ahrsworkt + e \end{aligned}$$

$$\begin{aligned} lincwage = {} & \beta_1 + \beta_2 union\_rev + \beta_3 black + \beta_4 asian + \beta_5 female + \beta_6 age + \beta_7 midwest + \beta_8 south \\ & + \beta_9 west + \beta_{10} private\_wkr + \beta_{11} public\_sector + \beta_{12} educ\_rev + \beta_{13} ahrsworkt + e \end{aligned}$$

To examine the factors influencing income in the United States, we estimate two OLS specifications. The first uses raw income (*incwage*) as the dependent variable, while the second uses the natural logarithm of income (*lincwage*), which allows coefficients to be interpreted as percentage differences and helps address skewness in the income distribution. The main independent variable in both models is *union_rev*, a dummy variable equal to 1 if the individual is a member of a union and 0 otherwise.

A broad set of control variables are included to account for demographic characteristics, regional differences, employment type, education, and weekly hours worked. These controls help isolate the relationship between union membership and income while offering a fuller understanding of wage variation.

Race is included using two binary variables: *black* and *asian*, with white individuals being the base group. Persistent racial wage gaps in the U.S. have been well documented—for example, the Economic Policy Institute reports that median hourly earnings of Black workers are roughly 24.4 percent lower than those of White workers, even after adjusting for education (Wilson and Darity Jr., 2022). Including race allows the model to test whether such disparities persist after conditioning on other characteristics.

Gender is measured using a dummy variable for *female*. Gender wage disparities remain substantial; according to the Pew Research Center, women earned about 85 cents for every dollar earned by men in 2024 (Fry and Aragão, 2025). Including gender enables the model to assess whether women still face a wage penalty after accounting for education, experience, union membership, and other factors.

Age is included as a continuous proxy for experience (*age*). Earnings generally rise with age as individuals acquire skills and seniority, though this relationship may flatten or decline later in life. Including *age* therefore controls for experience-related wage differences.

Geographic variation is captured through three regional dummy variables: *midwest*, *south*, and *west*, with the Northeast as the reference group. Wages can vary widely by region due to differences in cost of living, industry presence, and state-level labor laws. For example, the South generally has lower average wages and weaker union presence compared to the Northeast or West. According to the Bureau of Labor Statistics (Union Members, 2024), as of 2024 nearly all southern states had unionization rates below the national average. Including these regional variables helps account for geographic wage patterns and region-specific labor market dynamics.

The model also includes *private_wkr* and *public_sector* to distinguish between people working in the private sector and those in government jobs. These sectors have different wage structures, with public sector jobs often having stronger union presence. Research from the Congressional Budget Office found that federal workers tend to earn 23.2% more than private sector employees with similar qualifications, especially in lower-skilled positions (Burns). These variables therefore help account for wage differences arising from the employment sector rather than worker characteristics.

Additionally, *educ_rev* is included as a continuous measure of years of schooling, given its strong and well-established correlation with earnings. According to a Bureau of Labor Statistics report, workers with a bachelor's degree earn nearly twice as much, on average, as those with only a high school diploma (Employment Projections - Education Pays, 2024). Including *educ_rev* ensures that estimated wage differences are not simply driven by unequal access to schooling.

Lastly, *ahrsworkt* measures the average number of hours worked per week. This is important because people who work more hours generally earn more, regardless of other factors. Including this variable helps separate the impact of actual labor input from other wage determinants in the model, including education, union status, and race.

Multiple Linear Regression Models with Interaction Terms

Including interaction terms is essential when the effect of one independent variable may depend on the value of another. For example, the wage premium associated with union membership might differ based on race or gender. Without interaction terms, the model assumes that effects are the same across all groups, which can lead to inaccurate conclusions. By incorporating these terms, important group specific differences can be revealed and improve the model's accuracy.

Model 3

$$\begin{aligned} incwage = {} & \beta_1 + \beta_2 union\_rev + \beta_3 black + \beta_4 asian + \beta_5 female + \beta_6 age + \beta_7 midwest + \beta_8 south \\ & + \beta_9 west + \beta_{10} private\_wkr + \beta_{11} public\_sector + \beta_{12} educ\_rev + \beta_{13} ahrsworkt + \beta_{14} union\_rev\_black + e \end{aligned}$$

This model includes an interaction between union membership and being Black. It tests whether the wage premium from unionization is different for Black workers than for others. Past research shows that unions have historically played a big role in helping close wage gaps for Black workers (Wilson and Darity Jr., 2022). This model helps us see if that still holds true today, or if there are still differences in how much unionization helps.

Model 4

$$\begin{aligned} incwage = {} & \beta_1 + \beta_2 union\_rev + \beta_3 black + \beta_4 asian + \beta_5 female + \beta_6 age + \beta_7 midwest + \beta_8 south \\ & + \beta_9 west + \beta_{10} private\_wkr + \beta_{11} public\_sector + \beta_{12} educ\_rev + \beta_{13} ahrsworkt + \beta_{14} union\_rev\_female + e \end{aligned}$$

Model 4 looks at the interaction between union membership and being female. It's useful because women still tend to earn less than men on average, and unions could help reduce that gap. Including this interaction lets us test if women gain just as much from being in a union as men do, or if gender-based wage gaps continue even in unionized jobs. Pew Research found that union pay structures often help reduce wage inequality (Fry and Aragão, 2025), so this term is key for testing that idea.

Model 5

$$\begin{aligned} incwage = {} & \beta_1 + \beta_2 union\_rev + \beta_3 black + \beta_4 asian + \beta_5 female + \beta_6 age + \beta_7 midwest + \beta_8 south \\ & + \beta_9 west + \beta_{10} private\_wkr + \beta_{11} public\_sector + \beta_{12} educ\_rev + \beta_{13} ahrsworkt + \beta_{14} union\_rev\_south + e \end{aligned}$$

This model includes an interaction between union membership and living in the South. That matters because union laws and culture vary a lot by region, and the South generally has fewer protections for workers and lower union membership overall (Union Members, 2024). So this model tests whether unions are just as effective at raising wages in the South as they are in other regions, or if their impact is more limited in that context.

Model 6

$$\begin{aligned} \text{uhrsworkt} = {} & \beta_1 + \beta_2 \text{nchild} + \beta_3 \text{age} + \beta_4 \text{educclean} + \beta_5 \text{black} + \beta_6 \text{asian} + \beta_7 \text{amindian} \\ & + \beta_8 \text{midwest} + \beta_9 \text{south} + \beta_{10} \text{west} + \beta_{11} \text{single} + \beta_{12} \text{yg2015} \\ & + \beta_{13} (\text{nchild} \cdot \text{yg2015}) + \varepsilon \end{aligned}$$

Model 6 looks at whether the effect of union membership changes depending on if someone works in the private sector. Since unionization and bargaining power can look very different in private companies compared to public institutions, it's important to separate those out. This interaction helps us understand whether private-sector unions are delivering the same wage benefits that public-sector unions typically do.

Model 7

$$\begin{aligned} incwage = {} & \beta_1 + \beta_2 union\_rev + \beta_3 black + \beta_4 asian + \beta_5 female + \beta_6 age + \beta_7 midwest + \beta_8 south \\ & + \beta_9 west + \beta_{10} private\_wkr + \beta_{11} public\_sector + \beta_{12} educ\_rev + \beta_{13} ahrsworkt + \beta_{14} union\_rev\_public\_sector + e \end{aligned}$$

This model includes an interaction between union membership and working in the public sector. Jobs within the public sector have a higher unionization rate than jobs in the private sector do, showing that there is an existing correlation between the variables. So, this interaction term helps us figure out if unionized public-sector jobs lead to stronger wage premiums than other types of jobs. It's a helpful way to see if public institutions are doing a better job at protecting worker's pay through unions.

Model 8

Model 8 incorporates all five interaction terms simultaneously, allowing for a comprehensive analysis of how union wage effects vary across demographic and occupational subgroups when these factors are considered together. By including all interaction terms in one specification, it becomes possible to assess how the significance and direction of each effect change once overlapping influences are accounted for. This approach helps determine whether the relationships identified in previous models hold when demographic, regional, and sectoral interactions are jointly examined, providing a fuller picture of unionization's heterogeneous impact on income.

V. Description of Results

Model 1

Model 1 estimates the relationship between union membership and income, controlling for race, gender, age, region, sector, education, and hours worked. The model has an R^2 value of 0.1886, which implies that approximately 18.86% of the variation in income is explained by the multiple regression. The coefficient for *union_rev* is $721.82. This means that, holding all other variables constant, being a union member is associated with earning $721.82 less than a comparable non-union worker. In terms of semi-elasticity, it is approximately -1.11%, indicating that being part of a union is associated with a 1.11% decrease in income, when evaluated at the mean and holding all other variables constant. However, the coefficient was not statistically significant, with a t-statistic of -0.61, falling below the critical value at the 1% level.

Race had a notable effect, where being Black led to a $6,385.71 loss in income compared to White individuals, and this result is statistically significant. In contrast, being Asian was associated with a $2,582.24 increase in income. While this effect is only marginally significant, it may indicate a form of racial gap, which aligns with data from Wilson's Economic Policy Institute report focused on exploring Black-White disparities in the labor market. Similarly, gender showed an effect: women earned $18,601.40 less than males, demonstrating a persistent gender wage gap. Age had a positive influence, where each additional year resulted in an increase of $566.89 in annual income.

Regional differences were also observed relative to the Northeast: living in the Midwest corresponded to earning $5,492.11 less per year, and living in the South corresponded to earning $6,030.79 less. No statistically significant difference was found for the West relative to the Northeast.

The type of employment also had an impact. Those employed in private companies earned $16,263.19 more than government workers. Conversely, individuals in the public sector earned $16,062.53 more than those in other industries. Education had the most significant effect on income, with a t-value of 70.02 and a coefficient of $9,073.98, indicating that an additional year of schooling is associated with a $9,073.98 increase in annual income, highlighting the strong monetary return on human capital. Finally, weekly hours worked also mattered; each additional hour increased annual income by $1,161.49.

Model 2

Model 2 is a multiple log-linear regression model. The R^2 value of this model is 0.362, which implies that 36.20% of the variance of *linc_wage* is explained by the independent variables. This highlights the importance of adding more control variables to reduce omitted variable bias.

The coefficient β_2 for *union_rev* is 0.1093719 and is statistically significant at the 1% level. The semi-elasticity is approximately 10.902, indicating that being a union member is associated with a 10.9% increase in income when evaluated at the mean.

The regression model showed that race had an observable effect on income. Being Black showed a wage penalty of approximately 10.9% compared to White individuals, while being Asian implied a 4.75% wage premium. Additionally, sex had a significant impact on income, with females earning approximately 23.6% less than male counterparts.

Age also had a notable impact on income, with a 1.13% increase for every additional year in age. Regional differences were another factor affecting income levels. Compared to individuals living in the Northeast, those in the Midwest earned 6.43% less, and those in the South earned 6.38% less. As for the West, no statistically significant difference was found between individuals living in that region and those in the Northeast.

Moreover, the type of employment had a notable impact on wages. Individuals employed in the public sector experienced a 25.5% wage premium, whereas those in the private sector experienced a 14.4% wage premium. Education was another variable with a substantial effect on wages: for each additional year of education, wages increased by 12.4%. Lastly, hours worked per week had a modest but statistically significant effect on income, with each additional hour worked associated with approximately a 3.01% increase in wages.

Multiple Regression Models with Interaction Terms: Model 3, 7, 8

The models below include the interaction terms discussed previously. In attempting to focus on the interaction terms specifically, interpretations of previous variables introduced remain the same, regardless of the change in their values. Moreover, due to a lack of statistical significance, models 4, 5, and 6 will not be analyzed.

Model 3

Model 3 includes an interaction term between union membership and race through the variable *union_rev_black*. This was constructed by multiplying the two dummy variables *union_rev* and *black*. The interaction term, as reported in Table 6, was found to be statistically significant at the 5% level, with a coefficient of $7,460, holding other variables constant. This indicates that unionized Black workers earn higher wages compared to individuals affected by those two variables separately. This suggests that being part of a union is monetarily beneficial for Black individuals. The marginal wage difference supports this theory, as the change in wage for someone who is part of a union varies depending on whether they are Black or not. The wage premium for those who are both union members and Black is $5,956.18,[1] compared to their White unionized counterparts, who receive $1,503.48[2] less in wage premium. These results align with Wilson's (Wilson and Darity Jr., 2022) findings, which suggest that wage premiums are larger for Black workers than for White ones.

[1] $Difference_{union\_rev=1,black=1}$ = -1,503.48 + 7,459.66 = 5,956.18.
[2] $Difference_{union\_rev=1,black=0}$ = -1,503.48.

Model 7

Model 7 introduces the interaction term between union membership and the public sector industry (*union_rev_public_sector*). This interaction is created by combining the dummy variables *union_rev* and *public_sector*. As reported in Table 6, the interaction term coefficient is $8,053.24, which is statistically significant at the 5% level. This means that individuals who are both part of a union and work in the public sector experience an $8,053.24 change in income compared to what would be predicted from the separate impacts of unionization and public sector employment alone. Holding other variables constant, the predicted marginal effects were calculated. A unionized worker in the public sector is predicted to earn approximately $6,338.77,[3] more than a non-union public worker. Among workers outside the public sector, union membership is associated with a $1,713.49[4] decrease in income relative to non-union workers. This implies that union wage premiums are more substantial in government employment environments compared to other industries.

Model 8

Model 8 introduces all five interaction terms simultaneously, providing an opportunity to thoroughly examine the variables' impact on wage and salary income. In this model, with all variables controlled for, the coefficient for *union_rev* decreases to $4,774.07 and becomes statistically significant at the 10% level, whereas previously it had remained consistently insignificant.

At first glance, the average coefficient on *union_rev* suggests a decrease in income. However, this interpretation is misleading: once interaction terms are included, the effect of union membership varies across groups, so the single *union_rev* coefficient no longer represents an overall effect—it applies only to the baseline category. Additionally, all variables (excluding interaction terms) continue to be statistically significant at varying levels, except for *west*, which has remained insignificant in all lin-lin models.

When comparing interaction models (3, 5, and 7) with Model 8, coefficients remain of the same sign throughout. However, the Model 5 interaction became statistically significant at the 10% level, which points to a weaker monetary return for union members in the South.

The interaction between union membership and race (*union_rev_black*) remains positive and statistically significant, increasing its coefficient by $1,023. Furthermore, the interaction term *union_rev_public_sector* saw an increase of $2,499 compared to Model 7 and increased in statistical significance from the 5% to the 1% level. This further supports that public sector unionization provides the highest income returns.

VI. Conclusion

This paper sought to investigate whether union membership affects wages and salary income from 2021 to 2024, taking into account how demographic variables such as race, gender, region, and employment type influence this relationship. Although these estimates are not causal, they show that—holding other variables constant—the association between union membership and earnings is heterogeneous across demographic groups. Results confirmed that union membership is positively correlated with considerable salary premiums in general. This is particularly evident with the interaction term *union_rev_black*, which shows that Black union members benefit further from wage premiums compared to the base group, White individuals. This suggests that unions may play a role in reducing the racial wage gap, aligning with Wilson's findings (Wilson and Darity Jr., 2022), which indicate that unionization reduces wage inequality through wage compression.

However, the analysis could not establish that unionization helped reduce the gender wage gap, given the lack of statistical significance of the interaction term *union_rev_female*. Regardless of statistical significance, the

[3] $Difference_{union\_rev=1,public\_sector=1}$ = -1,713.48 + 8,053.24 = 6,338.76.
[4] $Difference_{union\_rev=1,public\_sector=0}$ = -1,713.48.

analysis presented evidence that males still gain further from unionization than females, consistent with Budd's findings (Budd and Na, 1999). Furthermore, a substantial wage premium was highly associated with union membership in the public sector compared to other industries, as shown in Model 7. In contrast, employment type was found to have no statistical significance when testing the interaction term *union_rev_private_wkr* individually or with all interaction terms in Model 8.

Regarding regional impact, the interaction term *union_rev_south* presented significance at the 10% level in Model 8. For this specification, wage premiums were significantly lower in the South compared to the base group, showing regional differences in union effectiveness. This reinforces the U.S. Bureau of Labor Statistics' report indicating significantly lower unionization rates in the South (Union Members, 2024).

With this in mind, this study encountered several limitations. The nature of the cross-sectional data limited the study's ability to assess changes over time. For example, union membership could lead to higher wages as tenure increases. Moreover, the CPS sampling design may have introduced bias due to potential clustering and autocorrelation among individuals from the same household, violating the OLS assumption of independent observations. Additionally, multicollinearity in the interaction terms may have affected their statistical significance in Model 8.

A possible avenue for future research could involve investigating how non-wage benefits, such as healthcare or paid leave, trade off with wages. Future studies could also explore how different industries—particularly high- versus low-skill sectors—differ in union wage premiums. Furthermore, it would be relevant to analyze how wages and employment respond to labor policy changes across jurisdictions, and whether union coverage mitigates or amplifies those effects, for example, by providing wage protection or influencing unemployment risk. In terms of macroeconomic conditions, future research could assess how shocks affect employment and wages, and the extent to which unionization mediates these impacts.

Overall, further research can help clarify who benefits most from unionization and under what conditions, enabling more effective targeting of organizing efforts and benefit design.

References

Blanchflower, D. G., & Bryson, A. (2003, September). *What effect do unions have on wages now and would 'what do unions do' be surprised?* (Tech. rep. No. w9973). National Bureau of Economic Research (NBER). https://www.nber.org/papers/w9973

Budd, J. W., & Na, I.-G. (1999, October). The union membership wage premium for employees covered by collective bargaining agreements [SSRN Abstract No. 180090]. https://ssrn.com/abstract=180090

Burns, M. (2024, July). BLS Report Shows Government Workers Make 40% More Than Private Industry Employees. *American Society of Employers*. https://www.aseonline.org/News-Events/Articles/bls-report-shows-governmentworkers-make-40-more-than-private-industry-employees

Card, D. (2001). The effect of unions on wage inequality in the u.s. labor market. *Industrial & Labor Relations Review*. https://www.jstor.org/stable/2696012

Congressional Budget Office. (2017, April). *Comparing the compensation of federal and private-sector employees, 2011 to 2015* (tech. rep. No. 52637). Congressional Budget Office. https://www.cbo.gov/publication/52637

Flood, S., et al. (2025). IPUMS CPS: Version 13.0 [Dataset]. https://doi.org/10.18128/D030.V13.0

Fry, R., & Aragão, C. (2025, March). Gender Pay Gap in U.S. Has Narrowed Slightly over 2 Decades. *Pew Research Center*. https://www.pewresearch.org/short-reads/2025/03/04/gender-pay-gap-in-us-has-narrowed-slightly-over-2-decades/

Reilly, K. T. (1996). Does union membership matter? the effect of establishment union density on the union wage differential. *The Review of Economics and Statistics*. https://www.jstor.org/stable/2109803

Shierholz, H., et al. (2024, January). Workers Want Unions, but the Latest Data Point to Obstacles in Their Path. *Economic Policy Institute*. https://www.epi.org/publication/unionmembership-data/

U.S. Bureau of Labor Statistics. (2024, August). Employment Projections - Education Pays. *U.S. Bureau of Labor Statistics*. https://www.bls.gov/emp/chart-unemployment-earnings-education.htm

U.S. Bureau of Labor Statistics. (2025, January). *Union members - 2024* (tech. rep.) (News Release). U.S. Bureau of Labor Statistics. https://www.bls.gov/news.release/pdf/union2.pdf

Wilson, V., & Darity Jr., W. (2022, March). Understanding Black-White Disparities in Labor Market Outcomes Requires Models That Account for Persistent Discrimination and Unequal Bargaining Power. *Economic Policy Institute*. https://www.epi.org/unequalpower/publications/understanding-black-white-disparities-in-labor-market-outcomes/

VII. Appendix

Table 1: Description of Variables

Variable Name	Variable Description	Variable Type
incwage	Respondent's total pre-tax wage and salary income	Continuous
lincwage	Log of respondent's total pre-tax wage and salary income	Continuous
union_rev	Dummy variable: 1 = union member, 0 = non-member	Dummy
sex	Respondent's sex: Female or Male	Categorical
race	Respondent's race: Black, White, or Asian	Categorical
region	Respondent's region: Northeast, South, Midwest, West	Categorical
south	Dummy variable: 1 = South, 0 = other regions	Dummy
west	Dummy variable: 1 = West, 0 = other regions	Dummy
midwest	Dummy variable: 1 = Midwest, 0 = other regions	Dummy
private_wkr	Dummy variable: 1 = private sector, 0 = other	Dummy
public_sector	Dummy variable: 1 = public sector, 0 = other	Dummy
female	Dummy variable: 1 = female, 0 = male	Dummy
black	Dummy variable: 1 = black, 0 = other	Dummy
asian	Dummy variable: 1 = asian, 0 = other	Dummy
ahrsworkt	Total hours worked in previous week	Continuous
age	Age of respondent (15–85)	Continuous
union_rev_black	Interaction: 1 = union member and black	Dummy
union_rev_female	Interaction: 1 = union member and female	Dummy
union_rev_south	Interaction: 1 = union member and South	Dummy
union_rev_private_wkr	Interaction: 1 = union member and private worker	Dummy
union_rev_public_sector	Interaction: 1 = union member and public sector	Dummy

Table 2: Summary Statistics of Variables

VARIABLES	N	Mean	SD	Min	Max
age	40,956	42.56	14.49	15	85
ahrsworkt	40,956	38.45	11.93	1	99
incwage	40,956	64,829	77,981	0	2.100e6
educ_rev	40,956	14.24	2.803	0	21
midwest	40,956	0.199	0.399	0	1
south	40,956	0.366	0.482	0	1
west	40,956	0.274	0.446	0	1
private_wkr	40,956	0.835	0.371	0	1
public_sector	40,956	0.0492	0.216	0	1
union_rev	40,956	0.111	0.314	0	1
female	40,956	0.490	0.500	0	1
black	40,956	0.100	0.301	0	1
asian	40,956	0.0660	0.248	0	1
lincwage	39,956	10.70	0.981	1.609	14.56

Table 3: Summary Statistics of Variables According to Union Membership

VARIABLES	N (union_rev=0)	Mean (union_rev=0)	N (union_rev=1)	Mean (union_rev=1)	Diff. in Mean
age	36,400	42.28	4,556	44.77	2.49
ahrsworkt	36,400	38.17	4,556	40.67	2.50
incwage	36,400	64,285	4,556	69,181	4,869
educ_rev	36,400	14.18	4,556	14.71	0.53
midwest	36,400	0.200	4,556	0.195	-0.005
south	36,400	0.384	4,556	0.227	-0.157
west	36,400	0.265	4,556	0.350	0.085
private_wkr	36,400	0.877	4,556	0.500	-0.377
public_sector	36,400	0.0367	4,556	0.149	0.1123
female	36,400	0.490	4,556	0.487	-0.003
black	36,400	0.0999	4,556	0.104	0.0041
asian	36,400	0.0656	4,556	0.0687	0.0031
lincwage	35,456	10.67	4,500	10.92	0.25

Table 4: Linear-Linear Regression Models without Interaction Terms

VARIABLES	**Model 1**
union_rev	-721.8 (1,177)
black	-6,386*** (1,187)
asian	2,582* (1,427)
female	-18,601*** (711.4)
age	566.9*** (24.20)
midwest	-5,492*** (1,170)
south	-6,031*** (1,057)
west	-1,282 (1,101)
private_wkr	16,263*** (1,162)
public_sector	16,063*** (1,878)
educ_rev	9,074*** (129.6)
ahrsworkt	1,161*** (30.03)
ahrsworkt_sq	
Constant	-134,236*** (2,815)
Observations	40,956
R-squared	0.189

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

Table 5: Log-Linear Regression Models without Interaction Terms

VARIABLES	Model 2
union_rev	0.109*** (0.0132)
black	-0.109*** (0.0135)
asian	0.0475*** (0.0161)
female	-0.236*** (0.00804)
age	0.0113*** (0.000275)
midwest	-0.0643*** (0.0132)
south	-0.0638*** (0.0119)
west	-0.0148 (0.0124)
private_wkr	0.144*** (0.0131)
public_sector	0.255*** (0.0211)
educ_rev	0.124*** (0.00147)
ahrsworkt	0.0301*** (0.000343)
Constant	7.298*** (0.0321)
Observations	39,956
R-squared	0.362

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

Table 6: Multiple Regression Models with Interaction Terms

VARIABLES	Model 3	Model 4	Model 5	Model 6	Model 7	Model 8
union_rev	-1,503 (1,237)	-1,516 (1,585)	-26.76 (1,337)	-448.9 (1,820)	-1,713 (1,255)	-4,774* (2,734)
black	-7,245*** (1,258)	-6,387*** (1,187)	-6,365*** (1,187)	-6,383*** (1,187)	-6,377*** (1,187)	-7,339*** (1,260)
asian	2,582* (1,427)	2,577* (1,427)	2,572* (1,427)	2,582* (1,427)	2,562* (1,427)	2,536* (1,427)
female	-18,597*** (711.3)	-18,778*** (749.6)	-18,610*** (711.4)	-18,604*** (711.6)	-18,576*** (711.4)	-18,809*** (750.1)
age	566.3*** (24.20)	566.8*** (24.20)	566.7*** (24.20)	567.0*** (24.20)	567.4*** (24.20)	565.7*** (24.20)
midwest	-5,489*** (1,170)	-5,478*** (1,171)	-5,457*** (1,171)	-5,486*** (1,171)	-5,486*** (1,170)	-5,448*** (1,171)
south	-6,007*** (1,057)	-6,020*** (1,057)	-5,774*** (1,083)	-6,022*** (1,058)	-5,997*** (1,057)	-5,602*** (1,085)
west	-1,249 (1,101)	-1,283 (1,101)	-1,268 (1,101)	-1,276 (1,101)	-1,256 (1,101)	-1,227 (1,101)
private_wkr	16,224*** (1,163)	16,348*** (1,168)	16,302*** (1,163)	16,368*** (1,278)	16,015*** (1,167)	15,387*** (1,336)
public_sector	16,019*** (1,878)	16,208*** (1,888)	16,117*** (1,878)	16,061*** (1,878)	13,359*** (2,221)	12,772*** (2,298)
educ_rev	9,074*** (129.6)	9,070*** (129.7)	9,076*** (129.6)	9,074*** (129.6)	9,081*** (129.6)	9,083*** (129.7)
ahrsworkt	1,162*** (30.03)	1,162*** (30.03)	1,162*** (30.03)	1,161*** (30.03)	1,161*** (30.03)	1,161*** (30.03)
union_rev_black	7,460** (3,622)				8,483** (3,714)	
union_rev_female		1,667 (2,228)		2,386 (2,296)		
union_rev_south			-2,851 (2,600)		-4,398* (2,668)	
union_rev_private_wkr			-465.3 (2,368)	3,070 (2,714)		
union_rev_public_sector			8,053** (3,538)	10,552*** (4,022)		
Constant	-134,117*** (2,815)	-134,183*** (2,816)	-134,401*** (2,819)	-134,332*** (2,856)	-134,051*** (2,816)	-133,405*** (2,887)
Observations	40,956	40,956	40,956	40,956	40,956	40,956
R-squared	0.189	0.189	0.189	0.189	0.189	0.189

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

References

Blanchflower, D. G., & Bryson, A. (2003, September). *What effect do unions have on wages now and would 'what do unions do' be surprised?* (Tech. rep. No. w9973). National Bureau of Economic Research (NBER). https://www.nber.org/papers/w9973

Budd, J. W., & Na, I.-G. (1999, October). The union membership wage premium for employees covered by collective bargaining agreements [SSRN Abstract No. 180090]. https://ssrn.com/abstract=180090

Burns, M. (2024, July). BLS Report Shows Government Workers Make 40% More Than Private Industry Employees. *American Society of Employers*. https://www.aseonline.org/News-Events/Articles/bls-report-shows-governmentworkers-make-40-more-than-private-industry-employees

Card, D. (2001). The effect of unions on wage inequality in the u.s. labor market. *Industrial & Labor Relations Review*. https://www.jstor.org/stable/2696012

Congressional Budget Office. (2017, April). *Comparing the compensation of federal and private-sector employees, 2011 to 2015* (tech. rep. No. 52637). Congressional Budget Office. https://www.cbo.gov/publication/52637

Flood, S., et al. (2025). IPUMS CPS: Version 13.0 [Dataset]. https://doi.org/10.18128/D030.V13.0

Fry, R., & Aragão, C. (2025, March). Gender Pay Gap in U.S. Has Narrowed Slightly over 2 Decades. *Pew Research Center*. https://www.pewresearch.org/short-reads/2025/03/04/gender-pay-gap-in-us-has-narrowed-slightly-over-2-decades/

Reilly, K. T. (1996). Does union membership matter? the effect of establishment union density on the union wage differential. *The Review of Economics and Statistics*. https://www.jstor.org/stable/2109803

Shierholz, H., et al. (2024, January). Workers Want Unions, but the Latest Data Point to Obstacles in Their Path. *Economic Policy Institute*. https://www.epi.org/publication/unionmembership-data/

U.S. Bureau of Labor Statistics. (2024, August). Employment Projections - Education Pays. *U.S. Bureau of Labor Statistics*. https://www.bls.gov/emp/chart-unemployment-earnings-education.htm

U.S. Bureau of Labor Statistics. (2025, January). *Union members - 2024* (tech. rep.) (News Release). U.S. Bureau of Labor Statistics. https://www.bls.gov/news.release/pdf/union2.pdf

Wilson, V., & Darity Jr., W. (2022, March). Understanding Black-White Disparities in Labor Market Outcomes Requires Models That Account for Persistent Discrimination and Unequal Bargaining Power. *Economic Policy Institute*. https://www.epi.org/unequalpower/publications/understanding-black-white-disparities-in-labor-market-outcomes/

Fall 2025 Issue IV Volume II