Macalester Journal of Economics

Vol. 33 - Summer 2024

Macalester Journal of Economics

Volume 33 - Summer 2024

Contents

Foreword

The Macalester College chapter of Omicron Delta Epsilon, the international honors society in economics, proudly edits the Macalester Journal of Economics (MJE) every year. This year's editors – Hufsa Ahmed ('24), Anna Durall ('24), Mahmoud Majdi ('24), and Emma Nguyen ('24) – have carefully selected eight papers on a variety of important topics that represent the interdisciplinary, high-quality research that our students have produced in the last academic year. The selected group of papers includes four term papers written as part of a 300-level course on econometrics and four capstone research papers written in three courses including International Economic Development, Open Economy Macroeconomics, and Mathematical Modeling. Inequality, poverty, and inter-generational mobility are at the forefront of much the current economic research and policy debates. This year's selection of papers mirrors this trend.

The lead article, authored by Bryson Berry, investigates the impact of parental educational attainment on child poverty in the context of Peru. Relying on a longitudinal cross-sectional survey of children from age 8 to 22, Berry provides further evidence of the importance of parental education levels in shaping a child's poverty. Specifically, Berry explores this relationship through poverty-related outcomes including a child's own educational attainment, literacy, and cognitive test scores as well as smoking and alcohol consumption. Overall, the evidence is mixed with stronger positive effects of a father's education on child outcomes relative to the mother's educational attainment.

Faye Dingle evaluates the efficacy of affordable housing on US homelessness. Using a two-way fixed effects regression model in conjunction with US-wide data on Low Income Housing Tax Credit (LIHTC) housing units and information on unhoused people, Dingle explores the affordable-housing-homelessness relationship. In general, Dingle finds that a rise in affordable housing reduces homelessness. While the direction of the overarching effect holds across most racial and gender groups, the magnitude systematically varies. Specifically, Dingle finds that white- and male-identifying homeless people seem to benefit the most from the LIHTC program.

Finally, Gavin Englestad develops a mathematically-inspired model of wealth inequality. Building on the simple Yard Sale Model (YSM), Englestad develops a model of wealth exchange between agents with heterogenous wealth levels and varying probabilities of winning wealth transfers. Englestad shows that the simplest-iteration of the model is able to replicate US wealth inequality quite well and that a more advanced version of the model – one that accounts for wealth groups among agents that influence wealth transfer probabilities – yields similar overall wealth inequality, but with a notable dominance of the wealthiest group.

Related to the issues of inequality, poverty, and intergenerational mobility are two studies authored by Sophie Biesterfeld and Ella DeMay. Using US county-level data on fertility and food insecurity rates, Biesterfeld, for example, considers the impact of food insecurity on fertility dynamics. Specifically, Biesterfeld develops an instrumental variables model that relies on the cost of a meal as well as vehicle/transportation access to isolate the causal effect of food insecurity on fertility. Biesterfeld finds a positive association between food insecurity and fertility rates along with significant regional disparities that align with the much studied human capital theory.

In contrast, DeMay evaluates the efficacy of infrastructure-based climate resilience in Bangladesh; a country that is subject to many natural calamities including intense, high frequency flooding. Specifically, DeMay investigates how post-disaster responses in spending on education and health are shaped by pre-disaster investments in durable home building materials. Utilizing a difference-in-differences empirical framework DeMay finds that investments in stronger, more durable home building materials insulates families against the negative education and health spending shocks other families experience after exposure to a natural disaster.

In addition to these five studies, the current issue of MJE encompasses one paper on the topic of international trade and macroeconomics, one paper in the realm of transportation economics and industrial organization, and one study at the intersection of sports and labor economics. Each of these additional articles raises a meaningful research question on an important economic and policy-relevant issue. Zoe Felsch, for example, evaluates whether tariff-based policies meaningfully effect the current account and whether they can be used to reverse the US current account deficit? To answer these questions Felsch expands upon the recently developed Tradable Non-Tradable model and integrates both the terms of trade and import tariffs. In line with recent real-world experiences during the US-China trade war, Felsch's analysis shows that tariff policy alone is not sufficient to overturn the large US current account deficit.

Tanya Nangpal moves the discussion from macroeconomic issues (i.e., oil price shocks) to the microeconomic ramifications on US airlines. Specifically, Nangpal evaluates whether changes in oil prices have heterogenous effects on budget versus full-services US airlines. Contrary to common beliefs, Nangpal finds that full-service airline stock prices are more sensitive to fuel price shocks than stock prices of budget airlines.

Rounding out this year's issue of MJE, Samina Stack researches the important issue of discrimination in US labor markets. Similar to prior work, Stack points to the difficulty in identifying and measuring discrimination in the context of unobservable worker characteristics, quality, and productivity. To overcome these persistent challenges, Stack considers the prevalence of discrimination in the context of the NBA labor market. In theory, the presence of many measurable statistics quantifying various facets of a player's skill level and productivity allows the researcher to attribute persistent wage differences to potentially discriminatory behaviors, such as wage gaps across equally skilled and productive domestic and foreign-born players. Empirically, Stack finds that foreign players command a significant wage premium over their equally-skilled domestic counterparts; a surprising result that is robust to numerous sensitivity analysis and piques curiosity for further research.

On behalf of my colleagues in the Economics Department, I am delighted to present the research of these talented students. I am confident that you will find it enlightening and be impressed by the value of a liberal arts education.

Felix Friedt
Associate Professor of Economics

Parent's Education on Outcomes Affecting Child Poverty: Evidence from Peru

Bryson Berry
ECON 426: International Economic Development Capstone (Advisor: Amy Damon)

Education has long been widely understood as an important determinant for escaping poverty. Individuals with low education levels have more difficulty finding better jobs and earning higher income. If these individuals become parents, their children will have significant disadvantages compared to children whose parents have higher incomes. This is due to the inter-generational transmission of poverty (ITP), which suggests that poverty is transmitted like the genetic transmission of traits. Latin American countries historically have had high levels of education and income inequality, which ITP relies on (Aldaz-Carroll & Moran, 2001; Behrman & Rosenzweig, 2002). Economic mobility and the transmission of human capital can therefore be evaluated through the ITP.

I use longitudinal cross-sectional data collected by Young Lives in Peru to explore the potential causal relationship between the education of children's parents and the future poverty outcomes of those children. The data is collected over 5 rounds. In round 1 of the survey, the children are 8 years old, and by round 5 they are 22. Because the surveyed children are still young by round 5, proxies and indirect measurements are used to determine the likelihood that a child will be in poverty. These proxies include specific educational attainment levels of the children and variables that indicate the well-being of a child. Evidence from earlier studies (Aldaz-Carroll & Moran, 2001; Arceo-Gomez & Campos-Vazquez, 2014; Kaestner, 1998) shows that these proxies can serve as useful measures for children escaping poverty.

This paper contributes to existing literature in two ways. First, it expands research on the ITP in Latin America. Most research on the effect of parents' education on children's outcomes has been conducted in the West. This is usually due to the availability of large datasets or the ease with which these studies can be conducted (Chevalier, 2004). Studying Peru is particularly beneficial because of Peru's persistent push for education reform and historically high rates of poverty. Second, I present robust results on the causal relationship between the outcomes of children due to their parents' education by approaching poverty

---

Bryson Berry is a senior with a major in Economics and a minor in Geographic Information Systems. Correspondence concerning this article should be addressed to bberry@macalester.edu.

outcomes as the result of multiple contributing factors. Poverty is a multifaceted issue that is frequently measured in economics literature through income or consumption. However, poverty has many determinants, some of which are the result of conditions that a person faces early in their life. Thus, by estimating childhood poverty using factors that contribute to the likelihood of a child being impoverished, a more complete picture of the mechanisms that contribute to poverty is drawn.

In the following few sections, I begin by reviewing the relevant economic literature, next, I give the economic framework, the empirical approach, and describe the data. Then, I present the main results as well as additional checks for robustness. Finally, I discuss the limitations of the study and offer a conclusion.

Literature Review

There is a breadth of economic literature on inter-generational poverty rates. Typically, economists study the relationship between parental income on the future incomes or poverty rates of their children (Chetty et al., 2018; Solon, 1999). There is also a large body of work on understanding the conditions that affect educational attainment outcomes of children (Escobal D'Angelo, 2012; Glewwe et al., 2017). Escobal D'Angelo et al. (2005) analyze the effects of economic shocks on changes in school attendance levels in Peru. They hypothesize that short-term economic shocks do not reduce school enrollment. Instead, they reduce family investment into human capital. As a result, a child may go to school less, or parents may enroll the child in a worse school. Thus, the quality of education is affected, rather than the quantity. If economic shocks affect the quality of education, then this might be an exogenous variable that can affect childhood poverty outcomes in the future. Their results show that income shocks cause parents to cut spending on education.

Two papers relevant to my study seek to identify the causal processes that contribute to inter-generational mobility, but these studies focus on Western countries. Chevalier (2004) expands upon previous research that uses school leaving age (SLA) with the goal of controlling for non-genetic endogeneity. Chevalier constructs a regression discontinuity that looks at the educational attainment of children before and after the British Education Act of 1972, which increased the minimum school leaving age across Great Britain. He accomplishes this by using schooling data from the British Family Resources Survey (FRS). Chevalier can control genetic factors by looking at educational attainment for children living with their natural parents and for children living with adoptive parents. He finds a positive effect of both parents' education on their children's educational attainment when focusing on biological parents. Furthermore, he concludes that the inter-generational effect is causal in nature.

Oreopoulos et al. (2006) expand upon Chevalier's study and others like it. However, they attempt to find a causal relationship between parental compulsory schooling and the probability that a student repeats a grade in the United States. The data they use is from IPUMS, which gives them access to hundreds of thousands of observations. This data allows them to obtain precise estimates of child and parental education levels. They use an instrument for minimum number of years required before a person can work. Using the instrumental variable estimates, their results suggest that an additional year of parental education reduces the probability that their child repeats a year by 2-4 percent. One of the weaknesses of this paper is that directly determining when a child repeats a grade is difficult due to state and time variation in minimum school age entry. As a result, Oreopoulos et al. (2006) classify grade repetition based on whether the child's educational attainment is below the median for their state. Additionally, there is the likelihood that their instrument is correlated with other state-level changes that affect school progress.

De Walque (2009) helps to answer the nature vs nurture question regarding the effect of parental education on children's educational attainment. Unlike Chevalier or Oreopoulos et al., De Walque's study is based in Rwanda. He focuses on two questions: the first is to estimate the effect of adoptive parents' education on the adoptive children's educational attainment. The second is to compare the effect of the adoptive parents' education to the education of the absent biological parents. To isolate the causal mechanisms, De Walque uses instruments on the exogenous variation of parents' schooling levels and compares biological children to adoptive children. Results from this study show that the education of the most educated adult female in the household has a significant positive effect on the educational attainment of the adoptive child. He finds that the effect of the adult male's education is less comparatively. He also finds that keeping the orphan child in their extended family also has positive effects. However, these findings are only true if the adoptive parent is a relative. This makes determining a direct causal relationship difficult as other endogenous factors must be considered as explanatory variables for the inter-generational effect of education.

Aldaz-Carroll & Moran (2001) study the effect of the inter-generational transmission of poverty (ITP) in Latin America–which, at the time of their paper, was largely neglected by this field of research. Rather than focusing on the economic environment that affects the household, Aldaz-Carroll and Moran focus on "family factors". They use a child's educational attainment as a proxy for whether the child will escape poverty. In their framework, if a child from a household whose parents have not completed primary education goes on to complete their secondary education, they are counted as having escaped poverty. Their results show that across 16 Latin American countries, additional years of mother's and father's education have positive statistically significant effects on the likelihood that a child completes secondary school. Additionally, they find that if the child is female, the probability that a child completes secondary school increases by 12 percent. Tansel (1997) also explores the ITP for children in Côte d'Ivoire and Ghana. She uses a Probit model to estimate primary school attendance and Tobit models for years completed of middle school and high school by gender. She instruments the Probit and Tobit models to correct for bias in the original un-instrumented models. Tansel finds that parent education has a significant influence on the educational attainment of both male and female children. However, she finds that father's education is more important than mother's education. In both papers, whether the parent is male or female effects the academic performance of the child. Furthermore, the gender of the child affects the emphasis that parents' place on their educational success.

Duarte et al. (2018) draw from both Tansel (1997) and Aldaz-Carroll & Moran (2001) in their study of the ITP in Spain. They also employ Becker's Quantity-Quality model but can study a larger sample by using data from the Spanish Statistical Institute. They specify a Probit model to study factors that after the probability that a child is able to attain a secondary level of education. Like previous papers, the study assumes socioeconomic conditions like parental education and gender affect the probability of achieving a secondary education level. They find evidence that the probability that a person has completed a secondary level of education is positively correlated with parents having completed that same level. While the findings of this paper are consistent with other empirical studies of the ITP, the authors lack robustness checks against other factors that were considered in previous literature.

Empirical Strategy and Data

Theoretical Framework

There are various theoretical approaches to understanding how education may affect childhood outcomes. Becker & Lewis (1973) developed the "quantity-quality" model to explain the trade-offs between investments in children. Children are treated as investments, so the number of children a woman has directly affects their outcomes due to the allocation of resources including time and income. Therefore, the model suggests that there is a negative relationship between women's education and the number of children they have (Currie & Moretti, 2002). It also suggests that the income elasticity of demand for children decreases as families (more specifically parents) make more income (Doepke, 2015). This implies then, that lower-income families have a lower income elasticity for child quality and a higher elasticity for quantity.

Another important model for the household allocation of resources is the intra-household distribution model. There are various formulations of this model. One of the earliest formulations of this was developed by Becker in 1964, who took a unitary approach

(Thomas, 1990). That is, all resources from members of a household are pooled together and used to jointly maximize the household welfare function. Birdsall (1991) presents an approach that emphasizes the effect of child birth order. This model assumes that expenditures across all children are not equal due to budget or market constraints. The result is that older children receive a larger allocation of resources than younger children. Furthermore, this constraint can lead to older children studying less and entering the workforce earlier (Ejrnæs & Pörtner, 2004). Similarly, Sen (1983) suggests that parental time and resources that are allocated to a child are determined by the number of siblings, the gender, and the relative age of the child. These models are important for studying and understanding inter-generational poverty transmission.

Empirical Strategy

Borrowing from elements from Aldaz-Carroll & Moran (2001), the first estimating equation uses a linear probability estimation to estimate the probability that a child completes secondary school. The value of the dependent variable is dichotomous, meaning it is one if the child completed secondary school and zero if the child did not. The first estimating equation is written as:

$$ED_i = \beta_0 + \beta_1 ED_{iMother} + \beta_2 ED_{iFather} + X_i + \gamma_h + \epsilon_i \tag{1}$$

Where the dependent variable is the probability that a child i has completed at least secondary school. This is a function of independent variables taken by 2016. The main independent variables are the mother's level of education ($ED_{iMother}$) and the father's level of education ($ED_{iFather}$). Unlike Aldaz-Carroll and Moran, I use parents' level of education rather than a binary variable that represents if a parent completed primary school because the mean level of education for parents is greater than primary school level. X_i is a vector of child-level controls (such as child age and gender) and γ_h is a vector of household-level controls including the age of the parents, number of household members and the wealth of the family. Previous studies tend to show that the educational level of mothers' is usually insignificant when both parents are considered or if the child is male (De Walque, 2009; Plug, 2004; Tansel, 1997). Fathers' education seems to have a stronger effect, especially if the child is male. The issue of mothers' education being insignificant is likely due to collinearity between mother's and father's education (Aldaz-Carroll & Moran, 2001). Potential endogeneity could also affect this result, which I will attempt to control for by considering shocks in the parent's lives that could have affected their educational attainment.

The primary limitation of LPM estimation is that it can return values above one and below zero, which does not make sense in the context of probability. Thus, I compare the linear probability estimation to a Probit estimation of the same independent variables. The Probit estimation is written as:

$$Prob(Secondary = 1) = F[\beta_0 + \beta_1 ED_{iMother} + \beta_2 ED_{iFather} + X_i + \gamma_h + \epsilon_i] \tag{2}$$

Additional estimating equations use child well-being proxies as estimates for childhood poverty. These proxies include a general well-being variable constructed by the Young Lives Survey, if a child is literate, if a child frequently smokes, and if a child has frequently consumed drugs or illicit substances. The proxies of well-being were chosen because there is empirical evidence that these proxies affect poverty outcomes. For instance, illicit substance use has an adverse effect on poverty outcomes (Kaestner, 1998; Pérez et al., 2018), especially in low-income or developing countries.Pérez et al. (2018) study the effects of parental education on students' drug use in Mexico and Argentina. They find that in Mexico, students whose parents had higher education had lower levels of drinking and smoking compared to students whose parents had lower education.

These results are not intended to establish causality between the education level of parents and the various proxies. It is unlikely that the proxy estimates can establish a causal relationship as these proxies are likely to be correlated with one another. Furthermore, it is likely there are significant external factors that could influence the likelihood that a child smokes or uses illicit substances than are correlated with the error term and that I cannot account for. Instead, like the results from equation (2), these results should help to explain the causal relationship between parental education and childhood poverty. The results from all of these estimating equations should allow me to establish the relationship between parental education and childhood poverty outcomes even with the constraint that I do not have access to direct poverty data such as consumption or income data.

Data

For this study, I use data from the Young Lives Longitudinal Survey of Childhood Poverty. The Young Lives survey is a cohort study that follows the lives of 12,000 children across four countries–Ethiopia, India, Peru, and Vietnam–over 15 years from 2001-2002 to 2016. The survey follows children in two cohorts: a younger and older cohort. The younger cohort follows 2,000 children in each country who were all under a year old at the time of the first round of the survey. The older cohort follows 1,000 children in each country who were all between 7 and 8 years old at the time of the first round. Data for children were collected from 20 sentinel sites that allowed Young Lives to monitor the same cohort of children over a long period.

This study specifically focuses on children in the older cohort in Peru. There are a few distinctions between the Young Lives survey in Peru compared to the other countries where data were collected. The survey of the older cohort in Peru is 714 children rather than 1,000. Peru's sample design was slightly different from other Young Lives locations. The sentinel sites were selected non-randomly to intentionally over-sample poorer children, thus the highest-ranking 5 percent of districts in Peru were excluded. As a result, 75 percent of the sample sites were considered "poor" compared to 25 percent considered "non-poor". However, the sampling of clusters in Peru was completely random, whereas in other countries this process was only semi-random.

For this study, I focus on round 5 of the Young Lives survey, but I use data from rounds 3 and 4 to construct my dependent variables and add additional controls. Table 1 shows attrition and the reasons for it between rounds 1 and 5. There is a 14.1 percent attrition rate. Most of the attrition comes from children either refusing to continue the survey (42 children) or children being untraceable (39 children). The attrition rate is a significant limitation since that dataset is already quite small for a study of this nature. All Young Lives Studies suffer from attrition; however, Peru suffers the most.

Table 1

Attrition Between Round 1 and Round 5

	Older Cohort	
Initial Sample Round 1 (2002)	714	
Died	6	0.8%
Refused	42	5.9%
Untraceable	39	5.5%
Living Abroad	19	2.7%
Interviewed in Round 5 (2016)	608	
Attrition	100*	14.1%

Note: The Young Lives survey does not include death within the measured attrition rates. Thus, the row showing attrition is 100 instead of 106 (excluding the 6 deaths).

Table 2 shows descriptive statistics for the child and family variables used in this study. The constructed dataset includes childhood and household-level data. The childhood and household data are from all five rounds of the Young Lives survey. Young Lives modifies the survey questions depending on the round of the data collection to reflect the changing stages of life of the children in the study. Round 3 of the survey asks more personal life questions to the children in the survey. By rounds 4 and 5, the children might be starting families, living independently or with different family members, or attending university or working full-time jobs. I account for this in two ways. First, the household characteristics used as controls come directly from round 3, when the child was 15. This guarantees that most of the children are still living at home and that the influences of their family members can be accurately controlled for. Second, to determine the highest grade level achieved by each child, data from rounds 3 and 4 were combined and a dummy variable was created that identified if a child completed secondary school. Data for the PPVT and math test scores comes from round 3 to remain consistent with the household-level controls. All other data on childhood outcomes is from round 5, as I am testing the final outcomes for the children.

Figure 1 in the appendix further explores the data, showing the proportion of children who complete at least secondary school by gender. About 3 percent more boys complete secondary school than girls which could be explained by multiple factors including parental preferences or shocks in girls' lives such as pregnancy at a young age (Arceo-Gomez & Campos-Vazquez, 2014; De Walque, 2009). Figures 2 and 3 in the appendix show the likelihood that a child completes secondary school given that the father or mother have completed primary school. There is a higher likelihood that the child completes secondary school if a child's mother did not complete primary school than if the father did not complete primary school. If the mother or father complete primary school, the likelihood that the child completes secondary school is roughly the same. These figures reveal a pattern which supports a possible policy implication which I outline later in the paper.

Table 2

Summary Statistics

	N	Mean	SD
Child Characteristics			
Child Sex (Male=1)	714	.541	0.499
Child's Age in Years	608	21.925	0.418
Parent Characteristics			
Father's Education	345	9.142	3.832
Father's Age in Years	355	51.904	8.189
Mother's Education	441	8.007	4.609
Mother's Age in Years	447	47.993	7.365
Household Characteristics (Round 3)			
Dummy for Bio-Father is Household Head	714	.64	0.480
Sex of Household Head (Male=1)	677	.789	0.408
Dummy for Bio-Parent is Household Head	714	.525	0.717
HH Members aged 0-5	677	.461	0.682
HH Members aged 6-12	677	.74	0.834
HH Members aged 13-17	677	.492	0.638
HH Members aged 18-60	677	2.492	1.126
HH Members aged 61+	677	.199	0.510
Wealth Index	675	.583	0.186
Test Scores (Round 3)			
PPVT Test Score	643	96.924	17.300
Math Test Score	671	13.139	5.722

Source: Young Lives Data Constructed (Modified for this paper)
Note: Data selection specified round 5 of the Young Lives survey. Some data from rounds 3 and 4 were modified to include in the round 5 selection in order to account for when child was living at home.

Results

Main Findings

Table 3 shows the regression estimates for equations 1.1 and 1.2. Columns 1 and 2 show the estimating equations with no controls. In both the LPM and Probit models, an additional level of mother's and father's education has a positive statistically significant effect on the probability that a child completes secondary school. These base equations establish that a relationship does exist and is worth exploring further. In columns 3 and 4, the controls are added. Now, the mother's education is no longer statistically significant when controlling for other child and household-level characteristics. The estimated effect of an additional level of father's education on the probability that a child completes secondary school falls to roughly 1.5 percent in the LPM model and 1.6 percent in the Probit model.

These results are consistent with previous literature. Aldaz-Carroll & Moran (2001) used a Logit model on a cross-sectional dataset on Lima, Peru that show similar results, finding that only father's education is statistically significant. More interestingly, these results support the conclusion found in Behrman & Rosenzweig (2002). Specifically, using robust reduced form OLS estimates on the effects of mother's and father's schooling children's schooling, they find that sometimes increasing mother's level of schooling reduces children's level of schooling. This supports the notion that mothers spending time away from home and not providing care for children has an adverse effect on their outcomes. Interestingly, a biological parent being the head of the household has a statistically significant increase on the probability that a child completes secondary school. Therefore, the influence of a biological parent is important to childhood outcomes especially if the parent is highly educated.

Table 3

Effect of Parents' Education on Completing at Least Secondary Education

VARIABLES	(1) LPM	(2) Probit	(3) LPM	(4) Probit
Father's Education	0.0307***	0.0322***	0.0153**	0.0163*
	(0.00761)	(0.00794)	(0.00765)	(0.00865)
Father's Age in Years			-0.00332	-0.00368
			(0.00409)	(0.00473)
Mother's Education	0.0127**	0.0134**	0.000477	0.00295
	(0.00624)	(0.00632)	(0.00621)	(0.00704)
Mother's Age in Years			-0.000799	0.000213
			(0.00437)	(0.00496)
Dummy for Bio-Father is Household Head			0.0928	0.0575
			(0.164)	(0.207)
Child is Male			-0.0276	-0.0371
			(0.0464)	(0.0545)
Child's Age (years)			-0.0388	-0.0382
			(0.0531)	(0.0583)
Sex of Household Head (Male=1)			0.0879	0.146
			(0.186)	(0.227)
Dummy for Bio-Parent is Household Head			0.0951**	0.0994*
			(0.0439)	(0.0509)
HH Members aged 0-5			-0.0290	-0.0262
			(0.0341)	(0.0390)
HH Members aged 6-12			-0.0525*	-0.0512
			(0.0291)	(0.0330)
HH Members aged 13-17			-0.0507	-0.0574
			(0.0373)	(0.0431)
HH Members aged 18-60			0.00619	0.0117
			(0.0253)	(0.0311)
HH Members aged 61+			0.0166	0.00793
			(0.0657)	(0.0795)
Wealth Index			0.799***	0.840***
			(0.155)	(0.179)
Observations	329	329	325	325

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Note: Probit columns show marginal results.

Tables 4 and 5 check the robustness of these results by running regressions on the effect of mother's and father's education separately. With no controls, mother's education has a 2.5 percent marginal effect in both the LPM and Probit models. Father's education has a higher marginal effect of 3.8 percent. When including all controls, mother's education becomes statistically insignificant. The marginal effect of father's education remains statistically significant, although now to the 5 percent level, and decreases to between 1.6 and 1.8 percent.

Table 4

Effect of Mother's Education on Child Completing at Least Secondary Education

VARIABLES	(1) LPM	(2) Probit	(3) LPM	(4) Probit
Mother's Education	0.0256***	0.0255***	0.00757	0.00971*
	(0.00470)	(0.00486)	(0.00507)	(0.00564)
Dummy for Bio-Mother is Household Head			0.0504	0.0692
			(0.129)	(0.137)
Observations	441	441	436	436
Child Controls	No	No	Yes	Yes
Household Controls	No	No	Yes	Yes

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Note: Probit columns show marginal results. Dummy for if mother is head of household is included in household controls.

Table 5

Effect of Father's Education on Child Completing at Least Secondary Education

VARIABLES	(1) LPM	(2) Probit	(3) LPM	(4) Probit
Father's Education	0.0373***	0.0384***	0.0158**	0.0182**
	(0.00623)	(0.00675)	(0.00687)	(0.00800)
Dummy for Bio-Father is Household Head			0.0298	-0.0108
			(0.154)	(0.199)
Observations	345	345	339	339
Child Controls	No	No	Yes	Yes
Household Controls	No	No	Yes	Yes

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Note: Probit columns show marginal results. Dummy for if father is head of household is included in household controls.

To further evaluate the effect of parents' education on children's poverty outcomes, I use proxies for poverty. These proxies have been shown in previous studies to be effective for measuring the likelihood that a child will be in poverty. Table 6 shows LPM and Probit results for the three selected proxies–child literacy, smoking, and alcohol consumption. Columns 1 and 2 show results for the probability of the child being literate. Every level of father's education increases the probability of a child being literate by 1.6 percent in the LPM model and 1.2 percent in the Probit model. Both results are statistically significant. Father's level of education is not significant for the probability that the child smokes or drinks alcohol, however, the signs are positive which is surprising. For all columns except column 6, Mother's education negatively effects the probability of a child being literate, if a child smokes, or if a drinks alcohol. However, none of these results are statistically significant. Thus, other than for child literacy, I cannot assert that parent's education has a significant effect on child poverty through these proxies.

Table 6

Effect of Parental Education on Child Poverty Proxies

	(1)	(2)	(3)	(4)	(5)	(6)
	Child is Literate		Child Smokes		Child Consumes Alcohol	
VARIABLES	LPM	Probit	LPM	Probit	LPM	Probit
Father's Education	0.0164***	0.0119***	0.00313	0.00293	0.00648	0.00167
	(0.00538)	(0.00416)	(0.00697)	(0.00683)	(0.00507)	(0.0151)
Father's Age in Years	0.000886	-0.000130	0.00269	0.00254	0.00486*	0.00240
	(0.00288)	(0.00215)	(0.00373)	(0.00357)	(0.00269)	(0.0215)
Mother's Education	-0.00281	-0.000935	-0.00402	-0.00399	-0.00257	0.000350
	(0.00437)	(0.00315)	(0.00567)	(0.00583)	(0.00406)	(0.00381)
Mother's Age in Years	-0.00564*	-0.00327	-0.00509	-0.00458	0.00657**	-0.00338
	(0.00307)	(0.00214)	(0.00398)	(0.00374)	(0.00288)	(0.0303)
Dummy for Bio-Father is Household Head	-0.0773	-0.0606	0.0291	0.0476	0.0519	0.258
	(0.115)	(0.0861)	(0.149)	(0.159)	(0.110)	(10.56)
Child is Male	-0.0347	-0.0255	0.234***	0.240***	0.140***	0.0881
	(0.0326)	(0.0265)	(0.0423)	(0.0421)	(0.0306)	(0.789)
Child's Age (years)	-0.0461	-0.0313	-0.0223	-0.0223	0.0256	0.0114
	(0.0373)	(0.0247)	(0.0484)	(0.0479)	(0.0363)	(0.103)
Sex of Household Head (Male=1)	0.117	0.0934	-0.168	-0.168	-0.0709	-0.268
	(0.131)	(0.0953)	(0.169)	(0.173)	(0.125)	(10.47)
headpbior3	-0.0369	-0.0283	-0.0138	-0.00950	-0.0343	-0.0197
	(0.0309)	(0.0218)	(0.0400)	(0.0374)	(0.0315)	(0.177)
HH Members aged 0-5	0.00858	0.0121	0.0136	0.00731	0.0348	0.0105
	(0.0240)	(0.0177)	(0.0311)	(0.0288)	(0.0225)	(0.0946)
HH Members aged 6-12	-0.0421**	-0.0289**	-0.000514	0.00161	-0.0334*	-0.0205
	(0.0205)	(0.0147)	(0.0265)	(0.0249)	(0.0193)	(0.184)
HH Members aged 13-17	-0.0280	-0.0157	-0.0164	-0.0188	-0.0234	-0.0187
	(0.0263)	(0.0193)	(0.0340)	(0.0336)	(0.0247)	(0.168)
HH Members aged 18-60	-0.0154	-0.00938	0.0193	0.0181	0.00490	0.00256
	(0.0178)	(0.0129)	(0.0230)	(0.0219)	(0.0167)	(0.0247)
HH Members aged 61+	-0.00730	0.00959	0.0638	0.0461	0.0229	0.00574
	(0.0462)	(0.0422)	(0.0599)	(0.0531)	(0.0430)	(0.0548)
Wealth Index	0.258**	0.173**	0.129	0.122	0.0178	-0.0147
	(0.109)	(0.0809)	(0.141)	(0.139)	(0.107)	(0.140)
Observations	325	325	325	325	307	307

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Note: Probit columns show marginal effects.

Test Scores

Estimating the effect of parent's education on children's cognitive test scores from round 3 serves as a good robustness check for Table 3. This is because if parent's education improves children's performance on cognitive tests, it would make sense that children would also perform well academically, holding other factors constant. The two tests administered by Young Lives were the Peabody Picture Vocabulary test (PPVT) and a math test. The PPVT test administered to the older cohort in round 3 in Peru was modified specifically for Latin America and contains 125 items (Cueto & León, 2012). This test is supposed to be able to measure vocabulary into adulthood, making it a good test to use for all ages. The math test contained two sections. One section contained 20 items that dealt with addition, subtraction, multiplication, division, exponents, and fractions. The other section contained 10 items meant to assess mathematical problem-solving. See Cueto & León (2012) for more details on the tests.

The raw scores from the tests were used in the OLS estimates in Table 7. Column 1 shows the results of parent's education on child's PPVT test score. Each additional level of father's education increases a child's test score by 0.8 points, while each additional level of mother's education increases a child's test score by 0.7 points. Both results are statistically significant to the 1 percent level. Column 2 shows the results of parent's education on child's math test score. Here, only the father's education is statistically significant, increasing a child's scores by roughly 0.4 points per level of education. Given that each parent's education–especially father's education–increases test scores, it is feasible to conclude that each parent's education is, at least to some extent, causally linked to the probability that a child completes secondary school.

Table 7

Effect of Parents' Education on Test Scores (Round 3)

	(1)	(2)
	PPVT	Math
VARIABLES	OLS	OLS
Father's Education	0.827***	0.383***
	(0.271)	(0.0948)
Mother's Education	0.735***	0.103
	(0.230)	(0.0771)
Constant	71.26	23.24
	(46.25)	(14.55)
Observations	307	322
R-squared	0.383	0.298
Child Controls	Yes	Yes
Household Controls	Yes	Yes

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

Note: The PPVT test had 125 questions each worth 1 point. The Math test had 30 questions with a point recorded for each correct answer.

Shocks in Parents' Education

From 1980 to 2000 Peru experienced a civil war primarily between indigenous Peruvians and the central government. Between 1982 and 1997, the Tupac Amaru Revolutionary Movement (MTRA) contributed to a significant amount of violence during the conflict. Many of the parents of the Young Lives children were very young, often school-aged, during this conflict. Therefore, this conflict is an exogenous factor that could have potentially affected the educational attainment of the parents. To test this, I construct a treatment group that consists of parents who were primary school-aged (6-12 years old) at the time of the conflict, separated by mother and father. I interact these treatment groups with the respective mother and father education variables in Tables 8 and 9.

Columns 1 and 2 of Table 8 present statistically significant results for the base models for the effect of father's education on the probability of a child completing secondary school. In the LPM model, the probability increases by about 3.3 percent per level of father's education and the probability increases by 3.4 percent in the Probit model. Surprisingly, if a father was in primary school during the conflict, the likelihood that a child would complete secondary school increases by roughly 0.8 percent in both models. Columns 3 and 4 present models with child and household-level controls. The marginal effect of a father's level of education decreases to roughly 1.5 and 1.7 percent respectively but remains significant. Columns 5 and 6 add mother's age and level of education. Neither father's nor mother's level of education has a significant marginal effect. The effect of father's education if he was in primary school during the conflict is still positive.

Table 9 shows results if the mother experienced the conflict. Columns 1 and 2 also report statistically significant results for the base models. Each level of mother's education increases the probability that a child completes secondary school by about 3.2 percent in both the LPM and Probit models. However, unlike in Table 8, when interacting with the conflict variable, the effect decreases by 1.9 and 1.8 percent respectively. This result is what I expected to see in table 8. In columns 3 and 4, when including child and household level controls, the marginal effect remains statistically significant, but decreases to 1.4 percent in the LPM model and 1.6 in the Probit model. The interaction with the conflict decreases the effect to roughly 0 and 0.2 percent. Columns 5 and 6 add father's age and level of education. In these columns, mother's level of education is no longer statistically significant. Furthermore, the interaction now makes the effect of mother's education negative on child's probability of completing secondary school. However, Father's level of education is statistically significant and has a positive marginal effect of 1.5 and 1.6 percent respectively.

Table 8

Probability of Child Completing at least Secondary Education Given Father Experienced Shock during Primary School

VARIABLES	(1) LPM	(2) Probit	(3) LPM	(4) Probit	(5) LPM	(6) Probit
Father's Education	0.0329***	0.0342***	0.0148*	0.0166*	0.0130	0.0133
	(0.00747)	(0.00812)	(0.00798)	(0.00948)	(0.00872)	(0.0100)
Father's Age in Years	-0.00702	-0.00702	-0.00505	-0.00582	-0.00375	-0.00467
	(0.00428)	(0.00441)	(0.00424)	(0.00494)	(0.00500)	(0.00571)
Mother's Education					0.000579	0.00289
					(0.00624)	(0.00704)
Mother's Age in Years					-0.00107	0.000124
					(0.00441)	(0.00497)
Father was in Primary School 1982-97 During Conflict	-0.156	-0.155	-0.0462	-0.0718	-0.0864	-0.108
	(0.154)	(0.159)	(0.152)	(0.170)	(0.155)	(0.171)
Father's Education * Conflict	0.00846	0.00792	0.00310	0.00474	0.00790	0.00930
	(0.0142)	(0.0154)	(0.0139)	(0.0163)	(0.0142)	(0.0164)
Observations	345	345	339	339	325	325
Child Controls	No	No	Yes	Yes	Yes	Yes
Household Controls	No	No	Yes	Yes	Yes	Yes

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Note: Probit columns show marginal results. Column 5 uses the interaction between father's education and if father was in primary school during the Peruvian Conflict.

Table 9

Probability of Child Completing at least Secondary Education Given Mother Experienced Shock during Primary School

VARIABLES	(1) LPM	(2) Probit	(3) LPM	(4) Probit	(5) LPM	(6) Probit
Mother's Education	0.0321***	0.0323***	0.0137**	0.0157**	0.00268	0.00528
	(0.00629)	(0.00671)	(0.00641)	(0.00739)	(0.00761)	(0.00878)
Mother's Age in Years	0.00386	0.00415	0.00288	0.00385	0.000679	0.00161
	(0.00373)	(0.00391)	(0.00367)	(0.00419)	(0.00473)	(0.00537)
Father's Education					0.0152**	0.0161*
					(0.00766)	(0.00864)
Father's Age in Years					-0.00277	-0.00326
					(0.00417)	(0.00479)
Mother was in Primary School 1982-97 During Conflict	0.188**	0.182*	0.220**	0.219**	0.0824	0.0789
	(0.0956)	(0.0962)	(0.0907)	(0.0996)	(0.101)	(0.110)
Mother's Education * Conflict	-0.0147	-0.0148	-0.0144	-0.0144	-0.00455	-0.00492
	(0.00953)	(0.00995)	(0.00902)	(0.0105)	(0.0101)	(0.0120)
Observations	441	441	436	436	325	325
Child Controls	No	No	Yes	Yes	Yes	Yes
Household Controls	No	No	Yes	Yes	Yes	Yes

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
Note: Probit columns show marginal results. Column 5 uses the interaction between mother's education and if mother was in primary school during the Peruvian Conflict.

Conclusion

In this paper, I attempt to establish a causal relationship between parent's education and the poverty outcomes of their children by analyzing data from the Young Lives Longitudinal Study of Peru. The results of this analysis show that the education of fathers has a more significant effect on children's poverty outcomes when using secondary school completion as a proxy for poverty. In the primary specification, this is true even when controlling for child and household-level factors. The effect of fathers' education is also significant for fathers born during the Peruvian conflict, although this is only the case when mothers' education is not accounted for. Children's cogitative test scores are also positively and significantly affected by fathers' education. Mothers' education frequently has no effect or sometimes a negative effect on children's poverty outcomes, which could be explained by mothers serving more important roles as caretakers in Peruvian households. Neither mothers' nor fathers' education has a significant effect on any poverty proxies except for child literacy. Given that Peru continues to struggle with a high poverty rate, policymakers should continue to improve incentives for people to earn higher levels of education so that the effects can be passed on to their children.

One of the most glaring limitations of this paper is my number of observations. Having only 714 observations significantly hinders my ability to make broad conclusions about the effect of parental education. Studies completed by Chevalier (2004) and Oreopoulos et al. (2006) use tens and hundreds of thousands of observations respectively. A panel dataset would have also been preferred. However, the only constructed panel dataset for Peru has a similarly small number of observations. Access to good panel dataset in Latin America for the purpose of studying ITP are simply not very prevalent.

Moreover, I was limited by the inability to use income or consumption as measurements for poverty. In this paper, I use many proxies to tell a story about the likelihood that a child will be in poverty, but I cannot estimate if a child is in poverty. This analysis would be stronger if I could combine the methods that I use in this paper with estimates where the dependent variable was income or consumption compared to an international poverty level. I was limited by the age of the children in my dataset. A future more robust study should use income/consumption as well as multiple other predictors of poverty to tell a more complete story. In addition, other controls that were not included could likely be biasing my results. Future studies should aim to include controls relating to children's study and work habits and community-level effects.

The additional proxies that I use to support my conclusions regarding the effects of parental education on child poverty outcomes are useful by limited. An additional factor that contributes to child poverty teenage pregnancy. The survey data contained some information regarding whether girls became pregnant and at what age, but this data was significantly limited. Of course, there are disparities between gender when looking at education and poverty rates, especially in developing countries such as Peru. Evaluating the specific factors that contribute to these disparities would have strengthened my study and helped to inform further policy implications.

References

Aldaz-Carroll, E. & Moran, R. (2001). ESCAPING THE POVERTY TRAP IN LATIN AMERICA: THE ROLE OF FAMILY FACTORS. *Cuadernos de economía*, 38(114). https://doi.org/10.4067/S0717-68212001011400003

Arceo-Gomez, E. O. & Campos-Vazquez, R. M. (2014). Teenage Pregnancy in Mexico: Evolution and consequences. *Latin American Journal of Economics*, 51(1), 109–146. https://www.scielo.cl/pdf/laje/v51n1/art04.pdf

Becker, G. S. & Lewis, H. G. (1973). On the Interaction between the Quantity and Quality of Children. *Journal of Political Economy*, 81(2), S279–S288. http://www.jstor.org/stable/1840425. Publisher: University of Chicago Press

Behrman, J. R. & Rosenzweig, M. R. (2002). Does Increasing Women's Schooling Raise the Schooling of the Next Generation? *The American Economic Review*, 92(1), 323–334. http://www.jstor.org/stable/3083336. Publisher: American Economic Association

Birdsall, N. (1991). Birth order effects and time allocation. *Research in population economics*, 7, 191–213.

Chetty, R., Friedman, J. N., Hendren, N., Jones, M. R., & Porter, S. R. (2018). The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility. *National Bureau of Economic Research Working Paper Series*, No. 25147. https://doi.org/10.3386/w25147. featured in NBER digest on 2019-01-02

Chevalier, A. (2004). Parental Education and Child's Education: A Natural Experiment. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.553922

Cueto, S. & León, J. (2012). Psychometric Characteristics of Cognitive Development and Achievement Instruments in Round 3 of Young Lives. *Young Lives*, 114. https://www.younglives.org.uk/publications/psychometric-characteristics-cognitive-development-and-achievement-instruments-round-3

Currie, J. & Moretti, E. (2002). Mother's Education and the Intergenerational Transmission of Human Capital: Evidence from College Openings and Longitudinal Data. *National Bureau of Economic Research Working Paper Series*, No. 9360. https://doi.org/10.3386/w9360

De Walque, D. (2009). Parental Education and Children's Schooling Outcomes: Evidence from Recomposed Families in Rwanda. *Economic Development and Cultural Change*, 57(4), 723–746. https://doi.org/10.1086/598764

Doepke, M. (2015). GARY BECKER ON THE QUANTITY AND QUALITY OF CHILDREN. *Journal of Demographic Economics*, 81(1), 59–66. https://www.jstor.org/stable/26422362. Publisher: Department of Economics, Universite Catholique de Louvain

Duarte, R., Ferrando-Latorre, S., & Molina, J. A. (2018). How to escape poverty through education?: intergenerational evidence in Spain. *Applied Economics Letters*, 25(9), 624–627. https://doi.org/10.1080/13504851.2017.1352073

Ejrnæs, M. & Pörtner, C. C. (2004). Birth Order and the Intrahousehold Allocation of Time and Education. *The Review of Economics and Statistics*, 86(4), 1008–1019. http://www.jstor.org/stable/40042985. Publisher: The MIT Press

Escobal D'Angelo, J. (2012). *Multidimensional poverty and inequality of opportunity in Peru: taking advantage of the longitudinal dimension of Young Lives*. Young Lives, Oxford Department of International Development, University of Oxford.

Escobal D'Angelo, J., Saavedra Chanduví, J., & Suárez, P. (2005). *Economic shocks and changes in school attendance levels and education expenditure in Peru*. Number no. 13 in Working paper / Young Lives. Young Lives, Save the Children UK. OCLC: ocm59719938.

Glewwe, P., Krutikova, S., & Rolleston, C. (2017). Do Schools Reinforce or Reduce Learning Gaps between Advantaged and Disadvantaged Students? Evidence from Vietnam and Peru. *Economic Development and Cultural Change*, 65(4), 699–739. https://doi.org/10.1086/691993

Kaestner, R. (1998). Does Drug Use Cause Poverty? *National Bureau of Economic Research Working Paper Series*, No. 6406(published as Robert Kaestner. "Does Drug Use Cause Poverty?," in Frank J. Chaloupka, Michael Grossman, Warren K. Bickel and Henry Saffer, editors, "The Economic Analysis of Substance Use and Abuse: An Integration of Econometric and Behavioral Economic Research" University of Chicago Press (1999)). https://doi.org/10.3386/w6406

Oreopoulos, P., Page, M., & Stevens, A. (2006). The Intergenerational Effects of Compulsory Schooling. *Journal of Labor Economics*, 24(4), 729–760. https://doi.org/10.1086/506484

Plug, E. (2004). Estimating the Effect of Mother's Schooling on Children's Schooling Using a Sample of Adoptees. *American Economic Review*, 94(1), 358–368. https://doi.org/10.1257/000282804322970850

Pérez, A., Osman, A., Peña, L., Abad-Vivero, E. N., Hardin, J. W., Sargent, J., Thrasher, J. F., & Mejía, R. (2018). Parent Education and Substance Use among Latin-American Early Adolescents. *Health behavior and policy review*, 5(3), 91–99. https://doi.org/10.14485/HBPR.5.3.9. Place: United States

Sen, A. (1983). Economics and the Family. *Asian Development Review*, 01(02), 14–26. https://doi.org/10.1142/S0116110583000081

Solon, G. (1999). Intergenerational Mobility in the Labor Market. *Handbook of Labor Economics*, volume 3, 1761–1800. Elsevier. https://doi.org/10.1016/S1573-4463(99)03010-2

Tansel, A. (1997). Schooling Attainment, Parental Education, and Gender in Côte d'Ivoire and Ghana. *Economic Development and Cultural Change*, 45(4), 825–856. https://doi.org/10.1086/452309

Thomas, D. (1990). Intra-Household Resource Allocation: An Inferential Approach. *The Journal of Human Resources*, 25(4), 635–664. https://doi.org/10.2307/145670. Publisher: [University of Wisconsin Press, Board of Regents of the University of Wisconsin System]

Appendices

Figure 1

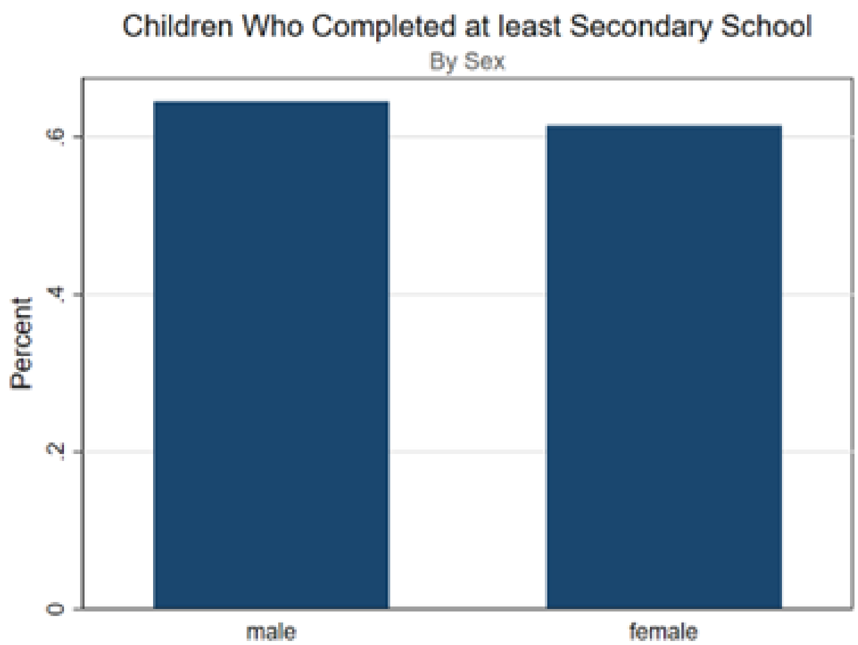

Figure 2

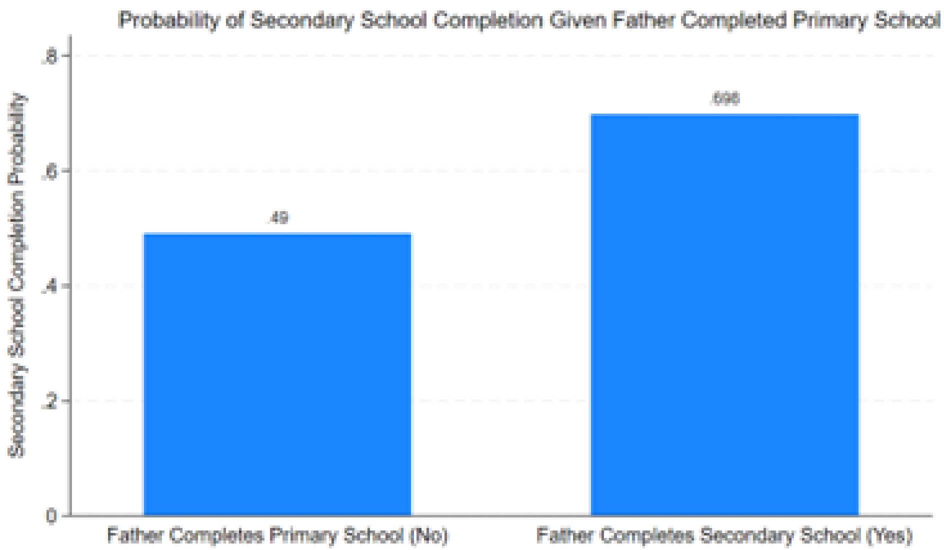

Figure 3

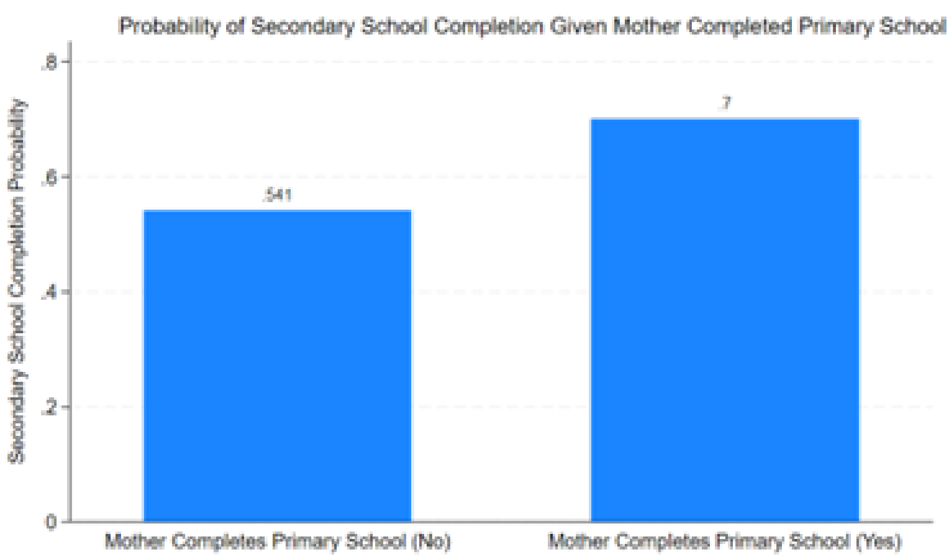

Food for Thought: Understanding the Impact of Food Insecurity on Fertility Dynamics in the U.S.

Sophie Biesterfeld

ECON 381: Introduction to Econometrics (Advisor: Amy Damon)

American abundance isn't just a concept—it's a way of life. For over 100 years, the United States has stood as a world power, boasting ample wealth and resources from its free-market economy. Despite the consumerist culture and over-consumption stereotypes that plague the U.S., over 44 million Americans did not have enough to eat or did not have access to healthy food in 2023 (Rabbitt et al., 2023). How does this paradox between abundance and food insecurity persist within one of the wealthiest nations on earth?

The USDA defines food security as "access by all people at all times to enough food for an active, healthy life." It is one of the most critical public health issues in the United States, hindering the prosperity of its people. Food insecurity in the U.S. and other developed countries poses a unique case: rather than being a direct result of civil conflict, crop failures, or inadequate infrastructure, Americans who are hungry simply don't have enough money to buy food (The Global Giving Team, 2021). In addition to this, about 10% of the approximately 65,000 census tracts in the United States were classified as food deserts, regions where people have limited access to healthy and affordable food (Dutko et al., 2012).

Literature Review

Food insecurity bears many economic and social costs to American society such as productivity losses, education and social service expenses, lost economic potential, and decreases in community welfare. Without question, the facet of American life that faces the greatest burden of food insecurity costs is the healthcare sector. A 2019 study from Berkowitz et al. (2019) estimated that food insecurity costs the U.S. healthcare system an additional $53 billion annually by triggering chronic diseases and fueling emergency room visits, hospitalizations, and readmissions. A more recent study from Palakshappa et al. (2023) found that food-insecure families in the U.S. paid an average of $2,500 more annually in healthcare costs than families with sufficient food.

Sophie Biesterfeld is a junior with a major in Economics and minors in Environmental Studies and Statistics. Correspondence concerning this article should be addressed to sbiester@macalester.edu

Unsurprisingly, the effects of food insecurity are not felt equally across different groups. Racial minorities, low-income people, and women disproportionately experience food insecurity. Even though national food insecurity rates fell substantially between 2010 and 2021, a recent study from Gundersen (2023) found that these more vulnerable groups saw smaller declines in food insecurity rates and were "left behind." This is backed by a study from Flores-Lagunes et al. (2018) which found that Black households are more likely than white households to experience a greater intensity of food insecurity. Additionally, women living alone and single-mother households with children are more likely to be severely food insecure than nuclear family households (USDA, 2019). Women also experience exacerbated adverse health effects related to food insecurity. Their poor health can impact fertility, creating long-term repercussions on population dynamics and demographic trends. A study using data from California found that nutrition during pregnancy not only influences the current health condition of women and infants, but also plays an important role in the health condition of children and adults in the future (Braveman et al., 2010).

There is an extensive body of research to support the correlation between food insecurity and negative health outcomes for women. For this reason, I wanted to explore the relationship between food insecurity and fertility in the United States. In 2006, a study by Cook et al. (2006) found that U.S. households suffering from food insecurity are more likely to have children with poorer general health. One year later, Carmichael et al. (2007) found that food insecurity during pregnancy was associated with higher risks of birth defects in California. Analysis from Grilo et al. (2015) echoed these sentiments, stating that poor nutrition increased the risk for pregnancy complications such as gestational diabetes, preeclampsia, and fetal macrosomia, leading to worse birth outcomes including shorter gestations and lower birth weights.

While many researchers have examined the impact of food insecurity on pregnant women, few have looked at the possibility of food insecurity inducing fertility outcomes. One of the few studies that investigates a causal relationship is from DiClemente et al. (2021) which analyzed fertility preferences among Tanzanian women during times of food insecurity. The results indicated that women who experience household hunger had a preference to delay or avoid pregnancy. An adjacent study by Alam and Pörtner (2018) found that the likelihood of pregnancies and childbirth was significantly lower and contraception use was significantly higher for Tanzanian households that experienced crop losses. Based on these results, I was curious how this relationship might fare in a drastically different part of the world.

Despite strong evidence that links food insecurity and fertility, there are various other interdependent influences. In 2020, Bijlsma and Wilson (2020) note that fertility rates in the UK are influenced by many external factors such as education, employment, and marriage. There are also more abstract influences on fertility that are harder to measure. Barrett et al. (2020) found that the desired family size of a household in Sub-Saharan Africa is positively related to the average family size in the community. This is evidence that fertility can not only be influenced by private desires, but societal mores as well. There is also a strong and pertinent income effect in this relationship. A 2016 study by Murthy (2016) found that low-income families are more likely to be food insecure, but they are also more likely to postpone medical care or underuse medicine because of budget constraints. Because food insecurity and income are so closely correlated, it will be hard to differentiate which fertility effects are due to income and which are due to food insecurity. In my regression analysis, I will do my best to control for these factors to cipher out the specific effect of food insecurity on fertility.

Economics of the Problem

The health literature on food and fertility is fairly straightforward. Food security improves women's overall health and nutrition, making them better equipped to sustain a pregnancy to term. If babies are carried to term when a mother is unhealthy, this can lead to negative health outcomes for the child. Introducing economic theory entangles the conversation

The baseline, neoclassical theory views fertility as a result of individual choice-making. Women will weigh the costs and benefits of having children, and continue to have them until the marginal benefit equals the marginal cost. In this theory, a high demand and low supply of food will induce women to stop having children they cannot feed, so long as they can control it. It will also prompt households to allocate scarce resources away from child-rearing activities, such as healthcare and education, and towards obtaining food. Additionally, because of the perceived impacts of giving birth while in poor health, this may also restrict women from choosing to have children while food insecure.

On the flip side, human capital theory suggests that fertility decisions are influenced by the trade-off between the quality and quantity of children. Especially in high-income families, parents may opt for fewer children, enabling them to invest more time and money per child. This will potentially increase each child's human capital and future earning potential. As parents see these returns on investment increase, they will continue to invest in their existing children, rather than having more. There are also strong income and education effects in this relationship. As one's income grows, the income effect predicts that people will begin to demand and consume more. This close link between income and food results in both variables having a similar impact on fertility choices. In the same vein, women with more education have a higher opportunity cost of bearing children in terms of lost income. In two-parent households, more educated women are better able to support themselves and have more bargaining power over family size, which is validated by the household bargaining model (Echevarria and Merlo, 1999). These women also have more robust sex education surrounding pregnancy prevention and pregnancy risks. This link between education and fertility results in both variables having a similar relationship with food insecurity.

While there are competing economic theories on the relationship between food insecurity and fertility, I believe the latter theory is supported by the U.S. context due to the robust social safety nets that counteract the neoclassical theory. Women experiencing food insecurity and economic instability may choose to have more children as a form of social security or insurance against economic shocks. Social programs such as the Child Tax Credit and the Earned Income Tax Credit provide financial assistance to low and middle-income families with children. There are also programs like the Supplemental Nutrition Assistance Program (SNAP) and the Special Supplemental Nutrition Program for Women, Infants, and Children (WIC) that offer nutritional support for pregnant women, new mothers, and young children. Due to the existence of these programs, as well as human capital theory and education/income effects, I hypothesize a positive relationship, meaning U.S. counties with lower rates of food insecurity will induce lower rates of fertility.

Data Description

I am examining panel data with 6,126 observations of women ages 15-60 across U.S. counties from 2010-2017 (NPS, 2023). The data for my independent variable of interest, food insecurity rate (FIR), comes from the non-profit organization Feeding America through their annual "Map the Meal Gap" study. Feeding America provides FIRs for every U.S. county, which they estimate based on household questionnaires, median income, unemployment, homeownership, and disability prevalence. The data for one of the instrumental variables, "cost per meal", also comes from Feeding America. This is calculated by adjusting the national average cost per meal by a relative food cost index to derive a county-level estimate.

The data for my dependent variable of interest, fertility, comes from the American Community Survey (ACS), conducted by the United States Census Bureau. All estimates are at the county-level and look at the 15-50 year old age range among women. The ACS provided one-year estimates for the total female population, as well as the female population who gave birth. With these values, I generated the fertility outcome variable: the percentage of women who gave birth in the past year. I used the same generation method to find various demographic percentages among these women such as the percentage of women who are married, foreign-born, employed, living below 100% of the poverty level, and among seven different age groups. These were all used as controls in my regression. Data for the other two instrumental variables also comes from the ACS, which looks at the number of cars per household. Per the US Census, all this data is collected in survey form from a random sample of addresses in each state.

When checking for robustness in my results, I also controlled for the number of abortions per county. This data comes from the Guttmacher Institute which provides abortion estimates for each US state. I disaggregated these state-level values by generating proportions based on similar related county-factors such as population and fertility.

Table 1

Descriptive Statistics

Variable	Obs.	Mean	Std. Dev.	Min.	Max.
FIR	6126	13.99	3.628	3.4	28.3
fertility	6126	5.446	1.524	0.85	16.002
poverty	6126	18.039	7.023	1.639	54.439
married	6126	46.536	6.589	21.771	68.387
labor	6126	72.459	6.103	41.513	99.395
foreign	6126	10.993	9.025	0.285	55.284

The FIR exhibits a sample mean of approximately 14% with a standard deviation of roughly 3.6, indicating notable variability in FIRs across counties. The range of FIR spans from a minimum of 3.4 to a maximum of 28.3, highlighting the diversity in outcomes observed within the dataset. Similarly, the poverty rate displays a sample mean of about 18% with a standard deviation of 7.023, reflecting substantial variation in poverty levels among counties. The fertility rate exhibits a mean of 5.4% with a standard deviation of about 1.5 indicating minimal variability in fertility across counties. These descriptive statistics provide valuable insights into the distribution of my variables of interest, laying the groundwork for subsequent empirical analyses.

Table 2

FIR Sample Average Comparison

	Below Average	Above Average	Total
N	3,178 (51.9%)	2,948 (48.1%)	6,126 (100.0%)
fertility	5.304 (1.423)	5.598 (1.611)	5.446 (1.524)
poverty	14.259 (5.478)	22.114 (6.176)	18.039 (7.023)
married	48.501 (5.578)	44.417 (6.930)	46.536 (6.589)
labor	73.972 (5.872)	70.829 (5.926)	72.459 (6.103)
foreign	11.498 (9.481)	10.449 (8.475)	10.993 (9.025)

Table 2 splits the counties into two groups: above and below the sample average FIR of roughly 14%. The below-average group holds about 52% of the sample observations while the above average holds about 48%. This symmetry suggests there are not large outliers on either end that significantly drag the average up or down. The values for the subsequent variables indicate the mean and standard deviation for the given group. Across the board, it is evident that the cluster of counties below the FIR average has a lower average fertility and poverty rate with higher rates of marriage and labor force participation. This relationship is consistent with the theory of income effect, since income and fertility have a positive correlation and food insecurity is closely correlated with income.

Table 3

Poverty Rate Sample Average Comparisons

	Below Average	Above Average	Total
N	3,199 (52.2%)	2,927 (47.8.%)	6,126 (100.0%)
FIR	12.008 (2.809)	16.156 (3.154)	13.990 (3.628)
fertility	5.350 (1.441)	5.550 (1.603)	5.446 (1.524)
married	49.237 (5.624)	43.584 (6.294)	46.536 (6.589)
labor	73.788 (5.324)	71.007 (6.554)	72.459 (6.103)
foreign	11.690 (9.232)	10.232 (8.733)	10.993 (9.025)

Table 3 uses the same method but instead splits the groups by the sample average poverty rate of about 18%. Displayed is the mean of each group with the standard deviation in parenthesis. Similarly to Table 2 the two groups have a fairly even split of observations, suggesting there are no large outliers related to poverty. On average, the counties clustered below the poverty rate average experience lower rates of fertility and food insecurity with higher rates of marriage and labor force participation. Given the income effect, the raw data statistics are unsurprising.

Figure 1

FIR vs. Fertility

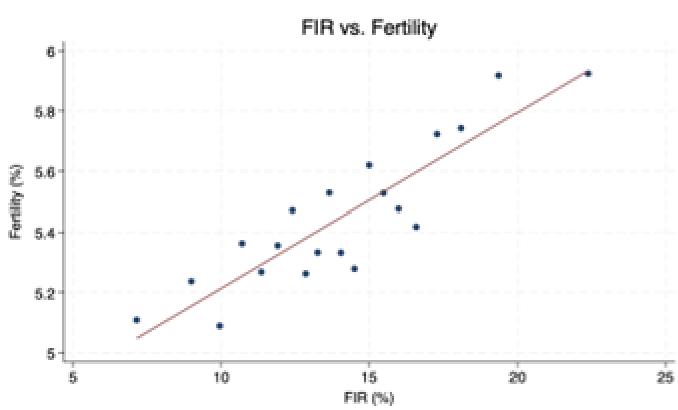

Figure 2

Poverty vs. Fertility

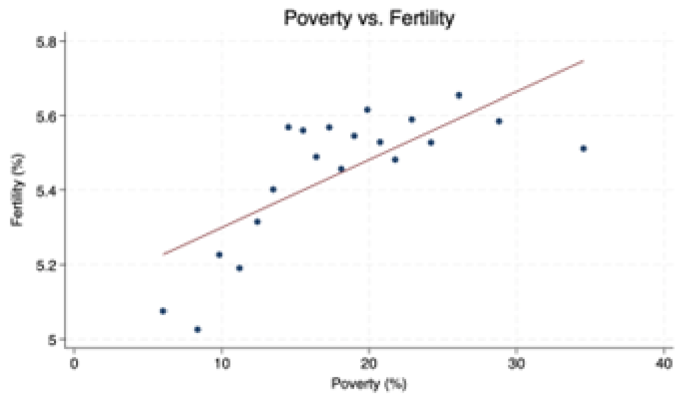

Figure 1 visualizes the trend in the raw data between food insecurity and fertility. In order to reduce clutter and better conceptualize the relationship, I created a bin scatterplot. This method groups the counties into equal-sized bins based on the x-axis variable, using the mean in each bin to create a scatterplot of data points. Figure 1 reveals a distinct positive linear trend between food insecurity and fertility. A similar relationship is also evident in Figure 2 which details poverty and fertility. Both figures further solidify the income effect theory.

Figure 3

Histogram and density plot of FIRs for every year

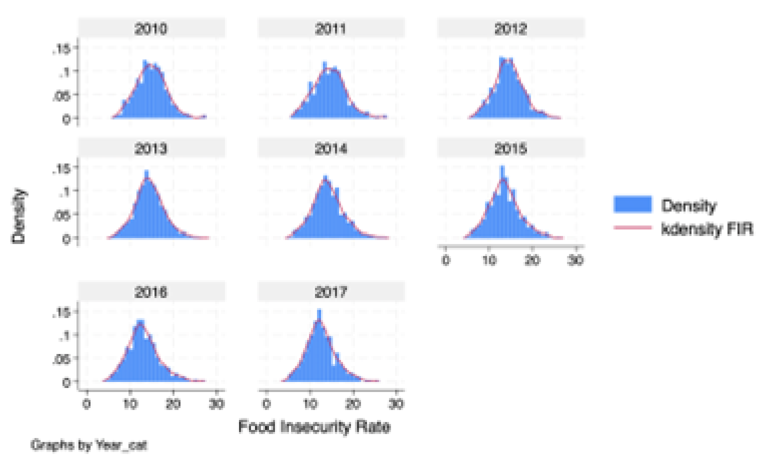

Figure 4

Histogram and density plot of Fertility for every year

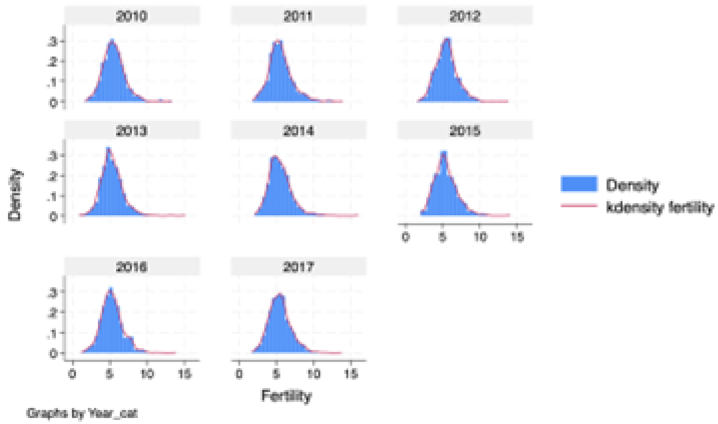

Figure 3, pictured above, combines a histogram and density plot of FIRs for every year. FIRs for each year are normally distributed with means of around 12-14%. This signifies that no single year is drastically skewing the FIR total sample mean. Figure 4 exhibits the same visualizations for fertility across each year. Fertility rates for each year are slightly skewed to the right with sizable tails. This suggests that every year, there are consistently a few outlier counties that have very high fertility rates of around 10-15%. Both figures demonstrate that there is uniformity across the years of interest when it comes to my independent and dependent variables. These observed patterns are less likely to be spurious or driven solely by anomalies in specific years. Having this consistency across time for both variables will strengthen my empirical exploration of a causal relationship.

Empirical Strategy

My regression will attempt to discern a causal relationship between food insecurity and fertility. I am using an instrumental variable (IV) regression with the two-stage least squares technique. The instruments are cost per meal in USD, number of households that do not own a vehicle, and number of households that own one vehicle. Stage one will use my instruments to predict the independent variable. This predicted value is then used in stage two analysis instead of the original values.

I am employing two-way fixed effects in hopes of curbing the impact of unobserved confounding variables that are constant across counties and time periods. This will be especially helpful to parse out the noise regarding how sentiments towards women and childbirth differ across time and communities. Rather than using robust standard errors, I am using VCE clustering which gathers my data into 823 clusters by county. Since the within-county observations are likely to be correlated, this method will allow for a more efficient estimation of standard errors by explicitly modeling the covariance structure within clusters.

First Stage IV

$$\begin{aligned} FIR_{it} = \pi_0 &+ \pi_1 CostPerMeal_{it} + \pi_2 noVehicle_{it} + \pi_3 oneVehicle_{it} + \pi_4 (LMD * FIR)_{it} \\ &+ \pi_5 Poverty_{it} + \pi_6 Married_{it} + \pi_7 Labor_{it} + \pi_8 Foreign_{it} + \pi_9 Age15to19_{it} \\ &+ \pi_{10} Age20to24 + \pi_{11} Age25to29_{it} + \pi_{12} Age30to34 + \pi_{13} Age35to39 \\ &+ \pi_{14} Age40to44_{it} + \pi_{15} Age45to50_i t + \alpha_i + \gamma_t + \mu_{it} \end{aligned} \tag{1}$$

The three instruments will be used to predict food insecurity values that will then be regressed with fertility. Cost per meal acts as a conduit to food prices and inflation over the observed years. It is directly related to the affordability of food, as higher costs will increase food insecurity. The issue here is that there is not much variation in meal prices across counties, making it hard to discern a statistically significant relationship between cost per meal and food insecurity, despite its relevance.

Vehicle and transportation access is another variable that is commonly used to evaluate food insecurity in economic research (Ploeg et al., 2015). Affordable and nutritious food options may be located in areas that are further away or less accessible by other modes of transportation. Despite their strong relevance, they are ultimately weak instruments due to the existing pathways of connection to fertility. Therefore, I cannot say for certain that vehicle ownership is uncorrelated with fertility. These shortcomings will be discussed later in the paper. While I have established these three instruments are weak in their own regard, I will still implement them in hopes of mitigating some endogeneity and improving the bias of my estimators.

Second Stage IV

$$\begin{aligned} Fertility_{it} = \beta_0 &+ \hat{\beta_1} FIR_{it} + \beta_2 (LMD * FIR)_{it} + \beta_3 Poverty_{it} + \beta_4 Married_{it} + \beta_5 Labor_{it} \\ &+ \beta_6 Foreign_{it} + \beta_7 Age15to19_{it} + \beta_8 Age20to24_{it} + \beta_9 Age25to29_{it} + \beta_{10} Age30to34_{it} \\ &+ \beta_{11} Age35to39_{it} + \beta_{12} Age40to44_{it} + \beta_{13} Age45to50_{it} \eta_i + \theta_t + \mu_{it} \end{aligned} \tag{2}$$

My final regression is intended to obtain consistent evidence of a causal relationship between my variables of interest. Ideally, this will help me draw reliable conclusions related to public health and the American economy.

Results

Table 4

OLS Regression Results

VARIABLES	*(1) No Lag* fertility	*(2) With Lag* fertility
L.FIR		0.0659**
		(0.0332)
FIR	0.0376*	0.000201
	(0.0221)	(0.0367)
Constant	4.149*	4.331
(2.223)	(2.929)	
Observations	6,125	5,158
R-squared	0.372	0.393

Notes: *, **, and *** indicate significance at the 10, 5, and 1 percent levels. Controls and TWFE added.

The first OLS regression finds weak but suggestive evidence of a positive relationship between FIR and fertility. It estimates that a 10 percentage point increase in food insecurity rate is associated with a roughly 0.38 percentage point increase in fertility rate. Using the sample average as an example, this is reflected as a fertility rate increase from 5% to 5.38%. However, because fertility choices are made roughly nine months before birth, I suspect we might see a stronger effect using a one-year lag on food insecurity rates. This means that food insecurity in 2010 is actually impacting fertility in 2011. In the second OLS regression, I lagged the independent variable and found an increase in its effect on the dependent variable. Due to the unbalanced nature of the panel, approximately 1,000 counties were dropped. Now, a 10 percentage point increase in food insecurity is associated with a roughly 0.7 percentage point increase in fertility rate. These results are more trustworthy, with an increased significance of 5%. The lagged model also exhibits a slightly higher R-squared compared to the model without a lag. While these values are less important for economic interpretation and determining causality, this increase could indicate that the lag improves the overall explanatory power of the model.

There is still an acute concern about the feedback loop between the independent and dependent variables. While varying levels of food insecurity can influence reproductive decisions, fertility rates can also strain available resources, which may lead women to delay or limit childbearing. This endogeneity makes it challenging to establish a clear causal relationship between the two variables. In an attempt to battle some of this endogeneity, I ran an instrumental variable regression.

Table 5

IV Regression Results

VARIABLES	*(1) No Lag* fertility	*(2) With Lag* fertility
L.FIR		-0.906
		(0.886)
FIR	0.415*	0.804
	(0.238)	(0.732)
Constant	0.419	4.640
(3.188)	(2.890)	
Observations	6,125	5,158
Number of FIPS	823	798

Notes: *, **, and *** indicate significance at the 10, 5, and 1 percent levels. Controls and TWFE added.

The addition of instruments induced a substantial change in coefficient estimates for regressions 3 and 4. No-vehicle households and one-vehicle households were significant at the 1% level, while cost per meal had a statistically insignificant coefficient. Regressions 3 and 4 have F statistics of about 261 and 280 respectively, far above the threshold for a strong IV (see Appendix A). The magnitude of the F stat indicates greater overall explanatory power of the instruments and provides evidence against the null hypothesis. Knowing that the instruments performed moderately well and substantially impacted coefficient estimates, I assert that these instruments untangled some of the biased estimates and unreliable statistical inferences that were present in my OLS regressions. However, in the IV process, the lagged model lost all significance. Regression 3 finds that a 10 percentage point increase in food insecurity is associated with a 4.15 percentage point increase in fertility rate.

Despite weak but suggestive evidence of a positive relationship, I was still doubting these results. Due to the substantial variation in governmental policies and regulations enacted at the state and local levels, I have trouble believing that a relationship holds uniformly across the entire country. To further investigate this, I looked into the region-specific effects of food insecurity by creating a dummy interaction variable for the Lower Mississippi Delta (LMD) region. This area consists of various counties surrounding the Mississippi River in the southern United States (National Park Service, 2024). It is also home to some of the strictest abortion laws and the highest food insecurity rates in the nation. I was curious to see how this combination of socioeconomic factors would influence the relationship between my variables of interest.

Table 6

IV Regression with LMD Dummy Interaction

VARIABLES	(5) fertility
FIR	0.416*
	(0.237)
LMD	-0.411*
	(0.220)
Constant	0.592
	(3.123)
Observations	6,125
Number of FIPS	823

Notes: *, **, and *** indicate significance at the 10, 5, and 1 percent levels. Controls and TWFE added.

Table 6 details the final regression results which continue to find weak but suggestive evidence of a positive relationship between FIR and fertility. The FIR coefficient effectively stays the same, finding that a 1 percentage point increase in food insecurity is associated with a 0.411 percentage point increase in fertility rate. These results may seem negligible, but they have significant real-world impact. In 2017, the average population of women ages 15-50 in a county was about 75,000, excluding extreme outliers. Assuming the average fertility rate of about 5%, the results of this regression indicate an additional 308 women giving birth in a year from just a 1 percentage point increase in FIR.

The addition of the LMD dummy interaction extracts the relationship between FIR and fertility in this specific region. Both coefficients are similar in magnitude, but the LMD variable indicates a negative relationship, with fertility decreasing by 0.408 percentage points for every 1 percentage point increase in FIR. These results paint an interesting picture of regional differences across the United States. The coefficient sign flip makes it clear that not all regions operate with the same positive relationship between FIR and fertility that is seen at the national level.

While these results technically hold significance, there is still a 10% chance that the observed relationship between FIR and fertility is due to random sampling variability rather than a true relationship in the population. The 10% significance level is not very stringent, but it still suggests that the observed relationship is likely to be meaningful. In order to check for robustness in my results, I implemented a few tests. I began by logging three control variables and my results came out effectively the same (see Appendix Table B1). I also added a control for abortions per county which also did not largely impact my results (see Appendix Table B2).

Limitations

Despite weakly significant results, a myriad of limitations prevent these results from being internally valid. As previously noted, all three instruments are correlated with the independent variable of interest, making them relevant candidates for an IV regression. However, the cost per meal instrument was ultimately statistically insignificant, despite boasting stronger evidence of exogeneity. On the other hand, the two instruments for vehicle access both had statistical significance, but still violated the exclusion restriction assumption due to the pathways between vehicle ownership and fertility. To put this into context, owning a car can expand one's access to birth control and maternal care and therefore influence fertility decisions. This direct relationship makes it impossible to discern the causal effect of the endogenous variable because the instrument is capturing variation in fertility that is not due to FIR. Due to the weakness of the instruments, these results are still plagued with bias as a result of existing reverse causality.

In addition to endogeneity, there is a sizable amount of omitted variable bias, which occurs when influential variables are left out of the analysis. A few examples of omitted variables that also impact fertility are educational attainment, religious beliefs, abortion clinics, and other health factors. Failing to account for these variables limits the interpretation of my results as casual. In order to improve the robustness, future research should aim to incorporate a more comprehensive set of control variables.

The residual vs. fitted plot of my final regression revealed some interesting findings (see Appendix C). The residual trend is downward sloping, with data points clustered along the line. Based on this trend in residuals, the conditional mean of the error term given the independent variable is not zero, violating the first Gauss-Markov condition. Ideally, the error term should capture random fluctuations or noise that cannot be explained by the independent variables included in the model. However, this residual trend indicates that there is omitted variable bias in the model. Another plausible explanation is that once I controlled for other factors, it revealed that the true relationship between food insecurity and fertility is non-linear.

It is also imperative to acknowledge the potential existence of sampling error that comes with using survey data. While the ACS is employed nationwide, it does not mandate a response. Due to the sensitivity surrounding health, some women may not feel comfortable reporting on their fertility. The direction in which this bias may affect my results is ambiguous, but the bias itself remains present. Additionally, there is a threat of estimation error in the data from Feeding America. As per their technical report, data was first analyzed at the state level before these coefficient estimates were used with the same variables for every county. Due to the nature of empirical work, the measurement of these food insecurity estimates may be imprecise. The Technical Report also encourages users to exercise caution when comparing estimates over time, especially when differences are small since they may not be statistically different. That being said, the magnitude of those changes may be relatively large and potentially meaningful.

Conclusion

In conclusion, these findings reveal a weak but suggestive positive relationship between food insecurity and fertility at the national level. This is consistent with my hypothesis which draws on health literature and economic theories of human capital and income effects. However, the inclusion of the LMD dummy variable uncovers regional disparities. More importantly, it begins to highlight the significance of socioeconomic factors in shaping reproductive behaviors. Further research should look to expand upon these region-specific effects.

This paper attempts to contribute to the ongoing discourse surrounding public health and social welfare policies. My research underscores the importance of addressing food insecurity as a critical public health issue and women's issue. Policies aimed at reducing food insecurity not only alleviate economic hardship, but also have the potential to influence reproductive health outcomes among women. It is also critical to introduce the idea of race into this study, as food and healthcare access varies greatly among racial groups. By further exploring these dynamics, research like this can help to inform evidence-based policies that will improve the well-being of American women and their communities.

References

Alam, S. A. and Pörtner, C. C. (2018). Income shocks, contraceptive use, and timing of fertility. *Journal of Development Economics*, 131:96–103.

Barrett, S., Dasgupta, A., Dasgupta, P., Adger, W. N., Anderies, J., van den Bergh, J., Bledsoe, C., Bongaarts, J., Carpenter, S., Chapin, F. S., Crépin, A.-S., Daily, G., Ehrlich, P., Folke, C., Kautsky, N., Lambin, E. F., Levin, S. A., Mäler, K.-G., Naylor, R., Nyborg, K., Polasky, S., Scheffer, M., Shogren, J., Jørgensen, P. S., Walker, B., and Wilen, J. (2020). Social dimensions of fertility behavior and consumption patterns in the Anthropocene. *Proceedings of the National Academy of Sciences of the United States of America*, 117(12):6300–6307. Publisher: National Academy of Sciences.

Berkowitz, S. A., Basu, S., Gundersen, C., and Seligman, H. K. (2019). State-Level and County-Level Estimates of Health Care Costs Associated with Food Insecurity. *Preventing Chronic Disease*, 16:180549.

Bijlsma, M. J. and Wilson, B. (2020). Modelling the Socio-Economic Determinants of Fertility: A Mediation Analysis using the Parametric g-formula. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 183(2):493–513.

Braveman, P., Marchi, K., Egerter, S., Kim, S., Metzler, M., Stancil, T., and Libet, M. (2010). Poverty, Near-Poverty, and Hardship Around the Time of Pregnancy. *Maternal and Child Health Journal*, 14(1):20–35.

Carmichael, S. L., Yang, W., Herring, A., Abrams, B., and Shaw, G. M. (2007). Maternal Food Insecurity Is Associated with Increased Risk of Certain Birth Defects ,. *The Journal of Nutrition*, 137(9):2087–2092.

Cook, J. T., Frank, D. A., Levenson, S. M., Neault, N. B., Heeren, T. C., Black, M. M., Berkowitz, C., Casey, P. H., Meyers, A. F., Cutts, D. B., and Chilton, M. (2006). Child Food Insecurity Increases Risks Posed by Household Food Insecurity to Young Children's Health. *The Journal of Nutrition*, 136(4):1073–1076.

DiClemente, K., Grace, K., Kershaw, T., Bosco, E., and Humphries, D. (2021). Investigating the Relationship between Food Insecurity and Fertility Preferences in Tanzania. *Maternal and Child Health Journal*, 25(2):302–310.

Dutko, P., Ver Ploeg, M., and Farrigan, T. (2012). Characteristics and Influential Factors of Food Deserts. Economic Research Report 144, United States Department of Agriculture.

Echevarria, C. and Merlo, A. (1999). Gender Differences in Education in a Dynamic Household Bargaining Model. *International Economic Review*, 40(2):265–286.

Flores-Lagunes, A., Jales, H. B., Liu, J., and Wilson, N. L. (2018). The Differential Incidence and Severity of Food Insecurity by Racial, Ethnic, and Immigrant Groups over the Great Recession in the United States. *AEA Papers and Proceedings*, 108:379–383.

Grilo, S. A., Earnshaw, V. A., Lewis, J. B., Stasko, E. C., Magriples, U., Tobin, J., and Ickovics, J. R. (2015). Food Matters: Food Insecurity among Pregnant Adolescents and Infant Birth Outcomes. *The Journal of Applied Research on Children: Informing Policy for Children at Risk*, 6(2):4.

Gundersen, C. (2023). A reconsideration of food insecurity trends in the United States. *Applied Economic Perspectives and Policy*, page aepp.13412.

Murthy, V. H. (2016). Food Insecurity: A Public Health Issue. *Public Health Reports (1974-)*, 131(5):655–657. Publisher: Sage Publications, Inc.

NPS (2023). Counties within LMDI - Lower Mississippi Delta Region (U.S. National Park Service).

Palakshappa, D., Garg, A., Peltz, A., Wong, C. A., Cholera, R., and Berkowitz, S. A. (2023). Food Insecurity Was Associated With Greater Family Health Care Expenditures In The US, 2016–17: Study examines the relationship between food insecurity and family health care expenditures. *Health Affairs*, 42(1):44–52.

Ploeg, M. V., Dutko, P., and Breneman, V. (2015). Measuring Food Access and Food Deserts for Policy Purposes. *Applied Economic Perspectives and Policy*, 37(2):205–225. Publisher: [Oxford University Press, Agricultural & Applied Economics Association].

Rabbitt, M. P., Hales, L. J., Burke, M. P., Coleman-Jensen, A., and United States. Department of Agriculture. Economic Research Service (2023). Household food security in the United States in 2022. Technical report, Economic Research Service, U.S. Department of Agriculture, Washington, D.C.

The Global Giving Team (2021). How Much Would It Cost To End World Hunger? Get The Facts.

USDA (2019). USDA ERS - Food Security and Nutrition Assistance.

Appendix A

Table A. F Tests

Regression 3	Regression 4	Regression 5
F(22, 822) = 227.63	F(21, 797) = 280.42	F (22, 822) = 227.63
Prob <F = 0.0000	Prob <F = 0.0000	Prob <F = 0.0000

Appendix B: Robustness Checks

Table B1. Regression with logged variables

VARIABLES	(1) fertility
FIR	0.399*
	(0.239)
LMD	-0.394*
	(0.220)
Constant	-5.990
	(3.575)
Observations	6,125
Number of FIPS	823

Table B2. Regression with abortion control

VARIABLES	(1) fertility
FIR	0.404*
	(0.238)
LMD	-0.406*
	(0.219)
Constant	0.440
	(3.099)
Observations	6,125
Number of FIPS	823

Appendix C

Figure C. RVF Plot

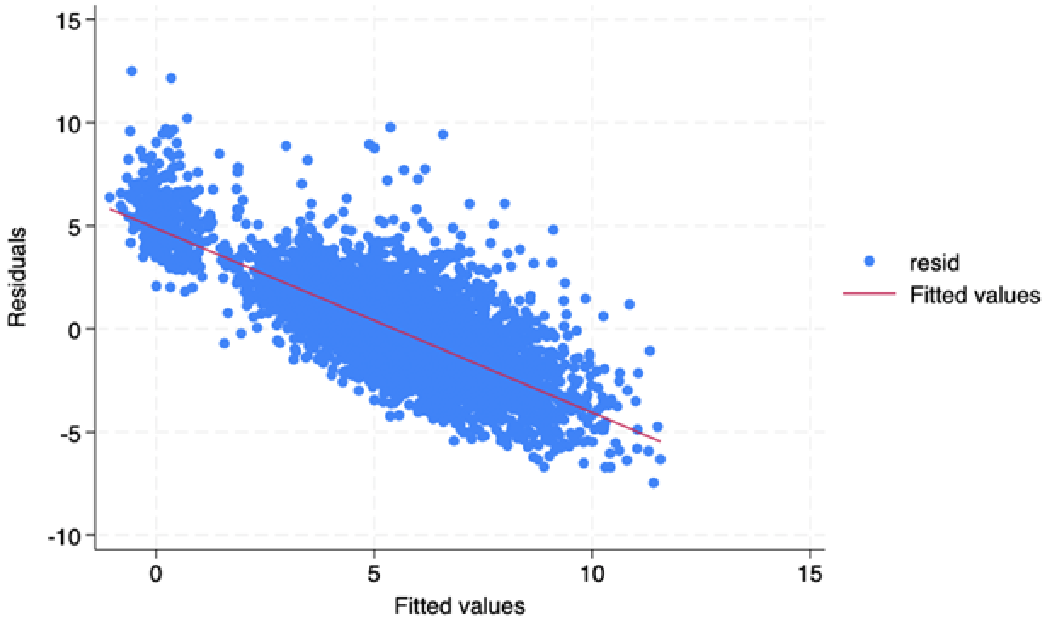

Infrastructure-Based Climate Resilience in Bangladesh

Ella DeMay
ECON 426: International Economic Development (Advisor: Amy Damon)

The scientific basis for human-caused climate change is well established and the consequences of this crisis are already being borne (Lee and Romero, 2023). Despite the fact that most climate change-causing greenhouse gas emissions originate from the Northern Hemisphere, the most immediate and drastic climate-induced calamities are frequently located in the Southern Hemisphere (Lee and Romero, 2023). Approximately one-third of Bangladesh experiences extreme flooding once in every ten years, and this frequency is expected to increase as global climate change escalates (Haque et al., 2022; Lee and Romero, 2023). Further, the negative impact of these calamities is centralized on the most vulnerable groups based on gender, age, health, and social status (Chisty et al., 2022). Due to its location, geography, and population's socioeconomic status, Bangladesh finds itself exceptionally susceptible to climate change risk (Bandyopadhyay and Skoufias, 2015; MEF, 2009). Human-driven climate change has already impacted Bangladesh, and regardless of the efforts to mitigate future change, the effects will worsen in this region for years to come (Haque et al., 2022; Lee and Romero, 2023; MEF, 2009).

The growing threat of climate change demands the urgent implementation of proven resilience strategies. Infrastructure damage caused by extreme weather events can be a serious setback for households and communities everywhere, especially those without durable infrastructure. A household's home is an important investment and vital tool for a group's well-being, capital accumulation, and development outlook (Karim and Noy, 2020). Climate shocks and consequent home repair costs set back households in terms of overall income, accumulated savings, and health and education spending (Karim and Noy, 2016). These setbacks can reduce wealth accumulation, loan eligibility, and vital education and health outcomes long-term (Karim and Noy, 2020, 2016). In the following paper, the effect of durable home building materials is analyzed through the examination of non-food spending and aggregate education and health spending fluctuations before and after extreme weather shocks. Literature states that spending on food will be relatively smooth even after a shock,

Ella DeMay is a senior with majors in Economics and Environmental Studies. Correspondence concerning this article should be addressed to edemay@macalester.edu.

while non-food spending follows a downward trend, most prominently due to health and education spending cuts (Karim and Noy, 2016). Health and education spending influence educational attainment and individual well-being down the line, and thus, variation in non-food spending will provide an indication of human impacts. With a "weak" housing material, non-food spending is expected to decrease relatively more than a "strong" material.

Literature Review

Bangladesh will experience more volatile weather and a higher frequency of extreme events in the coming years, and each of these events will impact human life, well-being, and national development. These impacts have been examined by researchers since the early 1970s, and these analyses have increased in frequency over the last five years. In order to discuss the effectiveness and resilience of various house-building materials, it is important to discuss the existing research on Bangladesh's climate adaptation efforts. Within this research, a variety of methods, indices, and weather shocks are used to begin to paint a picture of cyclones, floods, and storm surges and the consequent impacts they generate.

There are far-reaching consequences of severe weather events which occur in both short- and long-term time frames following a shock. Presently, a substantial proportion of homes in Bangladesh are made of mud, sticks, and grasses (*kutcha*), and suffer significant damage or complete demolition following severe weather events (Haque et al., 2022). The short run effects of severe weather events are incredibly impactful on vital infrastructure in Bangladesh. Beyond the immediate effects, Eskander and Barbier 2022 use a difference-in-difference design to investigate long-term effects of the 1970 Bhola cyclone in Bangladesh using geographical exposure and adulthood health outcomes. They conclude that cohorts from the most severely impacted regions face significant long-term health and education complications, consistent with the proposed hypothesis and previous literature (Eskander and Barbier, 2022). Further, this investigation establishes the positive influence that remedial infrastructure development programs and publicly funded aid can have. Following the establishment of the short- and long-term impacts of natural disasters, the importance of resilient infrastructure is clear. Gaining knowledge of the most resilient infrastructure building practices is thus imperative to supporting a healthy and empowered population in Bangladesh.

The Bangladeshi government, known as the People's Republic of Bangladesh (PRB), has acknowledged the climate threat and been pragmatic in its approach to confront it. This is likely due to the frequency of severe weather events, such as the 1970 Bhola cyclone, as well as the clear negative relationship of severe weather events to development and human well-being (Karim and Noy, 2020; Eskander and Barbier, 2022). In 2005, the PRB established the National Adaptation Programme of Action (NAPA) which aims to increase resilience in agriculture, infrastructure, and water systems as well as improve livelihood prospects, gender equality, and policy across the nation (MEF, 2005).

Within this research, the definition of resilience provided by United Nations Office for Disaster Risk Reduction (UNDRR), states: *"Resilience is defined as the ability of a system, community or society exposed to hazards to resist, absorb, accommodate, adapt to, transform and recover from the effects of the hazard in a timely and efficient manner, including the preservation and restoration of its essential basic structures and functions through risk management"* (Wannous and Velasquez, 2017). NAPA, with its mission of adaptation and resilience, was created in conjunction with the Green Climate Fund (GCF) who has since provided over 440 million dollars (USD equivalent) (Fund, 2019). In 2009, the Bangladesh Climate Change Strategy Action Plan (BCCSAP) was approved. This action plan helps clarify the uses of GCF money and expands the 2005 NAPA (MEF, 2005).

Karim and Noy 2020 examine the distribution of disaster risk reduction funding under the UNDRR. They utilize funding allocation data at the sub-district (*upazila*) level for the 2010-2011 and 2013-2014 fiscal years in combination with annual rainfall data (a proxy for flood risk) to examine the distribution of adaptation aid and risk prevention across Bangladesh. They find that funding was positively correlated with pre-existing social vulnerability and flood risk—the exact demographic which represents efficient fund usage. This provides evidence addressing concerns that funds would be politically allocated or implicitly skewed towards larger cities.

Since the adoption of government adaptation plans and the partnership with GCF, water-related resilience measures in Bangladesh have developed immensely. Following the 1970 Bhola cyclone, the PRB and related programs have led the construction of cyclone and flood shelters, the elevation of roads above sea level, the facilitation of community-based disaster preparedness, and the development of warning systems for cyclones, storm surges, and floods (MEF, 2005). These implementations—most notably the warning systems—have been declared successful programs by several research teams (Eskander and Barbier, 2022; Islam et al., 2020; Ali et al., 2018). To further determine the effectiveness of policies and adaptation strategies, a number of additional analyses have been completed.

Karim 2018 uses dually-identified treatment groups to measure disaster risk exposure among various populations. The dually-identified groups are defined geographically, using targeted rainfall data, as well as by self-selection into the treatment group (a method which poses its own host of drawbacks and benefits) (Karim, 2018). Haque et al. 2022 explored community flood resilience in the Jamuna floodplain, an area which is among the most at-risk in Bangladesh. Haque et al. 2022 creates a resilience index drawn from four categories of resilience—physical, institutional, economic, and social. Each of these factors contribute to a community's ability and efficiency of recovering from a severe flood. Some factors included are the proximity to a river, diversity of employment types, presence of education programs, and infrastructure strength (Haque et al., 2022). Through this inquiry, the team was able to conclude that diversity of employment types and durable building materials were useful adaptation solutions.

Despite the robust degree of research and effort within building Bangladesh's climate resilience, there is little research which specifically examines home infrastructure strength and overall household well-being through a non-food spending lens. Household infrastructure is one of many components which build a household's resilience, but it is nonetheless vital in creating a resilient, safe, and secure household environment.

Economic Theory

Karim 2018, which examines the household level response to frequent natural disasters, concluded that following a severe weather shock, reductions in income and expenditures will occur. They also note that these reductions are generally reflected in non-food categories, like health and education, which implies the possibility that natural disasters can reduce longer-term investment and trap households in cycles of poor health outcomes, reduced education levels, and poverty (Karim and Noy, 2016). This evaluation of non-food spending utilized by Karim and Noy 2016 is key to this analysis.

It is proposed that home building material has an influence on household resilience according to its "strength" or resistance to weather calamities. A stronger wall, made of brick or tin for example, will be more resilient to a weather event than kutcha (mud, sticks, and grasses), because it can take more damage before it is destroyed. Thus, households who have built with a more resilient material are more likely to maintain their income and non-food spending following a calamity, while those whose homes are made of a less strong material will need to invest in rebuilding efforts and see decreases in their other non-food spending.

Data Description

The question addressed in the following research measures the effectiveness of building material type on climate resilience using panel data on the basis of household. The data was obtained from the Bangladesh Climate Change Adaptation Survey (BCCAS), which was conducted by International Food Policy Research Institute (IFPRI). A description of the methods used will be found in the "Empirical Model" section.

The BCCAS survey was conducted in two rounds in 2010 and 2012. 800 households were surveyed, each uniquely identified by a household ID. This structure makes it possible to match round one and round two responses by household and construct a panel data set.

The surveys include a wide range of household and community level data, including demographic information, incidence rate of weather or non-weather shocks, categorical spending by month, and agricultural practices and adaptations. Of the 800 households, only 518 reported weather-based shock data. For the purposes of the following analysis, the wall material of the main dwelling, health and education spending, and non-food spending are identified to evaluate climate resilience by housing material.

Only seven homes in the entire BCCAS data set reported not being impacted by any of the 15 listed severe weather events between rounds one and two of the survey. This is likely because this survey was aimed at identifying climate adaptation strategies, and thus targeted communities who are likely to experience some sort of weather-based destruction over the course of several years. These severe weather events—titled "calamities" within this analysis—include the following severe weather events: flood, flash floods, droughts, salinity increase, sea level rise, frequent rainfall, temperature rising/sun intensity increasing, temperature variability, soil/river erosion, tornado or very high winds, cyclone, seasons changing, hailstorm, cold weave, and water level decrease. The dummy variable "unexpected calamities" used in the following analysis identifies the treatment group as households who experience one or more of the following events between rounds one and two of the survey: floods, flash floods, tornados/high winds, cyclones, or hailstorms. Of the 15 calamities surveyed, these five are used because of their unpredictability and tendency to cause significant infrastructure damage. The control group is thus made up of households which do not experience any of the five unexpected calamities.

The second dummy variable utilized in this analysis is characterized by severe home damage. Rather than identification based on the occurrence of a severe weather event, this dummy variable identifies households which experience damage caused by any of the immediate and severe calamities included in "unexpected calamities." Damage reports in the survey allowed household respondents four options: high loss, moderate loss, minor loss, or not applicable. The nature of this question causes potential selection bias within these results; people have a higher likelihood of remembering severe losses compared to minor losses, so it is possible that minor losses are under reported relative to major losses. Additionally, because this data are self-reported and relative, there is possible variation among the magnitude of realized damages if households define minor, moderate, and high losses differently. Thus, the inclusion of minor and moderate losses, while reasonable given the theory, are not part of this treatment group because of the variation in actual damage within these categories. Notably, due to data constraints, the number of observations in the treatment group of this dummy variable is only 27 households, creating 54 samples in the analyses where both survey responses are utilized.

Table 1 shows the primary variables of interest (non-food spending, education and health spending, and housing material strength) as well as demographic variables (total number of calamities experienced, number of rooms in primary homestead, and total asset value of household (HH)) by control and treatment of unexpected calamity occurrence. The independent variable, labeled “strong” or “weak” wall in Tables 1 and 2, is a dummy variable that signifies the strength of walls of a household’s main dwelling. Walls which are made of concrete, brick, cement, tin, or corrugated iron sheet are labeled as “strong,” while those made of mud, wood, bamboo, or other materials are “weak.” Table 2 shows equivalent information as Table 1, but it uses a different shock specification: severe damage.

Table 1

Data Description - Unexpected Calamities

Occurrence of an Unexpected Calamity:	No Unexpected Calamity	Unexpected Calamity		
	Control (0)	Treatment (1)	Total	P-Value
Observations	58 (11.6%)	443 (88.4%)	501 (100.0%)	
Average Monthly Non-Food Spending	3,667.655	4,900.469	4,757.748	0.446
Average Monthly Education & Health Spending	643.412	955.614	919.471	0.290
Average Number of Calamities Experienced	1.638	1.912	1.880	0.249
Average Number of Rooms	2.483	2.391	2.401	0.602
Average Total Value of HH	42,662.931	51,741.216	50,683.906	0.355
Strong Wall (0)	25 (43.1%)	211 (47.6%)	236 (47.1%)	0.516
Weak Wall (1)	33 (56.9%)	232 (52.4%)	265 (52.9%)	

Notes: Observations in this table include data from 2010 and 2012. All monetary values are in Bangladeshi Taka.

Empirical Model

The empirical strategy utilized in this analysis is a two-way fixed effects model with an interaction term between the weather-related shock variable and the independent variable of interest—wall material strength. Based on the conclusions of Karim and Noy 2020 and Karim and Noy 2016, households whose homes are built from weak wall materials are expected to experience a greater decrease in the dependent variable (non-food spending or health and education spending) than households with strong walls when they experience a severe weather shock.

Table 2

Data Description - Severe Damage

Occurrence of Severe Home Damage:	No Severe Loss Control (0)	Severe Loss Treatment (1)	Total	P-Value
Observations	79 (74.5%)	27 (25.5%)	106 (100.0%)	
Average Monthly Non-Food Spending	3,906.225	3,610.380	3,830.868	0.775
Average Monthly Education & Health Spending	620.998	605.142	616.959	0.922
Average Number of Calamities Experienced	2.633	1.667	2.387	0.054
Average Number of Rooms	2.139	2.000	2.104	0.453
Average Total Value of HH	48,251.646	45,026.400	47,476.346	0.778
Strong Wall (0)	44 (55.7%)	16 (59.3%)	60 (56.6%)	0.747
Weak Wall (1)	35 (44.3%)	11 (40.7%)	46 (43.4%)	

Notes: Observations in this table include data from 2010 and 2012. All monetary values are in Bangladeshi Taka.

Fixed effects for both place and time are utilized to isolate the phenomenon being identified (see equation 1 below). Village fixed effects will capture time-invariant observable and unobservable qualities of villages which influence the households in the surrounding area, such as geographically varied wealth disparities. Time fixed effects, in this case included by survey round, control for observable and unobservable qualities which change broadly across each round, like inflation. The utilization of two-way fixed effects, along with two additional household-specific time variant controls, allows this model to isolate the effect of wall material and calamity experience on non-food spending.

$$\begin{aligned} NonFoodSpending_i = {} & \beta_0 + \beta_1 WallMaterial_i + \beta_2 Shock_i \\ & + \beta_3 (Shock_i x WallMaterial_i) + Controls_i \\ & + i.village_i + i.round_i + \varepsilon_i \end{aligned} \tag{1}$$

In equation 1 above, *Shock* is either the "unexpected calamity" dummy or "severe damage" dummy, and the *Controls* include HH's total asset value, number of rooms in HH's home, and FE for round (1 or 2).

Analysis

Table 3 shows the results of the two-way fixed effects analysis among four combinations of dependent variables (education and health spending and non-food spending) and shock specifications (calamity occurrence or damage occurrence). The theory presented suggests that a household with weak walls who experiences a calamity will face larger consequences in terms of non-food spending than a household with stronger walls. The results shown in Table 3 generally indicate this is the case. The core regression of this analysis is column one of Table 3. The interaction term titled "Unexpected Calamity Occurrence # Weak Wall"' in this column signifies the change in the average household's education and health spending given that a household's main dwelling has weak walls and that the household experiences an unexpected calamity. The value of this parameter is -1,115 Bangladeshi Taka. In context, this means that the average household with weak walls who experiences a calamity will expect a decrease in their monthly education and health spending by about 1,115 Taka. This magnitude is both statistically and economically significant, given that average monthly spending of this sampled population was just under 5,000 Bangladeshi Taka (Table 1).

Two more important parameters of note are the "Unexpected Calamity Occurrence" and "Weak Wall" results for this same core model. Each of these parameters is insignificant. Within this analysis, this indicates that experiencing an unexpected calamity with a strong wall does not carry this same decrease in health and education spending, and that having a weak wall alone (without the occurrence of an unexpected calamity) does not correlate to a decrease in health or education spending. The ambiguity of these two parameters, as well as the interaction term of interest, all align with the theory outlined above, and suggest that wall strength can increase a household's resilience to calamities.

Column three, which analyzes education and health spending changes for households who reported severe calamity-related infrastructure damage, captures less robust results. The interaction term here ("Severe Damage by Calamity # Weak Wall") signifies the average change in education and health spending of a household given that they have weak walls and that they report experiencing severe damage between the two survey rounds. In this model, where a different shock specification is utilized, the parameter is negative but not significant. As mentioned previously, the "severe damage" dummy variable is self-reported and thus allows for potential data inconsistency. The reason for this lack of significance could be due to the smaller sample size or potential ambiguity caused by the self-reported nature of damage experienced by a household.

Non-food spending analyses (columns two and four in Table 3) provide additional robustness checks and are discussed in the "Robustness Checks and Limitations" section.

Table 3

Regression Results

Dependent Variable	Edu. & Health Spending	Non-Food Spending Loss	Edu. & Health Spending	Non-Food Spending
Shock Specification	Unexpected Calamity Occurrence	Unexpected Calamity Occurrence	Severe Damage	Severe Damage
Unexpected Calamity Occurrence	355.9	1,148		
	(320.3)	(1,040)		
Unexpected Calamity Occurrence # Weak Wall	-1,115***	-2,415*		
	(412.6)	(1,340)		
Severe Damage by Calamity			353.2	1,287
			(255.4)	(1,040)
Severe Damage by Calamity # Weak Wall			-371.8	-3,175**
			(391.6)	(1,594)
Weak Wall (Dummy)	507.9	820.3	-33.30	757.3
	(322.2)	(1,047)	(253.0)	(1,030)
Number of Rooms	363.7***	1,154***	203.5**	964.1**
	(89.82)	(291.8)	(95.47)	(388.7)
Total Asset Value of HH	0.005***	0.0291***	0.003**	0.028***
	(0.001)	(0.004)	(0.001)	(0.005)
Round 2 FE	449.8**	1,263**	345.7**	869.6
	(188.3)	(611.7)	(134.6)	(548.0)
Constant	128.0	-773.3	986.6**	1,466
	(559.9)	(1,819)	(487.3)	(1,984)
Observations	997	997	209	209
R-Squared	0.173	0.163	0.245	0.348
r2_a	0.134	0.124	0.098	0.221
F	4.436	4.129	1.663	2.733
Village FE	Yes	Yes	Yes	Yes

Standard errors in parentheses: ***p<0.01. **p<0.05. *p<0.1.

Robustness Checks & Limitations

To build the robustness of this analysis, two different overarching methods were used. One is the use of village fixed effects and the other includes alternate variable identifications. Village level fixed effects were used on all regressions in this analysis; the primary regressions can be found without village fixed effects in the appendix (Table 0.1). While household fixed effects would be possible, it causes extremely low variation in home material type because few households change wall material during the two-year period between survey one and two. Instead, village fixed effects are utilized to hold time-invariant variables constant at the village level. The second form of robustness analysis is by using several levels of specifications. The core specification of this analysis uses the dependent variable of education and health spending with the unexpected calamity occurrence shock. Three additional specifications are displayed in Table 3 and all but one align with the theory. The third column, discussed in the analysis section, indicates a result in the correct direction but does not carry significance. Beyond these alternative specifications, further analyses are shown in the appendix in Tables 0.1 and 0.2. They generally follow a similar pattern; however, some vary moderately from one another, especially in the case of the regressions which do not utilize village fixed effects.

Potential sources of endogeneity include the type of walls that households elect to maintain and the idea that weather events are not fully random. If weather shocks are not truly random—which is reasonable considering the nuanced spatial effects of climate and weather systems—then households who have experienced frequent calamities might opt to spend more on a more durable wall. If this phenomenon occurs, it introduces endogeneity into the picture. Additionally, because these data are collected by survey, there is potential bias in reporting of damage and calamity occurrence. There could be behavior trends across households which report more or less thoroughly due to recent investment in their homestead, their location within a village, or by their tendency to report differently because of their home material type.

Robustness and overall quality of this analysis could be further improved through several avenues. Ideally, a larger sample size and longer time frame would be utilized. If there was greater variation between the survey dates and more rounds, a more thorough analysis could be completed. Additionally, though much work was done to improve the identification of damage-causing calamities, analysis could benefit from clearer distinction of damage level and calamity severity, ideally through a method that is not self-reported.

Conclusions

This analysis indicates that resilience to climate change-related calamities is greater with a sturdier home building material among households in Bangladesh. It is important to note that this result is variable on the identification of calamity-occurrence and is shown most prominently in models which utilize the "severe loss" shock identifier. This result helps build the case for investment in more resilient infrastructure in the face of a growing climate crisis, which will continue to impact the well-being of Bangladeshi people. With further investigation of this topic, the PRB and supporting non-governmental actors will continue to develop more efficient adaptation systems. Through strengthening of existing homestead infrastructure, and creating new, resilient infrastructure, households will be able to retain more of their income to aid in their journey of improving health outcomes, education attainment, and holistic well-being.

References

Ali, F. M., Ingirige, B., and Zainal Abidin, N. A. (2018). Assembling and (Re)Assembling Critical Infrastructure Resilience in Khulna City, Bangladesh. *Procedia Engineering*, 212:832–839.

Bandyopadhyay, S. and Skoufias, E. (2015). Rainfall variability, occupational choice, and welfare in rural Bangladesh. *Review of Economics of the Household*, 13(3):589–634.

Chisty, M. A., Muhtasim, M., Biva, F. J., Dola, S. E. A., and Khan, N. A. (2022). Sendai Framework for Disaster Risk Reduction (SFDRR) and disaster management policies in Bangladesh: how far we have come to make communities resilient? *International Journal of Disaster Risk Reduction*, 76:103039. Publisher: Elsevier.

Eskander, S. M. and Barbier, E. B. (2022). Long-term impacts of the 1970 cyclone in Bangladesh. *World Development*, 152:105793. Publisher: Elsevier.

Fund, G. C. (2019). Bangladesh. Publisher: Green Climate Fund.

Haque, M. M., Islam, S., Sikder, M. B., and Islam, M. S. (2022). Community flood resilience assessment in Jamuna floodplain: A case study in Jamalpur District Bangladesh. *International Journal of Disaster Risk Reduction*, 72:102861.

Islam, S., Chu, C., and Smart, J. C. (2020). Challenges in integrating disaster risk reduction and climate change adaptation: Exploring the Bangladesh case. *International journal of disaster risk reduction*, 47:101540. Publisher: Elsevier.

Karim, A. (2018). The Household Response to Persistent Natural Disasters: Evidence from Bangladesh. *World Development*, 103:40–59.

Karim, A. and Noy, I. (2016). Poverty and natural disasters: A regression meta-analysis. *Review of Economics and Institutions*, 7(2).

Karim, A. and Noy, I. (2020). Risk, poverty or politics? The determinants of subnational public spending allocation for adaptive disaster risk reduction in Bangladesh. *World Development*, 129:104901. Publisher: Elsevier.

Lee, H. and Romero, J. (2023). IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. IPCC, Geneva, Switzerland. Technical report, Intergovernmental Panel on Climate Change (IPCC). Edition: First.

MEF (2005). National Adaptation Programme of Action (NAPA).

MEF (2009). Bangladesh Climate Change Strategy and Action Plan 2009.

Wannous, C. and Velasquez, G. (2017). United Nations Office for Disaster Risk Reduction (UNISDR)—UNISDR's Contribution to Science and Technology for Disaster Risk Reduction and the Role of the International Consortium on Landslides (ICL). In Sassa, K., Mikoš, M., and Yin, Y., editors, *Advancing Culture of Living with Landslides*, pages 109–115. Springer International Publishing, Cham.

Appendix
Additional Tables

Table 0.1

Table 3 analysis without village fixed effects. These regressions depict the importance of these locally based spatial controls in isolating spending effects from wall material type and shocks.

Dependent Variable	Edu. & Health Spending	Non-Food Spending Loss	Edu. & Health Spending	Non-Food Spending
Shock Specification	Unexpected Calamity Occurrence	Unexpected Calamity Occurrence	Severe Damage	Severe Damage
Unexpected Calamity Occurrence	363.6	833.8		
	(273.1)	(860.2)		
Unexpected Calamity Occurrence # Weak Wall	-549.5	-1,474		
	(386.7)	(1,218)		
Severe Damage by Calamity			473.9**	1,552**
			(191.4)	(776.7)
Severe Damage by Calamity # Weak Wall			-428.2	-2,368*
			(309.1)	(1,254)
Weak Wall (Dummy)	232.1	417.8	-86.65	166.1
	(273.6)	(861.5)	(150.7)	(611.4)
Number of Rooms	464.9***	1,255***	242.5***	1,125***
	(81.84)	(257.7)	(78.13)	(317.0)
Total Asset Value of HH	0.006***	0.031***	0.002**	0.027***
	(0.001)	(0.004)	(0.001)	(0.004)
Round 2 FE	449.9**	1,223**	356.0***	929.2*
	(192.9)	(607.4)	(131.8)	(534.9)
Constant	-670.2**	-145.6	-31.81	19.20
	(305.3)	(961.4)	(200.0)	(811.5)
Observations	997	997	209	209
R-Squared	0.095	0.139	0.149	0.270
r2_a	0.0891	0.133	0.123	0.248
F	17.23	26.56	5.874	12.43
Village FE	No	No	No	No

Standard errors in parentheses: ***p<0.01. **p<0.05. *p<0.1.

Table 0.2

Robustness checks among further shock specifications.

Dependent Variable	Edu. & Health Spending	Edu. & Health Spending	Edu. & Health Spending	Edu. & Health Spending
Shock Specification	Unexpected Calamity Occurrence	Unexpected Calamity Occurrence (Including Droughts)	Severe Damage	Any Damage (Minor, Moderate, or Severe)
Unexpected Calamity Occurrence	355.9			
	(320.3)			
Unexpected Calamity Occurrence # Weak Wall	-1,115***			
	(412.6)			
Unexpected Calamity Occurrence (with Droughts)		806.3		
		(506.1)		
Unexpected Calamity Occurrence (with Droughts) # Weak Wall		-750.7		
		(641.0)		
Severe Damage by Calamity			353.2	
			(255.4)	
Severe Damage by Calamity # Weak Wall			-371.8	
			(391.6)	
Any Damage by Calamity				-209.7
				(420.9)
Any Damage by Calamity # Weak Wall				-163.4
				(522.6)
Weak Wall (Dummy)	507.9	650.0	-33.30	-17.53
	(322.2)	(629.9)	(253.0)	(261.7)
Number of Rooms	363.7***	374.7***	203.5***	364.1***
	(89.82)	(90.10)	(95.47)	(90.59)
Total Asset Value of HH	0.005***	0.005***	0.003**	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)
Round 2 FE	449.8**	450.4**	345.7**	457.0**
	(188.3)	(188.9)	(134.6)	(189.0)
Constant	128.0	-415.4	986.6**	360.8
	(559.9)	(709.2)	(487.3)	(549.6)
Observations	997	997	209	997
R-Squared	0.173	0.169	0.245	0.167
r2_a	0.134	0.129	0.098	0.128
F	4.436	4.287	1.663	4.251
Village FE	Yes	Yes	Yes	Yes

Standard errors in parentheses: ***p<0.01. **p<0.05. *p<0.1.

Housing the Unhoused: How Affordable Housing Impacts Homelessness in the U.S.

Faye Dingle
ECON 381: Introduction to Econometrics (Advisor: Amy Damon)

Lack of affordable housing, homelessness, and rising prices for both renters and home buyers all characterize the housing crisis in America. Anacker (2019) states that housing has become less affordable in the past few decades because of increased restrictions and costs for developers, gentrification, and decreased government funding. She also argues these higher cost burdens can cause displacement and potentially homelessness. The number of unhoused[1] people, despite improving prepandemic, has risen above 2007 levels in the past few years (Department of Housing and Urban Development, 2023). This paper will attempt to answer the question: Is affordable housing effective in reducing homelessness and do these results hold for people of different races and genders?

There has been ample research investigating the relationship between affordable housing and homelessness, with most finding that increasing affordable housing can help to decrease homelessness. Many studies have investigated the Low-Income Housing Tax Credit (LIHTC) specifically. In one recent study, Liaw (2023) finds that one additional LIHTC unit is associated with a reduction of 2.8 unhoused people. She also asserts that the LIHTC increases the affordable housing supply. The placement of LIHTC units, however, appears to lack thorough planning. Shamsuddin & Cross (2020) found that Black and impoverished neighborhoods in Boston were no more likely to have higher concentrations of LIHTC units than better-off areas, conditional on the presence of LIHTC units. This suggests that LIHTC policies do not target non-White populations even if those groups may need more assistance.

I answer my question using annual data on the number of unhoused individuals in the U.S. and newly built LIHTC units. I then run two-way fixed effect regressions. Results show that more LIHTC units correlate with fewer unhoused people. These results

[1] I use the term "unhoused" instead of "homeless" when referring to individuals to mitigate negative connotations and to emphasize its link to structural issues rather than personal ones (Abrams, 2023). I continue to use the term "homelessness" when referring to the issue as a whole since "houselessness" carries its own connotations

Faye Dingle is a junior with majors in Economics and Mathematics and a minor in Geography. Correspondence concerning this article should be addressed to fdingle@macalester.edu.

held across almost all racial and gender groups but differed in magnitude and statistical significance between groups. White- and male-identifying people seem to benefit the most from the program, highlighting a policy shortcoming.

Context

In 1986, the federal government introduced the low-income housing tax credit program (LIHTC). The federal subsidy program gives tax incentives to developers creating or renovating affordable housing that targets people making at most 60% of the Area Median Income (AMI). This program accompanies other federal housing programs, such as the Section 8 Housing Choice Voucher, which directly subsidizes housing for low-income families (Usowski & Hollar, 2008). The average LIHTC project has around 68 units.

It is also important to acknowledge that while the LIHTC program plays an important role, it is not a paragon of homelessness solutions. Ellen et al. (2018) compared rental units to LIHTC units and found that LIHTC units are in neighborhoods with poorer labor markets, higher pollution, and lower school performance, but better public transportation. They also found that of the LIHTC tenants, people of color were found to live in less advantaged neighborhoods. Furthermore, Williamson (2011) examined Florida LIHTC residents and found that the majority are cost-burdened, with some being severely cost-burdened. These results also held across races, as White people did not fare better with cost burdens than people of color. Because of these findings, potential policy solutions should also account for impacts on quality of life for different demographics.

Literature

Homelessness in the United States is mainly caused by substance abuse, violence, and mental illness (Zhao, 2023). One of the main solutions to this issue is increasing the number of affordable housing units, and the primary factor driving the construction of affordable housing units is the LIHTC (Clark, 2016). Most studies find that affordable housing and homelessness have a negative correlation. They are also well linked. Quigley et al. (2001) find that housing supply and price mostly explain the variation in homelessness across cities. They also find that substantially increasing housing affordability and availability can reduce homelessness in the U.S.. This study matches Kim & Sullivan (2021), which finds that affordable housing units created by the LIHTC, combined with other homeless services, are effective in reducing homelessness compared to using homeless services alone.

Despite these positive results, an affordable housing approach may fall short in application. For example, Reid (2020) finds the costs to develop an additional LIHTC unit rose from $425,000 in 2016 to $480,000 in 2019. These cost barriers may disincentivize developers and make it hard to build enough units. Other issues include zoning and cost restraints, which believe will hinder a Biden Administration program called "Housing First," which prioritizes building construction to end homelessness. Affordable housing also may not be adding to the housing supply. Malpezzi & Vandell (2002) found no significant relationship between the number of LIHTC and other subsidized housing units and the overall housing supply, suggesting the LIHTC may be replacing rather than adding units. These setbacks bring into question whether affordable housing is the right approach to ending homelessness.

Alternatively, affordable housing may need to be paired with other solutions. However, there is a lack of consensus. Discrepancies in ideas may come from differing beliefs about what causes homelessness. Martin (2015) asserts that mental health is the leading cause of homelessness and that governments should allocate more resources to mental health. In contrast, Mansur (2002) finds that increased reliance on and awareness of housing subsidies can significantly improve homelessness conditions in California. Researchers have come to several different conclusions on the best approaches. Allen et al. (2022) argue that business communities can best support unhoused people through social, informational, and financial support rather than physical assets such as affordable housing or medical supplies. In another study, Loftus-Farren (2011) argues that tent cities can be instrumental in helping cities provide interim solutions for the unhoused population while they can get their feet on the ground. In contrast, Klammer & Scorsone (2022) find that providing temporary shelter may not be the best solution as it places the responsibility of relocating on unhoused individuals. To account for these differing and opposing solutions, it would be ideal to compare the effects of these variables versus affordable housing units.

To study this subject in depth, it is also important to acknowledge the roles that race and gender play in the homelessness crisis. Black and Indigenous people are the most overrepresented racial groups experiencing homelessness because of limited economic mobility, discrimination, and the criminal justice system (Olivet et al., 2021). Along these lines, Fusaro et al. (2018) found that homelessness was the highest for non-Hispanic Black people and Hispanic people of any race. They also found that the Black-White gap in homelessness was significant after controlling for variables such as gender, education, age, and more, but the Hispanic-White gap was not. Besides being overrepresented, people of color also face higher rates of racial stigma, which correspond to poorer psychological and physical health (Weisz & Quinn, 2018). Black people may also receive fewer services. Carter (2011) found that Black people were less likely to migrate to homeless services than White people. Because so many studies find that race and homelessness are related, there is a need to differentiate between races when studying homelessness. There is also a relationship between gender and homelessness. In one study, men more frequently reported unsheltered homelessness, and women who reported substance abuse were much more likely to experience unsheltered homelessness (Montgomery et al., 2017). These factors create nuance that must be examined to better understand the issue.

This paper will attempt to address the gaps in knowledge on the effects of the LIHTC units on minority groups. This may help with the lack of agreement on solutions to homelessness since it will distinguish groups in need.

Theory

In this paper, I am modeling how the number of affordable housing units affects the number of unhoused people by race and gender in a given year and state. Currently, there are disagreements as to what causes homelessness. Gender and social status may impact people's perceptions, with some attributing homelessness to personal circumstances and others blaming structural influences (Ljubotina et al., 2022). Assuming homelessness is caused partially by a lack of housing affordability, then increasing the amount of affordable housing units should decrease the total unhoused population. I hypothesize the parameter associated with the number of LIHTC units and total homelessness will be negative because I expect that increasing the number of affordable housing units will allow previously unhoused individuals to afford housing.

I also hypothesize that the results will be less significant for non-White individuals. This is because several factors drive a disproportionate amount of non-White people into homelessness, including housing discrimination, economic inequality, and the homeless response system (Fowle, 2022). Structural racism is also embedded in American history, affecting all aspects of life. Homelessness is no exception. Women also face significant challenges that can lead to homelessness such as intimate partner violence, the wage gap, discrimination, and a weak safety net (Bullock et al., 2020). There are also more males than females in the unhoused population (Department of Housing and Urban Development, 2023), so women and people of other genders could be less likely to be served. Because of this, I hypothesize that the results will be weaker for non-males.

Empirical Strategy

To test my hypothesis, I will use a two-way fixed effect model. For this model to be appropriate, I make three assumptions: the variables are independent and identically distributed, there is no perfect multicollinearity between any of my independent variables, and the expected value of the error term given any independent variable or entity fixed effect is zero. I will add an indicator variable for the state to control for entity fixed effects. This will give a different intercept for each state and control for any differences among states that do not change over time. For example, attitudes toward homelessness and state wealth are all unique to a state but do not change significantly over time. I will also add an indicator variable for the year to control for time fixed effects. This will control for things that affect all states the same but change year to year. This includes factors such as inflation, political party in charge, federal laws, recessions, and COVID-19. The regression will thus be:

$$y_{it} = \beta_0 + \beta_1 LIHTCunits_{it} + \alpha_i + \lambda_t + \upsilon_{it} \tag{1}$$

where y_i is the total unhoused population in a given state and year, and $LIHTCunits_{it}$ is the number of newly constructed LIHTC units in a given state and year. β_0 is the predicted unhoused population when no LIHTC units are built, holding state and year constant, and β_1 is the predicted change in the unhoused population when the number of LIHTC units increase by one, holding state and year constant. The term α_i represents entity fixed effects (state), and λ_t represents time fixed effects (year). The error term υ_{it} contains things that change over both time and entities. I will also use a Breusch-Pagan test for heteroskedasticity to determine whether to run the regressions with robust standard errors.

I will then check the robustness of my regressions. First, I will add a control variable, the number of Continuums of Care (CoCs), to represent homeless services in the U.S.. CoCs provide funding to help unhoused individuals through housing and support programs (Department of Housing and Urban Development, 2024). Since access to homeless services varies among states, changes over time, and could impact both the independent and dependent variables, it is included in the model. However, this may not be the best control variable since the number of CoCs may not change much yearly. I will also lag LIHTC units by one year, since newly constructed units in the previous year might have a higher impact on the current unhoused population. This technique is consistent with the model Kim & Sullivan (2021) use. The explanatory variable will then look like $LIHTCunits_{it-1}$. I will also run a regression by logging the total unhoused population for each state. This will be run with the same regression as before, except y_i represents the log of the total unhoused population for a given state and year. The change to the regression would account for changes in the growth rate of homelessness rather than level changes.

I will also run separate regressions for different demographic groups of unhoused individuals. The first characteristic I will examine is race, creating six regressions for the six categories of race in the HUD dataset. I will also run five more regressions for the different gender groups. If the direction of the beta is the same across all groups, I would know the direction of the beta in the general model is not caused by differences among race or gender groups. In other words, I would know that there is not a large negative effect from one group that overrides a positive effect from another group. The regressions would be in the same form as above, where y_{it} represents the total unhoused population for each racial and gender group.

Data

The data used in this project come from two sources: HUD Point-in-Time Counts and the Low Income Housing Tax Credit Database. The former provides information on the number of unhoused people on a state level from 2007-2023. This dataset also provides demographics, including gender and race, and the number of CoCs. The latter lists each LIHTC project, including the number of units (my explanatory variable), and I will sum the number of units per state for each year. I am using this to proxy affordable housing units in a given state. Though it is not a measure of every affordable housing unit, *The New York Times* estimated that 90% of affordable housing units built in the U.S. use the LIHTC (2012).

Table 1

Descriptive Statistics: LIHTC Units and Unhoused Populations

VARIABLES	Mean	SD
LIHTC Units	1,515.8	2,653.7
Total Unhoused	10,779.3	20,910.0
Total Unhoused by Race:		
White	4,833.8	10,316.2
Black, African American, or African	3,972.7	9,535.7
Asian or Asian American	125.6	452.7
American Indian, Alaska Native, or Indigenous	299.4	796.9
Native Hawaiian or Other Pacific Islander	156.1	451.9
Multiple Races	611.2	1,564.5
Total Unhoused by Gender:		
Female	3,936.7	7,975.1
Male	5,992.6	13,107.0
Transgender	44.1	174.2
Non-binary	28.8	111.2
Gender Questioning	12.6	43.8

Table 1 shows descriptive statistics for LIHTC and unhoused populations. The average number of LIHTC units across all states and years is 1,835. The average number of unhoused individuals is 10,779 people. Both figures have large standard deviations because of significant differences across states and years. White people had the highest number of unhoused people (4834), followed closely by Black people (3973). Other races had significantly lower averages, but they make up significantly smaller amounts of the total U.S. population. The average unhoused population was also significantly higher for men than for women, at 5993 and 3937, respectively. Other gender groups had much lower averages, and they also make up a small proportion of the total population.

Figure 1

Total Unhoused Population

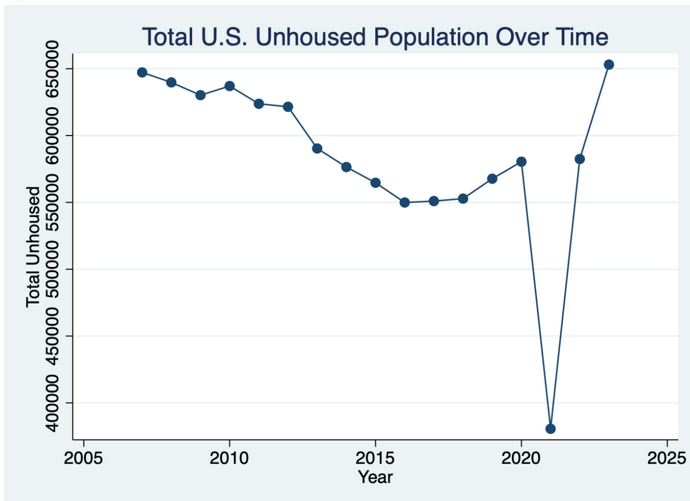

Figure 1 displays the total unhoused population in the U.S. from 2007 to 2023. The variable "Total Unhoused" refers to the total unhoused population in a given year across all states. Homelessness mostly decreased from 2007 to 2016 from about 650,000 people to 550,000 people before rising again from 2016 to 2020. There was a sharp decrease in 2021, dropping from about 575,000 people to about 375,000 people in 2021 and back up to 575,000 again in 2022. This coincides with the federal eviction moratorium, which prohibited landlords from evicting tenants for not paying rent and took effect in October 2020 and was implemented on and off until August 2021 (Liptak & Thrush, 2021). From 2022 to 2023, there was a rapid increase in homelessness, jumping above 2007 levels to over 650,000 people.

Figure 2 breaks down the total unhoused population by state, showing the states with the highest unhoused populations (California, New York, and Florida) and three other states. California, New York, and Florida have a significantly higher unhoused population, making the yearly changes for the other states nearly indiscernible. California's unhoused population also mirrors the total U.S. population patterns with a steady decrease, a steep drop in 2021, and then a sharp increase in homelessness. The other states do not follow this trend, possibly indicating that California is biasing the previous graph.

Figure 3 shows the newly built LIHTC units for the same states as Figure 2. In accordance with their unhoused populations, California, New York, and Florida build the most LIHTC units yearly. However, New York and California build roughly the same number of units despite California having a much larger unhoused population. From the

Figure 2

Total Unhoused Major States

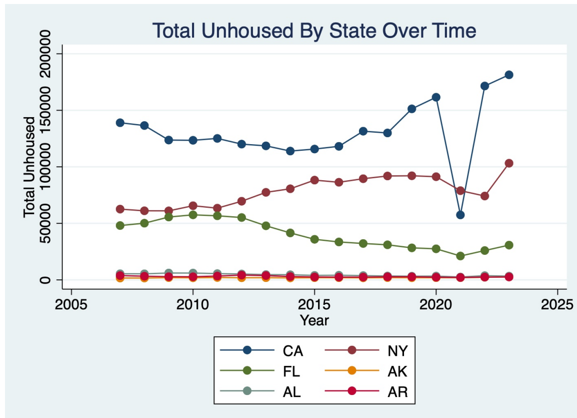

Figure 3

Average LIHTC Major States

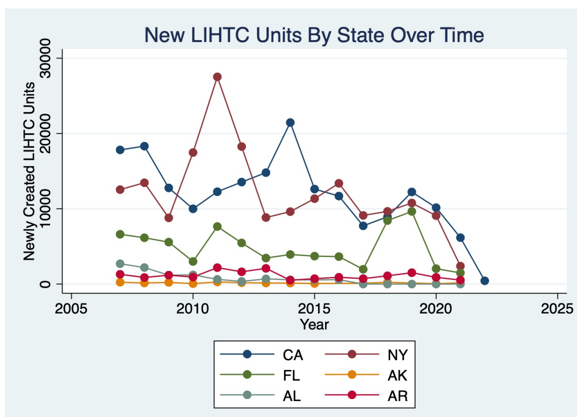

graph, it is unclear that an increase in LIHTC units corresponds to a decrease in the unhoused population. For example, there was a sharp increase in LIHTC units from 2010-2012 in NY, but the overall unhoused population steadily increased during this time and the years after.

Figure 4 is a bin scatter plot that shows the average number of LIHTC housing units placed in service and the total unhoused population for 2007-2023 across all states. The line of best fit shows that across all years, more newly created LIHTC units correspond to a higher yearly unhoused population. This is because places with higher unhoused populations are more likely to construct LIHTC units, so more LIHTC units does not actually cause more homelessness. This demonstrates why it is important to include entity and time fixed effects.

Figure 4

Total LIHTC vs Unhoused

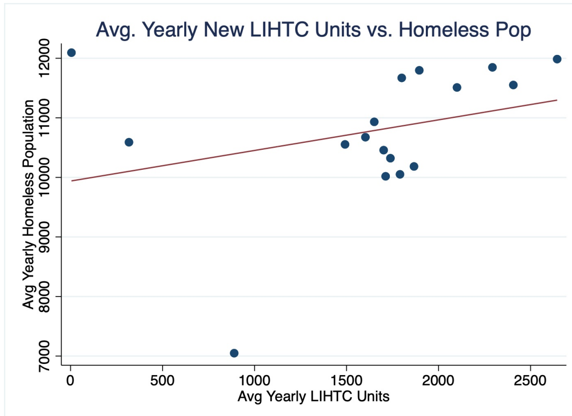

Figure 5

Total Unhoused by Race

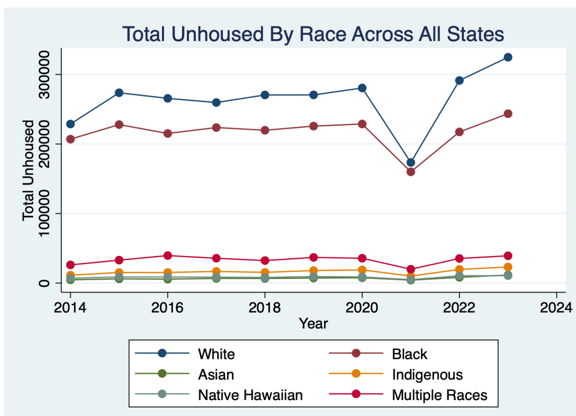

Figure 5 shows the unhoused population across all states, separated by race. Across all years, White people were the largest group of unhoused. However, the number of unhoused Black people is about 80% of the number of White people, despite the fact that they comprise roughly 13.6% of the U.S. population while White people account for roughly 75.5% (U.S. Census Bureau, 2024). People of multiple races and Indigenous people make up the next largest proportions, respectively.

Figure 6 shows the unhoused population by gender for the U.S. men make up the majority of the unhoused population, at roughly 60% in 2023 with women at 38%. Transgender, non-binary, and questioning individuals each comprised less than 1%. The data for Figures 5 and 6 range from 2014 to 2023, reflecting when HUD began collecting demographic data.

Figure 6

Total Unhoused by Gender

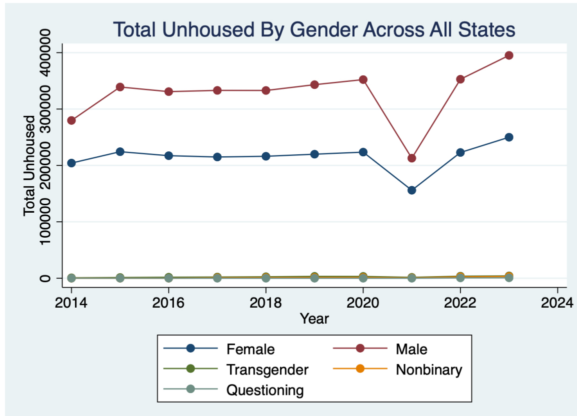

Results

This section will break down the regressions I ran, including their significance. Table 2 shows the effect of newly built LIHTC units on the total unhoused population level amount and growth rate. State and year were added as indicator variables, so each state and year has an intercept. All regressions are run with robust standard errors. The LIHTC coefficient for the first regression with just the state and year fixed effects was not statistically significant. With no newly built LIHTC units, the model predicts an average unhoused population of 3,376 people, holding state and year constant. This holds at the 1% confidence level. The second regression includes a control variable for the number of CoCs in a state and year, but neither coefficients for LIHTC units nor the number of CoCs were statistically significant. Because of this, the control is not used in any further regressions. This regression predicts an average unhoused population of 3,747 people given no LIHTC units, which was significant at the 1% level. The third regression lags the number of LIHTC units by one year. A one-unit increase in LIHTC units in the year before corresponds to a 0.607-person decrease in homelessness in the current year. This is significant at the 10% level. Since an average LIHTC project has around 68 units, a new project would correspond with a roughly 41-person decrease in the total unhoused population the following year on average. Given no LIHTC units, an unhoused population of 3,742 people is predicted, which is significant at the 1% level. The fourth regression had a coefficient of zero that was not statistically significant.

Table 3 shows regressions of the effect of the previous year's newly constructed LIHTC units on the unhoused population by race while still accounting for state and year fixed effects. It is important to note that racial demographics were only recorded from 2014 onward. A one-unit increase in LIHTC units in the previous year corresponds to a

Table 2

Regression Output

VARIABLES	(1) Two-way FE	(2) Two-way with control	(3) Two-way Lagged	(4) Two-way Logged
LIHTC units	-0.332	-0.327		-8.31e-07
	(0.390)	(0.396)		(6.29e-06)
Number of CoCs		-137.1		
		(298.1)		
Lagged LIHTC			-0.607*	
			(0.352)	
Constant	3,376***	3,747***	3,742***	7.614***
	(740.8)	(910.1)	(752.3)	(0.0582)
Observations	780	780	869	780
R-squared	0.958	0.958	0.953	0.972

Robust standard errors in parentheses

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

**The model included state and time fixed effects that were omitted from the table for clarity

0.813-person decrease in the unhoused population for White people, a 0.451-person decrease for Black people, and a 0.0791 person decrease for Indigenous people, holding state and year constant. These results held at the 5% confidence level for White people, and a 10% confidence level for Black and Indigenous people. Given no new LIHTC units in a given state and year, the model predicts an unhoused Indigenous population of 796 people, significant at the 1% level. The results for Asian and Native Hawaiian people and people of multiple races were statistically insignificant. White people benefit almost twice as much as Black people from the creation of an LIHTC unit. Although there are more White than Black unhoused people, White people account for far less than double the amount of Black unhoused people (Figure 5). The results also indicate that other races do not benefit much at all.

Table 4 shows regressions of the effect of LIHTC units in the previous year on the unhoused population by gender while accounting for state and year fixed effects. It is important to note that "non-binary" and "questioning" were not added as possible responses until 2022, explaining the discrepancy in observation numbers. These regressions show that a one-unit increase in LIHTC in the year before units corresponds to a 1.008-person decrease for males, a 0.427-person decrease for females, and a 0.0325-person decrease for transgender people, holding state and year constant. These results were significant at the 10% level for female people, the 5% level for male people, and the 1% level for transgender people. The reason why the results are the most significant for transgender people remains unclear. This

Table 3

Regression Output

VARIABLES	(1) White	(2) Black	(3) Asian	(4) Native Hawaiian	(5) Indigenous	(6) Multiple Races
Lagged LIHTC	-0.813**	-0.451*	-0.0723	-0.0313	-0.0791*	-0.0547
	(0.355)	(0.237)	(0.0485)	(0.0227)	(0.0404)	(0.0614)
Constant	342.2	128.5	19.28	66.20***	796.4***	100.4
	(607.1)	(382.8)	(28.02)	(23.72)	(63.98)	(143.6)
Observations	546	546	546	546	546	546
R-squared	0.925	0.963	0.793	0.889	0.874	0.865

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
*The model included state and time fixed effects that were omitted from the table for clarity

model predicts that given no LIHTC units, the total female unhoused population would be 778.5 people, which held at the 5% confidence level. These results show that an increase in LIHTC units the previous year is correlated with a much larger decrease in homeless for male people than female people. Although the population of male unhoused people is higher, unhoused female people account for more than half the number of male unhoused people (see Figure 6), yet benefit less than half the amount that male people do.

Table 4

Regression Output

VARIABLES	(1) Female	(2) Male	(3) Transgender	(4) Nonbinary	(5) Questioning
Lagged LIHTC	-0.427*	-1.008**	-0.0325***	-0.0136	0.00382
	(0.243)	(0.477)	(0.0100)	(0.00828)	(0.00824)
Constant	778.5**	698.9	-13.12	-1.806	-2.151
	(361.9)	(820.6)	(15.61)	(6.834)	(3.217)
Observations	546	546	546	384	109
R-squared	0.954	0.918	0.805	0.762	0.949

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1
**The model included state and time fixed effects that were omitted from the table for clarity

Limitations

Although the regressions accounts for time and entity fixed effects, I am still concerned with parts of my model affecting internal and external validity. One threat to internal validity is omitted variable bias, which would cause an issue if there were parts of the error term that affected both affordable housing and homelessness. One possible source of this is state real gross domestic product (not nominal GDP because time fixed effects account for inflation). However, since state real GDP is largely invariant over time for each state, it should mostly be accounted for by controlling for state fixed effects and was thus not included in this model. That being said, it is possible that slight variations over time still have an effect that is not being accounted for. Another possible issue is state population, which poses the same issue as state real GDP.

Another concern for internal validity is measurement error. Accurate measures of the unhoused population are nearly impossible to find, so I am concerned the discrepancies will affect the outcomes of the model. If a subsection of the unhoused population is systematically omitted from the data, it could lead to biased results. Twenty-nine observations of homelessness across all states and years are also labeled as missing. Almost all of these are a lack of data for American Samoa and Mississippi, which poses an issue if this is related to the total unhoused population. Mississippi also had a recorded unhoused population of 2 people in 2021, which I removed since the previous years had numbers in the thousands.

One threat to external validity is using LIHTC units as a proxy for affordable housing units. Although it is one section of affordable housing, affordable housing units that are not built using the tax credit are not included in this model. This means the results could not be generalized to affordable housing and instead just to LIHTC units. These results could also not be generalized to other countries. There are vast differences in affordable housing policy in other countries, as well as different cultures and attitudes around homelessness. To address these issues, I would perform a similar analysis in a random selection of countries.

Conclusion

The results of my regression analysis show a statistically significant effect of the number of LIHTC units in the previous year on the total unhoused population. The results are also economically significant because a roughly two-unit increase in LIHTC units in the previous year corresponds to a one-person decrease in the total unhoused population. Since this is not a one-to-one decrease, other populations besides the unhoused presumably benefit from affordable housing units. When examining the effects of affordable housing units on different races, there were statistically significant results for only White, Black, and Indigenous people. Increasing affordable housing units in the prior year correlated with a decrease in homelessness for all groups, but White people benefited far more than other races. The second-largest decrease was for Black people, but the effect was around half the amount for White people. For gender groups, male, female, and transgender people had statistically significant results. Increasing newly built LIHTC units the previous year corresponded to lower rates of homelessness for these groups, but male people benefited much more than other genders. Decreases in the unhoused population were the second highest for females, but the results were less than half the size of male people. Results for non-binary and gender questioning people were insignificant. The strongest statistical correlation across all demographic groups was for transgender people, but further research is necessary to explain this. Given my results, there is a policy failure for supporting non-White and non-male portions of the unhoused population.

References

Abrams, A. (2023). Is it OK to use the word ‘homeless’ – or should you say ‘unhoused’? https://www.theguardian.com/us-news/2023/jul/20/homeless-unhoused-houseless-term-history

Allen, R. E., Burns, C., & Subrahmanyan, S. (2022). Addressing homelessness with the "human ecology economics" framework: The role of the business community in california. 49(9), 1368–1386. https://www.proquest.com/econlit/docview/2775383620/8E438BE231544FE4PQ/2?parentSessionId=HZQymFRcB4Fgm2B5nUWFkR%2BG5WbzCj7zPmaAWuxgnxs%3D&sourcetype=Scholarly%20Journals. Num Pages: 19

Anacker, K. B. (2019). Introduction: housing affordability and affordable housing. 19(1), 1–16. https://doi.org/10.1080/19491247.2018.1560544. Publisher: Routledge _eprint: https://doi.org/10.1080/19491247.2018.1560544

Bullock, H. E., Reppond, H. A., Truong, S. V., & Singh, M. R. (2020). An intersectional analysis of the feminization of homelessness and mothers’ housing precarity. 76(4), 835–858. https://doi.org/10.1111/josi.12406. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/josi.12406

Carter, G. R. (2011). *From exclusion to destitution: Race, affordable housing, and homelessness.* https://doi.org/10.2139/ssrn.1808950

Clark, A. M. I. (2016). Homelessness and the crisis of affordable housing: The abandonment of a federal affordable housing policy. 25(1), 85–102. https://heinonline.org/HOL/P?h=hein.journals/jrlaff25&i=95

Department of Housing and Urban Development (2023). *PIT and HIC data since 2007.* https://www.hudexchange.info/resource/3031/pit-and-hic-data-since-2007

Department of Housing and Urban Development (2024). *Continuum of care (CoC) program eligibility requirements.* https://www.hudexchange.info/programs/coc/coc-program-eligibility-requirements

Ellen, I. G., Horn, K. M., & Kuai, Y. (2018). Gateway to opportunity? disparities in neighborhood conditions among low-income housing tax credit residents. 28(4), 572–591. https://doi.org/10.1080/10511482.2017.1413584. Num Pages: 20

Fowle, M. Z. (2022). Racialized homelessness: A review of historical and contemporary causes of racial disparities in homelessness. 32(6), 940–967. https://doi.org/10.1080/10511482.2022.2026995. Num Pages: 940-967 Place: Abingdon, United Kingdom Publisher: Taylor & Francis Ltd.

Fusaro, V. A., Levy, H. G., & Shaefer, H. L. (2018). Racial and ethnic disparities in the lifetime prevalence of homelessness in the united states. 55(6), 2119–2128. https://doi.org/10.1007/s13524-018-0717-0. Num Pages: 10

Kim, S. & Sullivan, A. A. (2021). Complementary policies for multidimensional problems: Does the low-income housing tax credit complement homeless services in the USA? 58(5), 881–1094. https://doi.org/10.1177/0042098020941688

Klammer, S. & Scorsone, E. (2022). Housing as a human right: A proposed alternative institutional structure in sacramento. 56(2), 523–529. https://doi.org/10.1080/00213624.2022.2063646. Num Pages: 7

Liaw, E. W.-K. (2023). *Essays on Low-Income Housing Policies*. UC San Diego. https://escholarship.org/uc/item/9hj753dv

Liptak, A. & Thrush, G. (2021). Supreme court ends biden's eviction moratorium. https://www.nytimes.com/2021/08/26/us/eviction-moratorium-ends.html

Ljubotina, O. D., Radovic, M. K., & Ogresta, J. (2022). Determinants of causal attributions of homelessness in croatia. 29(2), 163–190. https://doi.org/10.3935/rsp.v29i2.1972. Num Pages: 28 Place: Zagreb Publisher: Sveucliste Zagrebu, Pravni Fakulted-Univ Zagreb, Fac Law Web of Science ID: WOS:000886592200002

Loftus-Farren, Z. (2011). Tent cities: An interim solution to homelessness and affordable housing shortages in the united states. 99(4), 1037–1081. https://www.jstor.org/stable/23018618. Publisher: California Law Review, Inc.

Malpezzi, S. & Vandell, K. (2002). Does the low-income housing tax credit increase the supply of housing? 11(4), 360–380. https://doi.org/10.1016/S1051-1377(02)00123-7

Mansur, E. T. (2002). Examining policies to reduce homelessness using a general equilibrium model of the housing market. 52(2), 316–340. https://www.proquest.com/econlit/docview/56185753/987B2D3AD559439CPQ/16?sourcetype=Scholarly%20Journals. Num Pages: 25

Martin, E. J. (2015). Affordable housing, homelessness, and mental health: What heath care policy needs to address. 38(1), 67–89. https://www.jstor.org/stable/24459676. Publisher: SPAEF

Montgomery, A. E., Szymkowiak, D., & Culhane, D. (2017). Gender differences in factors associated with unsheltered status and increased risk of premature mortality among individuals experiencing homelessness. 27(3), 256–263. https://doi.org/10.1016/j.whi.2017.03.014

Olivet, J., Wilkey, C., Richard, M., Dones, M., Tripp, J., & Beit-Arie, M. (2021). Racial inequity and homelessness: Findings from the SPARC study. 693, 82–100. https://www.proquest.com/econlit/docview/2572868554/BC67FE4167E34124PQ/7?sourcetype=Scholarly%20Journals. Num Pages: 19

Quigley, J. M., Raphael, S., & Smolensky, E. (2001). Homeless in america, homeless in california. 83(1), 37–51. https://www.proquest.com/econlit/docview/56996315/987B2D3AD559439CPQ/8?sourcetype=Scholarly%20Journals. Num Pages: 15

Reid, C. (2020). The costs of affordable housing production: Insights from california's 9% low-income housing tax credit program.

Shamsuddin, S. & Cross, H. (2020). Balancing act: The effects of race and poverty on LIHTC development in boston. 35(7), 1269–1284. https://doi.org/10.1080/02673037.2019.1657072. Num Pages: 16

U.S. Census Bureau (2024). *Quickfacts: United states*. https://www.census.gov/quickfacts/fact/table/US/RHI225222

Usowski, K. & Hollar, M. (2008). Social policy and the u.s. tax code: The curious case of the low-income housing tax credit. 61(3), 519–529. https://doi.org/10.17310/ntj.2008.3.10

Weisz, C. & Quinn, D. M. (2018). Stigmatized identities, psychological distress, and physical health: Intersections of homelessness and race. 3(3), 229–240. https://doi.org/10.1037/sah0000093. Place: US Publisher: Educational Publishing Foundation

Williamson, A. R. (2011). Can they afford the rent? resident cost burden in low income housing tax credit developments. 47(6), 775–799. https://doi.org/10.1177/1078087411417078. Publisher: SAGE Publications Inc

Zhao, E. (2023). The key factors contributing to the persistence of homelessness. 30(1), 1–5. https://doi.org/10.1080/13504509.2022.2120109. Num Pages: 5 Place: Philadelphia Publisher: Taylor & Francis Inc Web of Science ID: WOS:000851593200001

A Simple Model of Wealth Inequality

Gavin Engelstad

MATH 432: Mathematical Modeling (Advisor: Will Mitchell)

Economic inequalities are a powerful and ever-present factor in people's lives. One's level of wealth determines where they live, who they know, and how healthy they are (Thomas et al., 2018). Therefore, understanding wealth dynamics can be vital to understanding how society as a whole functions.

In this paper, we explore the wealth distribution for individuals in the United States and build a remarkably simple model which approximates wealth inequality in the US. We also examine how wealth inequality can arise both at an individual level and at a more systemic group level.

Literature Review

Conventional Macroeconomic theory based on finding equilibria in a world of representative agents struggles to model heterogeneity in agent behavior and conditions and far-from-equilibrium interactions which change the system without adding massive amounts of complexity (Dosi & Roventini, 2019; Caiani et al., 2016). As an alternative, agent based models (ABMs) can be used to study the effects of these factors on the economy.

Although there are many robust published macro ABMs, very few are able to estimate inequality (Dawid & Gatti, 2018; Seppecher, 2012; Dawid et al., 2016). Instead, toy models like the Yard Sale Model (YSM), which rely on few parameters and simple agent interactions, do a better job (Xu, 2022).

In the YSM, there are many agents each of whom have some level of wealth. Each period, agents are paired up, and an agent in each pair is randomly selected to transfer wealth to the other. A transfer occurs based on a percentage of the net worth of the poorer agent, which varies based on whether the richer or poorer person is receiving the transfer (Chorro, 2016).

This paper will attempt to build a model similar to the YSM that relies on fixed transfer amounts and varying probability of winning the transfer, instead of the YSM where

---

Gavin Engelstad is a Junior with majors in Economics and Mathematics. Correspondence concerning this article should be addressed to gengelst@macalester.edu.

transfer amounts vary and probabilities are equal, based on wealth that similarly approximates the real world wealth distribution.

Economic Background

Model accuracy will be measured via comparisons to real world Lorenz Curves and Gini Coefficients.

A Lorenz Curve (Figure 1) is a line that maps wealth percentiles of the population to the total wealth owned by those at or below that percentile. The 45 degree line represents perfect equality, and the farther a Lorenz Curve curves away from this Line of Equality, the more unequal the wealth distribution is.

Figure 1

Example of a Lorenz Curve

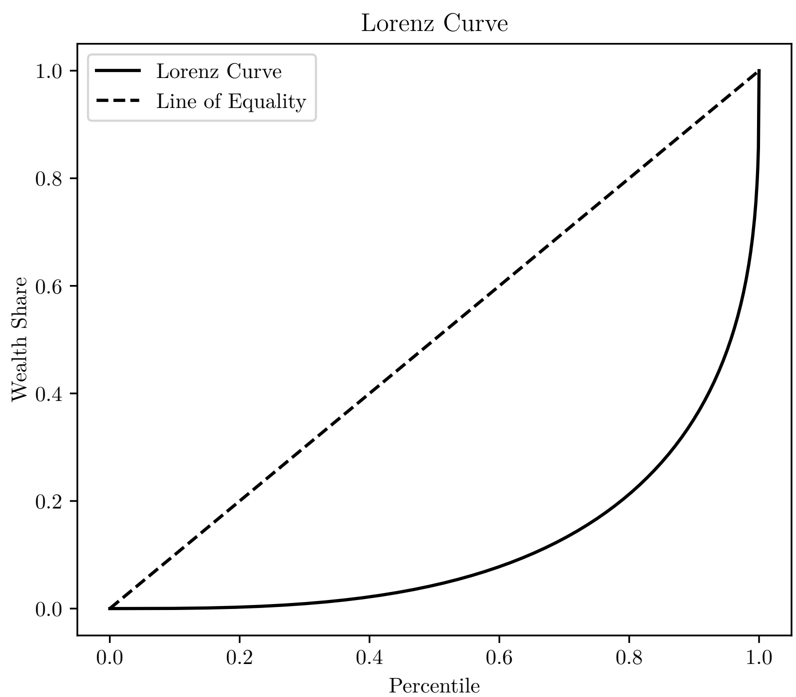

A Gini Coefficient is the ratio of the area between the Lorenz Curve and the Line of Equality and the area under the Line of Equality. It ranges from 0 to 1, where 0 means perfect equality, or that wealth is equally distributed, and 1 means perfect inequality, or that all wealth is controlled by one person.[1] Table 1 shows a handful of Gini Coefficient

[1] A Gini Coefficient above 1 is possible if you allow for negative net worth. This, however, is irrelevant to the model.

values around the world.

Table 1

Gini coefficients around the world. Data from Suisse (2022)

Country	Gini
United States	0.850
China	0.701
Uruguay	0.774
Argentina	0.809
United Kingdom	0.706
Germany	0.788

Region	Gini
Africa	0.879
Asia-Pacific	0.885
Europe	0.816
Latin America	0.858
North America	0.842
World	0.889

A Basic Model of Wealth Inequality

Model

Similar to the YSM, the model will consist of many agents exchanging wealth. Each period

1. Each agent will pair up with another one.

2. One agent in the pair will be randomly selected to transfer wealth to the other.

   Assuming i and j are paired agents with wealth levels w_i and w_j, the probability of i receiving a transfer from j is

$$p_{ij} = \begin{cases} 0.5 + \frac{\alpha}{2} & \text{if } w_i > w_j \\ 0.5 & \text{if } w_i = w_j \\ 0.5 - \frac{\alpha}{2} & \text{if } w_i < w_j. \end{cases} \tag{1}$$

3. A transfer of amount T will be initiated between the agents. If the agent selected to send the transfer has wealth less than T, the transfer will be for whatever amount of wealth the agent has to give.

p_{ij} isn't affected by the magnitude of the difference between w_i and w_j, only whether one is bigger than the other. In this case, the wealthier agent is α more likely to receive the transfer than the less wealthy one. Additionally, Step 3, where the transfer is completed, has a progressive redistributive mechanism built into how the model deals with agents who don't have enough wealth to afford a transfer. This gives a mechanism for wealth to trickle down, even when other factors in the model push wealth up.

The parameters in the model are described in Table 2.

Table 2

Model parameters

Parameter	Meaning
α	Difference in probability of richer and poorer agent receiving the transfer
$\overline{w}$	Average wealth for all agents
T	Transfer amount
n	Number of Agents in the model
t_f	Number of iterations

Model Behavior

When run, the model eventually reaches a steady state wealth distribution. Figure 2 shows the steady state Lorenz Curve for different α values. When α is higher, the richer person is more likely to receive the transfer, meaning the resulting distribution is more unequal. Similarly, a lower α, means the distribution becomes more equal.

Furthermore, the resulting steady state is stable. Figure 3 shows the steady state that results from different initial conditions for the wealth distribution. Even though both simulations start at opposite distributions (Figure 3a), the resulting steady states converge (Figure 3b). Visually, the Lorenz Curves look identical and the Gini Coefficients in the two simulations are within 1% of each other. It did, however, take 100 times longer for the unequal start to converge than the equal start.

When the same simulation is run for larger α values, the unequal initialization begins to converge faster than the equal one. However, even when $\alpha = 1$, a perfectly equal start never takes 100,000,000 periods to converge like the perfectly unequal one does in Figure 3.

Also, social mobility is extremely high in the model. Figure 4 shows the transition matrix between quartiles over 10,000 periods. It shows negligible difference in one's chance of ending up in a certain wealth quartile given any start quartile. This demonstrates that even in the steady state, individual agents experience significant variation in their wealth over time. Even though the wealth share for a given percentile stays mostly constant in the Lorenz Curve, the agents that make up that percentile changes significantly as time progresses.

Figure 2

Model steady state Lorenz Curves for different α values.

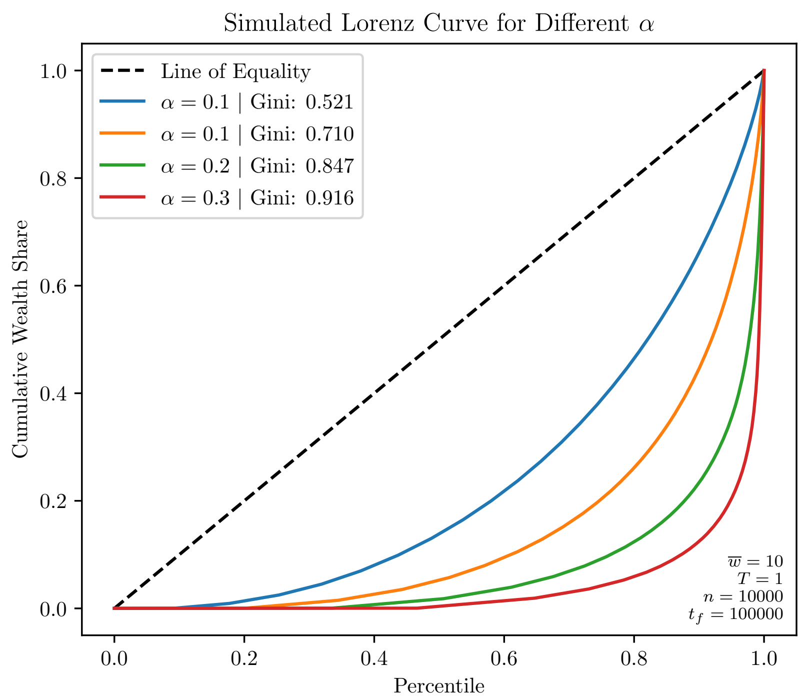

Data Fit

The most surprising part of this model is that despite being such a simple model, it fits real world wealth distributions to a remarkable degree.

Figure 5 shows that when $\alpha = 0.2$, the resulting simulation Lorenz Curve nearly perfectly matches points along the real Lorenz Curve for the US wealth distribution.[2] The model has a calculated R^2 of 0.997, which again suggests the model has an extremely accurate fit.

The Gini Coefficient of the simulation, 0.849, is also very similar to that of the United States in Table 1, 0.850. In fact, even at the steady state, the simulation Gini Coefficient varies by up to around 0.005 period to period. The real Gini value is well within the simulation margin of error.

[2]See Table A1 for a numerical representation of how close these points are to the simulation curve.

Figure 3

Steady states reached with different initial conditions

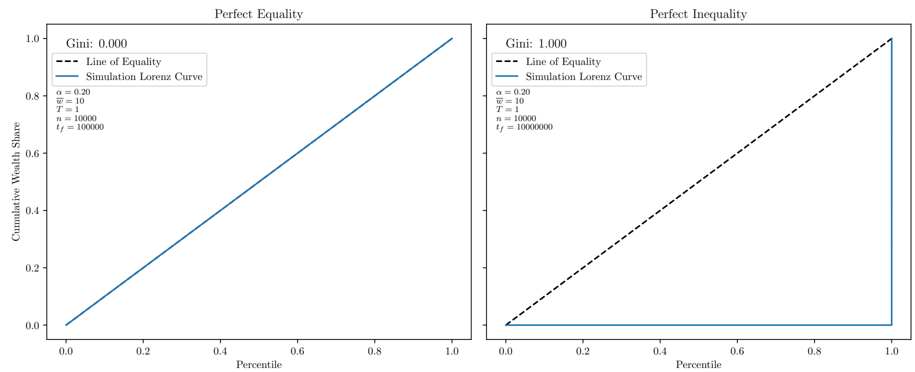

(a) *Perfect equality and perfect inequality initial conditions*

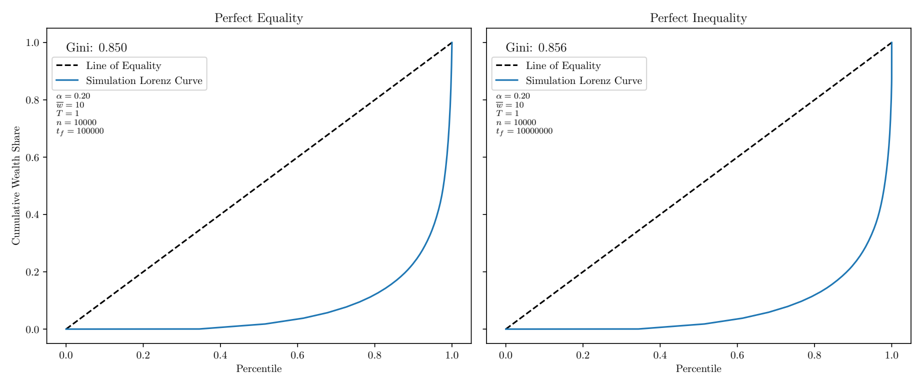

(b) *Perfect equality and perfect inequality steady states*

Figure 4

Transition matrix for probability of moving from the column quartile to the row quartile in 10,000 periods after reaching the steady state. Model parameters are the same as in Figure 3

$$
\begin{array}{c c}
 & \begin{array}{ccccc} Q_1 & Q_2 & Q_3 & Q_4 & Q_5 \end{array} \\
\begin{array}{c} Q_1 \\ Q_2 \\ Q_3 \\ Q_4 \\ Q_5 \end{array} &
\begin{bmatrix}
0.2265 & 0.1665 & 0.2065 & 0.213 & 0.1875 \\
0.177 & 0.2505 & 0.176 & 0.2095 & 0.187 \\
0.199 & 0.1915 & 0.2285 & 0.201 & 0.18 \\
0.1935 & 0.2125 & 0.204 & 0.187 & 0.203 \\
0.204 & 0.179 & 0.185 & 0.1895 & 0.2425
\end{bmatrix}
\end{array}
$$

Modeling Group Wealth Dynamics

Individual wealth isn't the only factor that contributes to systematic wealth inequality. This second model attempts to integrate the advantage one gains from being in a

Figure 5

Simulated Lorenz Curve with real 2019 US Lorenz Curve points overlayed. Data from Aladangady & Forde (2021)

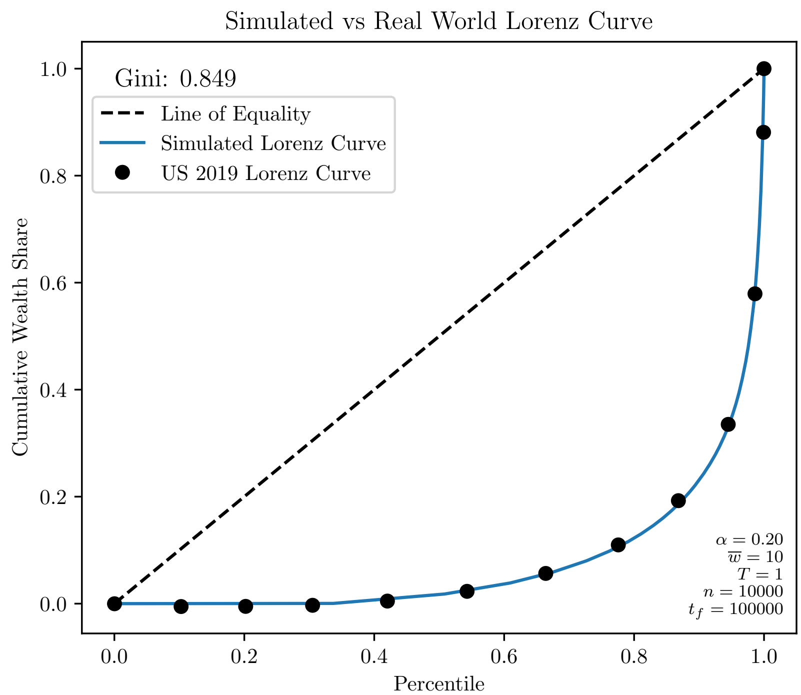

wealthy group into the model presented in the basic model section.

Model

The primary difference between this model and the basic model is that agents are now separated into groups. When deciding where to transfer wealth, the per capita average wealth of the groups agents are a part of, W_i, is also taken into account, not just the wealth of the agents involved in the transfer, w_i.

Each time-step,

1. Agents are paired up

2. In each pair, one agent is selected to receive a transfer and the other is selected to give a transfer.

Letting i and j be the paired agents, then

$$w_{ij} = \begin{cases} 1 & \text{if } w_i > w_j \\ 0 & \text{if } w_i = w_j \\ -1 & \text{if } w_i < w_j \end{cases} \tag{2}$$

and

$$W_{ij} = \begin{cases} 1 & \text{if } W_i > W_j \\ 0 & \text{if } W_i = W_j \\ -1 & \text{if } W_i < W_j. \end{cases} \tag{3}$$

The probability of i receiving a transfer from j given they're paired up is

$$p_{ij} = 0.5 + \frac{\alpha}{2} w_{ij} + \frac{\beta}{2} W_{ij}. \tag{4}$$

3. A transfer is initiated from the transferring agent to the receiving agent. Like in the basic model, if the transferring agent can't afford the transfer, the value of the transfer is capped at the wealth of the transferring agent.

The α parameter in Equation 4 represents the difference in probability of receiving the transfer based on your individual wealth while the β parameter represents the difference in probability from being a part of the wealthier group. Also, Equation 2 and 3 only take into account whether an agent or the group an agent is a part of has more, less, or the same amount of wealth, not the relative magnitudes, like Equation 1 in the basic model.

The parameters in the model are described in Table 3.

Table 3

Model parameters

Parameter	**Meaning**
α	Difference in probability of richer and poorer agent receiving the transfer
β	Difference in probability of agent in the richer group and agent in the poorer group receiving the transfer
$\overline{w}$	Average wealth for all agents
T	Transfer amount
n_g	(For each group g) Number of agents in the group
t_f	Number of iterations

Behavior

The grouped model yields a very similar steady state wealth distribution to the basic model. Figure 6 shows the steady state Lorenz Curve for the grouped model. Compared to the Lorenz Curve when $\alpha = 0.2$ in Figure 2, the shape and Gini Coefficient are very similar.

Figure 6

Lorenz Curve for the grouped model

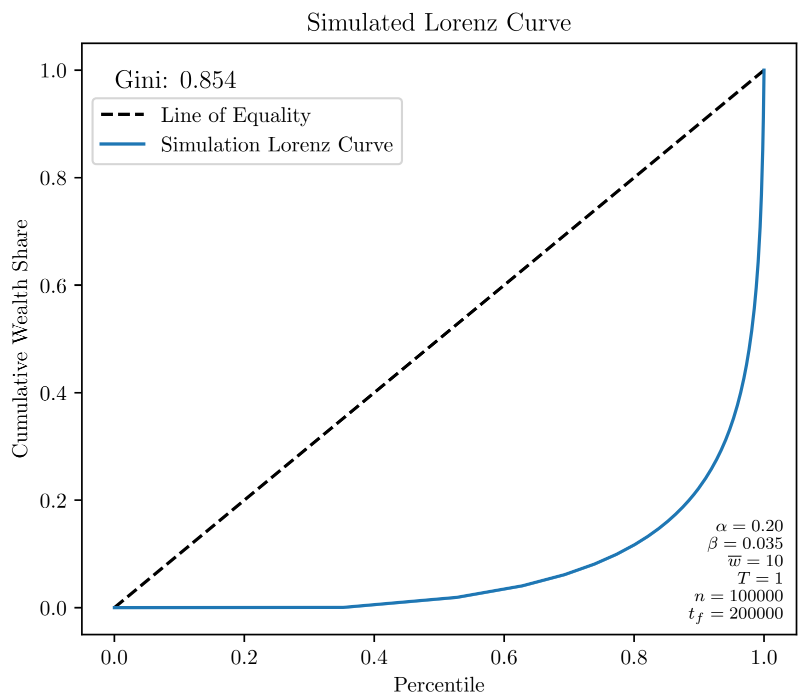

The groups in the model allow for comparison between wealth levels of each group. Figure 7 shows that when $\beta = 0$, the wealth shares held by each group approach the population shares for that group, even when the groups start with a disproportionate wealth distribution. Conversely, if $\beta > 0$, even by a small amount, the steady state wealth distribution leaves one group with a higher share of wealth than the others relative to their population shares.

By shading the Lorenz Curve, group disparities at each wealth share become apparent. Figure 8 shows that Group 1, the dominant group in the simulation, holds a very significant portion of the wealth, especially at higher percentiles. Interestingly, around the 95th percentile the share of wealth held by Group 2 and 3 stops increasing, suggesting that the top of the wealth distribution is only Group 1.[3]

[3]The other sudden changes in group slopes, namely at the 60th and 65th percentiles, I believe, are caused

Figure 7

Simulation wealth shares for each group

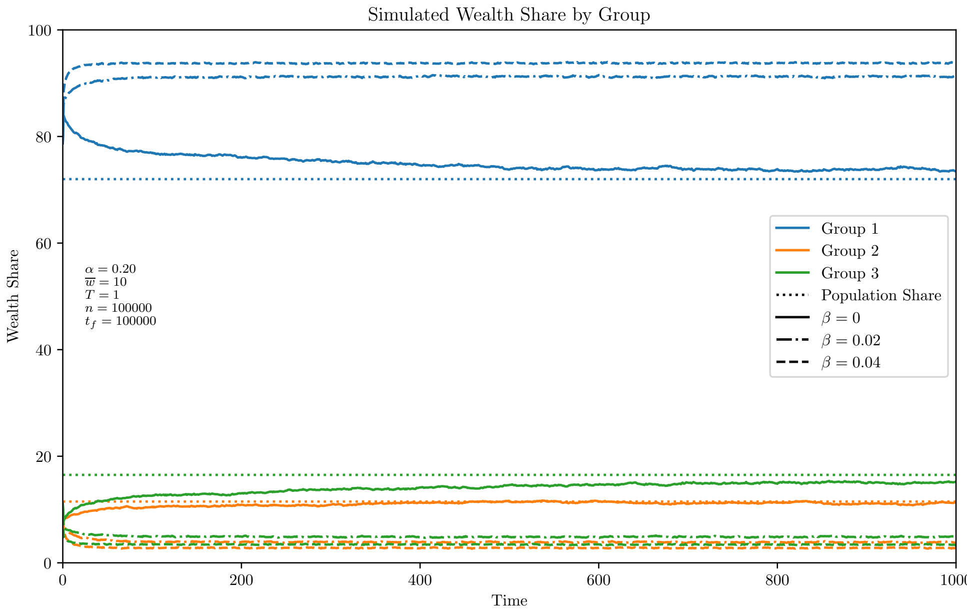

*At $t = 0$, agents in Group A has 2 more wealth than those in Group B and agents in Group B have 2 more wealth than agents in Group C. Otherwise, whichever group ends up dominant in the steady state can be unpredictable.

Data Fit

When group sizes are weighted to approximate the US racial distribution (Table 4),[4] $\alpha = 0.2$, and $\beta = 0.035$, the resulting distribution of wealth approximates that within the US.

Table 4

Race and group population shares

Racial Group	Simulation Group	Population Share
White (Non-Hispanic)	Group 1	0.720
Hispanic	Group 2	0.115
Black (Non-Hispanic)	Group 3	0.165

by the way the wealth share code deals with ties. In the case where wealth levels are equal, it puts Group 1 first, then 2, then 3. Weighting ties based on group shares at those values would likely get a smoother graph.

[4]Due to data constraints, all US data is normalized to only include White non-Hispanic, Black non-Hispanic, and Hispanic populations.

Figure 8

Lorenz Curve with group wealth shares overlayed. Shading represents the portion of that wealth share held by people in that group.

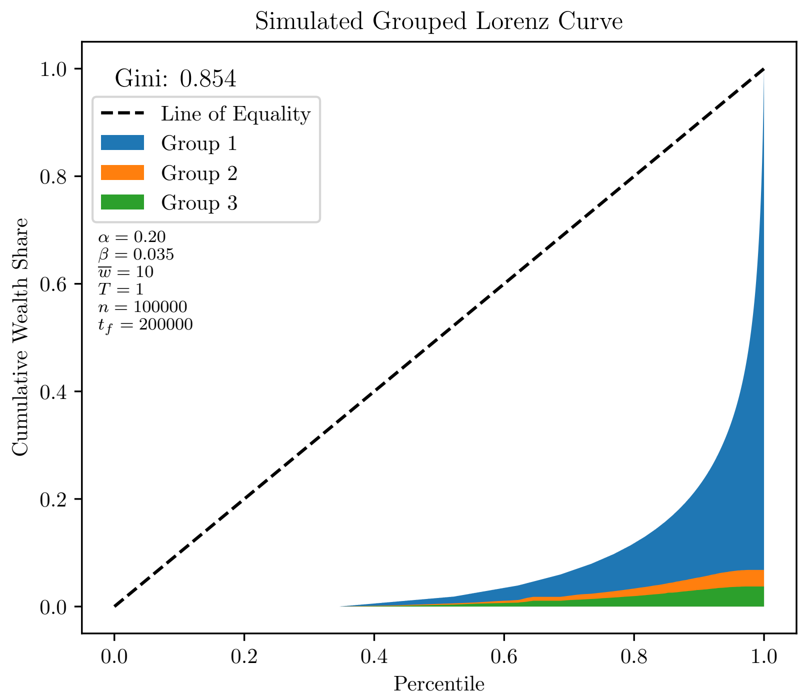

At an individual level, the resulting Lorenz Curve and Gini Coefficient is similar to that of the model presented in Section , so it resembles the real Lorenz Curve in a similar manner (Figure 9).[5]

At a group level, the simulation wealth shares also approximate that of US racial groups. Figure 10 shows that as time goes forward in the simulation, the steady state wealth distribution for groups approximates the real world.[6] This suggests that the addition of groups to the model allows the model to simultaneously approximate both real individual wealth dynamics and group wealth dynamics.

The ability to approximate grouped results is a direct result of the β parameter. Figure 7 shows that when $\beta = 0$, which is functionally identical to the basic model, the model converges to a proportional wealth distribution, which isn't what's observed in the real world.

At an individual level, however, the model is only moderately effective at replicating

[5]See Table A1 for the values and residuals numerically.
[6]See Table B1 for a numerical representation of this.

Figure 9

Simulated Lorenz Curve from the grouped model with real 2019 US points overlayed. Data from Aladangady & Forde (2021)

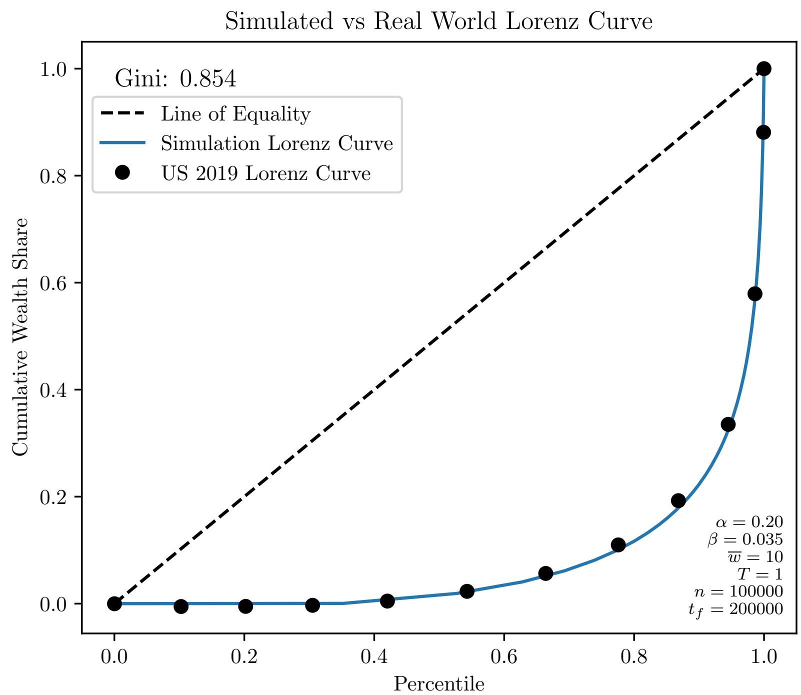

real phenomena. Figure 11 demonstrates that the plateauing behavior at the 95th percentile observed in the model doesn't exist in the real data. Instead, the real world data continues to smoothly increase its slope upward to the top percentile.

This effect is further demonstrated in Table 5. The R^2 for the model as a whole is still very high, though this could also be due to the fact that multiple groups are represented in the data meaning the mean is a very poor predictor. By group, the R^2, though always fairly high, gets substantially lower for groups 2 and 3, suggesting the model is a worse fit for these individuals.

Another way to evaluate the addition of parameters to a model is using the Akaike Information Criterion (AIC). The AIC value can be calculated using

$$\text{AIC} = n \ln\left(\frac{\text{RSS}}{n}\right) + 2k + 2$$

where RSS is the residual sum of squares, n is the number of points, and k is the number of parameters in the model. A lower AIC represents a better fitted model after accounting

Figure 10

Simulated wealth shares over time with US 2019 data overlayed. Data from Aladangady & Forde (2021)

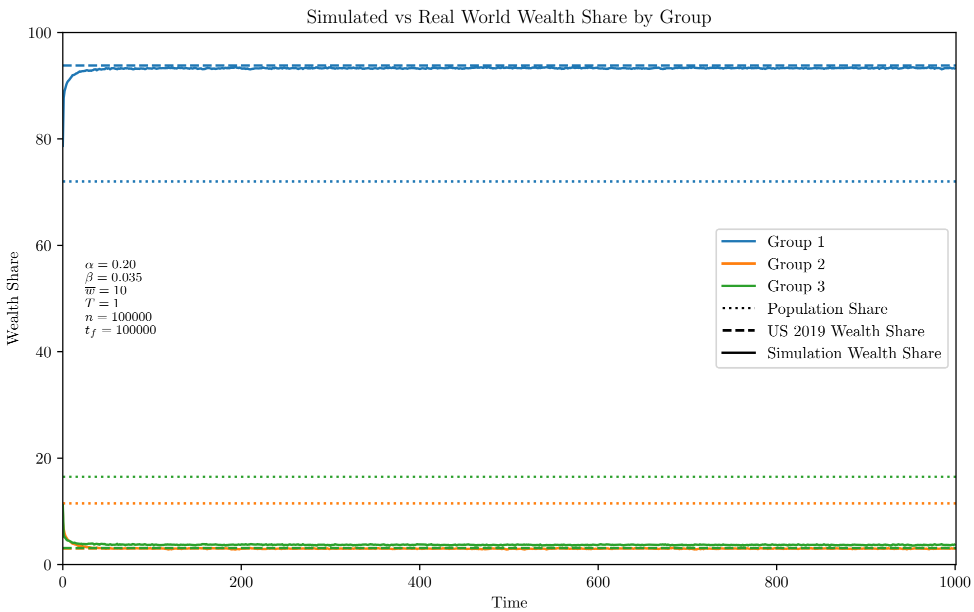

Table 5

R^2 values for the Figure 11. Each group is the R^2 just for that set of points, while the total R^2 is shown in the final row.

Group	$\mathbf{R^2}$
Group 1	0.998
Group 2	0.877
Group 3	0.805
Total	0.997

for the complexity added by a new parameter, so a model with a lower AIC is preferred Ledder (2023).

Based on the grouped data, the AIC of the basic model is -213.90 and the grouped model is -336.97, suggesting that despite the weaker fit at an individual level, especially for groups 2 and 3, the added parameter does better explain the real world data.

Conclusion

Overall, the models presented in this paper are able to successfully replicate parts of the wealth distribution of the United States. The basic model in Section closely ap-

Figure 11

Lorenz Curve with group wealth shares shaded overlayed with US 2019 data. Data from Aladangady & Forde (2021)

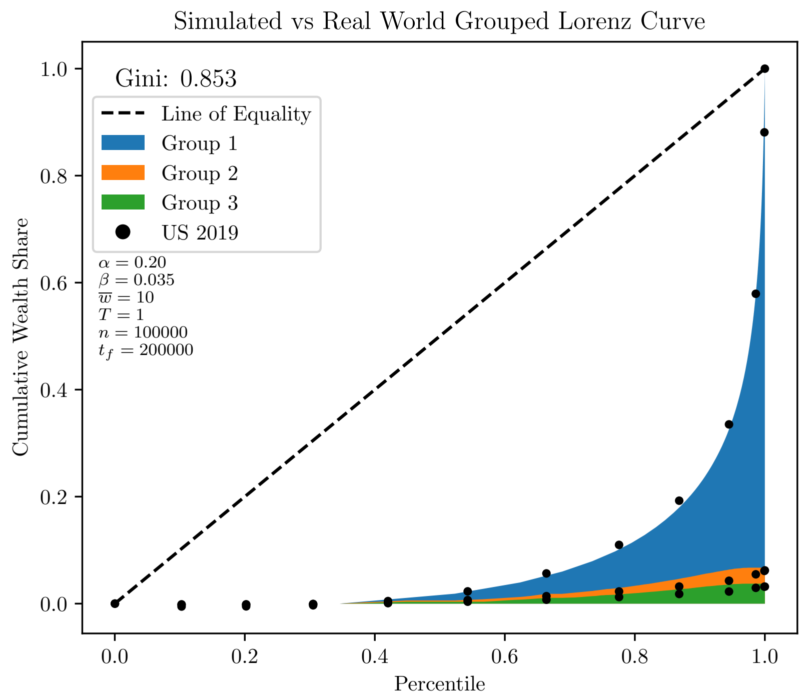

proximates the wealth distribution for individuals, shown by the Lorenz Curve and Gini Coefficient, of the US in 2019 and the grouped model in Section is able to approximate both the wealth distribution for individuals and racial groups, though struggles where those both interact.

As simple as they are, both models are remarkably accurate at simulating real world wealth dynamics, suggesting that more generally, wealth follows similar patterns where richer people experience advantages in the kinds of wealth transfers that occur every day.

The fact that the model didn't need to incorporate the relative wealth of agents[7] is especially interesting and may suggest that in the real world, class conflict can't be reduced to rich versus poor and includes complex interplays between every individual and the groups above and below them.

[7] and, in fact, has a worse fit when this is incorporated

Limitations

This paper presents very simple models for very complex phenomena. This has the advantage that it allows for better analysis, since what happens is clearer, avoids risks of overfitting Thomas et al. (2018), and requires less computational power, but does mean that many factors at play in the real world don't exist in the model. The model is perhaps an oversimplification of what it tries to represent.

Also, data limitations restricted the level of analysis that could be performed. Only 13 points along the real Lorenz Curve were used since that's what's available (Aladangady & Forde, 2021) and constructing a Lorenz Curve for better comparison would require access to restricted Survey of Consumer Finances (SCF) data (of Governors of the Federal Reserve Board, 2023).

Finally, this paper's analysis is only in the US, primarily due to difficulty finding Lorenz curve data for other countries. This reduces the context in which the model can be interpreted.

Further Research

Further research could work to improve the accuracy of Figure 11, potentially by going back to the standard YSM and modifying that to add group privilege instead of creating a whole new basic model or fitting better parameters, since the parameters presented in their paper were found within a only a handful of trial-and-error steps.

With access to SCF data, the accuracy of the model could be tested against millions of points along the Lorenz Curve, instead of the 13 in this paper. This would get better results regarding the accuracy and predictiveness of the model.

Finally, the model could be applied to a global scope and compared to other countries. Doing this would require data regarding the net worth of individuals in a country, but would allow for a more comprehensive analysis of the model. Similar to the SCF, however, this data in other countries doesn't allow public access (Authority, 2023).

References

Aladangady, A. & Forde, A. (2021). Wealth inequality and the racial wealth gap. *FEDS Notes. Washington: Board of Governors of the Federal Reserve System.*

Authority, F. C. (2023). Financial lives 2022 survey. *Financial Conduct Authority.*

Caiani, A., Godin, A., Caverzasi, E., Gallegati, M., Kinsella, S., & Stiglitz, J. E. (2016). Agent based-stock flow consistent macroeconomics: Towards a benchmark model. *Journal of Economic Dynamics and Control*, 69, 375–408.

Chorro, C. (2016). A simple probabilistic approach of the yard-sale model. *Statistics and Probability Letters*, 112, 35–40.

Dawid, H. & Gatti, D. D. (2018). Agent-based macroeconomics. *Handbook of computational economics*, 4, 63–156.

Dawid, H., Harting, P., van der Hoog, S., & Neugart, M. (2016). A heterogeneous agent macroeconomic model for policy evaluation: Improving transparency and reproducibility. pgs.

Dosi, G. & Roventini, A. (2019). More is different ... and complex!: The case for agent-based macroeconomics. *LEM Working Paper Series.*

Ledder, G. (2023). *Mathematical Modeling for Epidemiology and Ecology.* Springer Undergraduate Texts in Mathematics and Technology.

of Governors of the Federal Reserve Board, B. (2023). 2022 survey of consumer finances. *Board of Governors of the Federal Reserve System.*

Seppecher, P. (2012). Flexibility of wages and macroeconomic instability in an agent-based computational model with endogenous money. *Macroeconomic Dynamics*, 16(S2), 284–297.

Suisse, C. (2022). Global wealth databook 2022. *Credit Suisse Research Institute.*

Thomas, H., Mann, A., & Meschede, T. (2018). Race and location: The role neighborhoods play in family wealth and well-being. *American Journal of Economics and Sociology*, 77(3-4), 1077–1111.

Xu, S. (2022). Time-dependent analysis to yard-sale model of wealth distribution. *Diss. Tufts University.*

Appendices

Appendix A
Lorenz Curve Values Table

Table A1

US, Basic Model, and Grouped Model Lorenz Curve Values. US data from Aladangady & Forde (2021)

Population Share	US 2019	Basic Model	Grouped Model
0.0002	-0.0002	0.0000	0.0000
		(0.0002)	(0.0002)
0.1027	-0.0054	0.0000	0.0000
		(0.0054)	(0.0054)
0.2021	-0.0051	0.0000	0.0000
		(0.0051)	(0.0051)
0.3051	-0.0031	0.0000	0.0000
		(0.0031)	(0.0031)
0.4203	0.0048	0.0088	0.0078
		(0.0040)	(0.0030)
0.5429	0.0230	0.0249	0.0227
		(0.0019)	(-0.0003)
0.6640	0.0565	0.0551	0.0523
		(-0.0014)	(-0.0042)
0.7757	0.1098	0.1064	0.1015
		(-0.0034)	(-0.0084)
0.8682	0.1925	0.1860	0.1796
		(-0.0065)	(-0.01282)
0.9449	0.3351	0.3308	0.3258
		(-0.0043)	(-0.0093)
0.9861	0.5792	0.5842	0.5792
		(0.0050)	(-0.0001)
0.9993	0.8810	0.9439	0.9454
		(0.0629)	(0.0644)
1.0000	1.0000	1.0000	1.0000
		(0.0000)	(0.0000)
		$\alpha = 0.2$	$\alpha = 0.2$
		$\overline{w} = 10$	$\beta = 0.035$
		$T = 1$	$\overline{w} = 10$
		$n = 10000$	$T = 1$
		$t_f = 100000$	$n = 100000$
			$t_f = 100000$

Model difference from US 2019 values in parentheses

Appendix B
Race Wealth Shares Table

Table B1

US 2019 and Grouped Model wealth shares by race. US data from Aladangady & Forde (2021)

Racial Group	US 2019	Grouped Model
White (Non-Hispanic)	0.938	0.93375
		(-0.00425)
Hispanic	0.030	0.02995
		(-0.00005)
Black (Non-Hispanic)	0.031	0.03629
		(0.00529)
		$\alpha = 0.2$
		$\beta = 0.035$
		$\overline{w} = 10$
		$T = 1$
		$n = 100000$
		$t_f = 100000$

Model difference from US 2019 values in parentheses * US values shown post-normalization to get rid of the "other" category

Import Tariffs and Terms of Trade in the TNT Model: Addressing the US Current Account Deficit

Zoe Felsch

ECON 474: Open Economy Macroeconomics Capstone (Advisor: Mario Solis-Garcia)

A looming challenge facing the US at present is the economy's growing current account deficit. As theory implies this is unsustainable in the long-run (Schmitt-Grohé et al., 2022), economists and politicians alike question what policies might be enacted to address it, and if it can be mitigated at all.

For background, in the early 1980s the US switched from being a net external creditor to the world's largest external debtor. The US current account worsened to reach a figure of −\$647.2 billion in 2020, or −3.1 percent of the country's GDP, and it stands as the largest deficit in the world today. In theory, the question of the sustainability of a country's current account deficit depends on the economy's initial net international investment position; a perpetual deficit is only sustainable if the country begins the period as a net external creditor, which the US certainly is not. This theory outlined in the text of Schmitt-Grohé et al. (2022) therefore implies that the country's deficit is unsustainable. If gone unaddressed (if the deficit is not successfully matched with policy designed to improve the current account) we will meet economic catastrophe on a never-before-seen scale. Specifically, the US economy will meet a sudden stop, a forced current account reversal.

In an effort to tackle this among other goals, the Trump administration enacted a series of tariffs on thousands of US imports, with a focus on steel, aluminum, washing machines, and solar panel products. The first steps of this trade policy began in early 2018 with the implementation of a 25-percent tariff on steel imports and a 10-percent tariff on aluminum imports under Section 232 of the Trade Expansion Act. The scope of these particular tariffs and others like it were expanded upon over the course of the Trump administration, and many currently remain in place under the Biden administration. Trump's tariff policy is widely considered to be one of the largest tax increases in decades (York, 2023), and controversial at that. Not only are import tariffs generally disfavored by domestic consumers and firms, but their enactment also resulted in retaliatory tariffs

Zoe Felsch is a senior with majors in Economics and English. Correspondence concerning this article should be addressed to zfelsch@macalester.edu.

from key trading partners, like China, India, and Turkey. Interestingly enough, the former President outlined a plan in August of this year to expand this policy with what he called a "universal baseline tariff," should he be re-elected (Stein, 2023).

The theoretical logic behind import tariff policy is that cutting the economy's domestic consumption of foreign goods improves the trade balance and the current account (Schmitt-Grohé et al., 2022). This is unfortunately based on a simplistic understanding of the determinants of the current account; recent in-depth study concludes that the implementation of such policy has no effect on the current account. In particular, Schmitt-Grohé et al. (2022) use an intertemporal model of a small open economy called the TNT (Tradable Non-Tradable) model to draw said conclusion.

In this paper, I dive deeper and explicitly integrate the terms of trade and import tariffs into the TNT model in order to observe the impact of their movements on the current account. I view potential shocks to equilibrium in the model through a graphical analysis, then involve empirical evidence to allow for an analysis of the unique case of the US economy, and its challenges ahead. I finally conclude that my model reflects what I observe in the data, suggesting tariff policy does not and will not work to improve the current account, implying I must find an alternative approach to mitigating the risk of a sudden stop.

Literature Review

As one might expect, a wave of exploration and study of the determinants of the current account followed the aforementioned change in the US from surplus to deficit in the 1980s. In this paper, I take the heaviest influence from Sebastian Edwards (Edwards & Ostry (1987), Edwards (1987), Edwards (1988), Edwards (2004)), and from Schmitt-Grohé et al. (2022).

Engel & Kletzer (1986) examine the potential movements in saving and the current account in a small country using two models: the endogenous discount rate model developed by Uzawa (1968) and the uncertain lifetime model from Yaari (1965) and Blanchard (1985). Engel & Kletzer (1986) note the importance of using an infinite-horizon model, as it allows for the separation between short-run and long-run effects of policy, unlike a two-period model. The authors also stress the lesson that conclusions relating to policy drawn about the effects of tariffs can be incredibly sensitive to the structure of the model used. They conclude that the effect increased tariffs have on saving and the current account depend on a steady-state level of income and consumption through the distortion and expenditure effects. Generally, the Uzawa (1968) model and the Yaari (1965) model predict that tariffs incentivize saving, improving the current account.

Edwards (1987) uses a three-good, two-period general equilibrium model of a small open economy with utility-optimizing producers and consumers to analyze the determinants of the real exchange rate, and how internal shocks triggered by changes in import tariffs and the terms of trade impact the current account. The three goods that households consume are importables, exportables, and non-tradables. Edwards (1987) introduces temporary, permanent, and anticipated future movements in import tariffs to the model to allow for the observation of these changes on the equilibrium real exchange rate and the current account, which show dependence on intertemporal elasticities. This study sets the stage for that of Schmitt-Grohé et al. (2022) and influences how they view the current account as tied to the exchange rate, which I revisit later.

Edwards (1987) compares the effects of changes in the terms of trade with those in tariffs, using a model by Laursen & Metzler (1950). Specifically, Edwards (1987) shows that changes in import tariffs or the terms of trade affect the consumption rate of interest (CRI), intertemporal consumption decisions, the equilibrium real exchange rate and the current account. Edwards concludes that, in the general equilibrium model with free capital mobility, it is not possible to determine the effect temporary or permanent tariffs will have on the real exchange rate a priori. The effects of import tariffs on the current account are also undetermined because they depend on intertemporal price effects, initial expenditure on importables and non-tradables, and on the effect of the tariff on the real exchange rate.

Edwards & Ostry (1987) uses the same three-good, two-period general equilibrium model to look deeper into how today's real exchange rate and current account are affected by anticipated increases in future import tariffs. Edwards states it is likely that an anticipated tariff will worsen today's current account. This lesson has wider policy implications, as Edwards & Ostry (1987) writes that the mere discussion of the possibility of tariff policy by politicians will have adverse economic effects today.

The following year, Edwards (1988) narrows his focus to the impact that temporary terms of trade changes resulting from import tariffs might have on the real exchange rate and the current account. In this study, Edwards uses his three-good model to determine that there is a relationship between the terms of trade and the real exchange rate, but that the direction of movement in the equilibrium real exchange rate, and therefore the effect on the current account, cannot be determined by his model.

A key contribution in Edwards' study of the determinants of the current account is the precedent he sets in equating the real exchange rate and the relative price of non-tradables. This allows him to make the connection between movements in the real exchange rate and the current account. Schmitt-Grohé et al. (2022) build off Edwards' research, providing us an updated look at the relationship between the relevant variables through the scope of a two-good, intertemporal general equilibrium model of a small open economy with a government and utility-maximizing consumers. Schmitt-Grohé et al. (2022) suggest that a temporary fall in an economy's terms of trade caused by the imposition of import tariffs may lead to an improvement of the current account. They find that this effect can only take place under unique conditions, however, so tariffs are not a reliable policy tool for improving the current account.

Independent from analysis of the relationship between tariffs and the current account, Schmitt-Grohé et al. (2022) detail the theory behind sudden stops. They build off a foundation of understanding of the relationship between current account reversals and sudden stops outlined by Edwards (2004), who adds that restriction of capital mobility in the global financial market does not reduce the likelihood of a reversal, and that countries with more flexible exchange rate regimes are more prepared to accommodate shocks. Schmitt-Grohé et al. (2022) describe the hallmark economic consequences of a sudden stop—including a current account reversal, a contraction in aggregate demand, and depreciation of the real exchange rate—all of which they capture in their two-period, two-good TNT model.

Model and Theory

The Basics

I begin by specifying my definitions of the terms of trade, the trade balance, and the current account. A country's terms of trade[1] are the relative price of exports and imports, meaning the introduction of import tariffs to an economy is shown by a fall in the terms of trade. The trade balance[2] is the difference between the economy's endowment adjusted for the terms of trade and consumption, so there is a positive relationship between the terms of trade and the trade balance. I define a country's current account[3] as the sum of the trade balance and net investment income. Schmitt-Grohé et al. (2022) write that the trade balance and the current account move closely together, and I function under the same assumption in this paper.

Schmitt-Grohé et al. (2022) attest that a temporary fall in terms of trade may lead to an improvement in the current account. However, this conclusion is touch-and-go because it relies on caveats and interdependencies for it to hold. In other cases, like when the fall in the terms of trade is permanent, there is no effect on the current account. If there is an anticipated fall in future terms of trade, the trade balance and the current account worsen. It is also noteworthy that, in general, any tariff is theoretically welfare-decreasing.

[1] See Definition A1.
[2] See Definition A2.
[3] See Definition A3.

New Model

In this paper, I extend the TNT model presented by Schmitt-Grohé et al. (2022) to capture more direct effects of internal shocks to an economy induced by the imposition of import tariffs, specifically in relation to the economy's current account. The TNT model detailed by Schmitt-Grohé et al. (2022) is an intertemporal and intratemporal model of a small open economy with two goods; tradable goods can be imported and exported without limitation, and non-tradable goods must be produced and consumed domestically. The relative price used by Schmitt-Grohé et al. (2022) is that of non-tradables to tradables.

Through the lens of this primary model, import tariffs directly affect tradable goods (rather than non-tradables, for which the law of one price holds and the price remains unchanged). To distinguish between changes in consumption of imports and exports, however, I introduce a change to this two-good model to develop a new three-good model. This is similar to the general equilibrium models Edwards uses in his analyses, where the three goods are imports, exports, and non-tradables.

This model includes utility-optimizing consumers in addition to a government that collects tariffs and redistributes the revenue to consumers in the form of lump-sum transfers, denoted L_t. Households of consumers choose consumption of imports (C_t^M), exports (C_t^X), and non-tradables (C_t^N) based on their preferences and their budget constraints. I assume log preferences for households, and this is formally given in Definition A. Households have two budget constraints (for periods $t = 1, 2$), stating that income used equals income gained. The budget constraints for periods one and two are given below in Equations 1 and 2.

$$t = 1: \quad P_1^M C_1^M (1 + \tau_1) + P_1^N C_1^N + B_1 P_1^M = P_1^X Q_1^X + L_1 + B_0(1 + r_0) \tag{1}$$

$$t = 2: \quad P_2^M C_2^M (1 + \tau_2) + P_2^N C_2^N = P_2^X Q_2^X + L_2 + B_1 P_2^M (1 + r_1) \tag{2}$$

In this three-good model, consumption takes the form of a Cobb-Douglas function. It is a composite of imports and non-tradables, given by Definition A6. Households are endowed with (or produce) an output of Q_t^X. Tariffs are denoted τ_t, terms of trade are TT_t, and nominal prices of imports, exports, and non-tradables are P_t^M, P_t^X, P_t^N, respectively. For simplicity, I assume the economy enters the first period as neither a debtor or creditor, with $B_0 = 0$. I impose the transversality condition ($B_2 = 0$) because the economy cannot end the second period with foreign debt or credit extended.

The budget constraints illustrate the households' uses of income on the left-hand sides and the sources of income on the right. Specifically, $P_t^M C_t^M (1+\tau_t)$ stands for consumption of imports, adjusted for the price and tariff. $P_t^N C_t^N$ is consumption of non-tradables, adjusted for the price, and $B_1 P_1^M$ is any spending in the international capital market in period t=1, adjusted for the price of imports. $P_t^X Q_t^X$ is the income from exports, L_t is income from government lump-sum transfers to households, and the final terms on the right-hand sides are income received on the international capital market.

Dividing through Equations 1 and 2 by P_t^M to get the budget constraints in terms of imports, our item of focus, results in the new budget constraints (Equations 3 and 4).

$$t=1: \quad C_1^M(1+\tau_1)+p_1C_1^N+B_1=TT_1Q_1^X+\frac{L_1}{P_1^M}+\frac{B_0(1+r_0)}{P_1^M} \tag{3}$$

$$t=2: \quad C_2^M(1+\tau_2)+p_2C_2^N=TT_2Q_2^X+\frac{L_2}{P_2^M}+B_1(1+r_1) \tag{4}$$

This brings to light a new variable of interest–the relative price[4] of non-tradables to imports, denoted p_t. The relative price in our model describes the relationship between the price of imports and the price of non-tradables.

Merging the budget constraints (Equations 3 and 4) for both periods via B_1, I obtain the intertemporal budget constraint (IBC) (Equation 5), which describes the household's sources and uses of income over both periods.

$$TT_1Q_1^X+\frac{L_1+B_0(1+r_0)}{P_1^M}+\frac{\frac{TT_2Q_2^X+L_2}{P_2^M}}{1+r_1}=C_1^M(1+\tau_1)+p_1C_1^N+\frac{C_2^M(1+\tau_2)+p_2C_2^N}{1+r_1} \tag{5}$$

The left-hand side of the IBC stands for the household's sources of income over both periods, or lifetime wealth. It includes the economy's endowments of exports, adjusted for terms of trade, lump-sum transfers from the government and bond income, adjusted for the price of imports, and the interest rate. The right-hand side of Equation 5 stands for the household's uses of income, which includes consumption of imports, adjusted for tariffs, consumption of non-tradables, adjusted for the relative price, and the interest rate.

To ascertain the conditions in which the household maximizes utility subject to the IBC, I use the LaGrangian to identify the first-order conditions with respect to the exogenous variables (C_t^M and C_t^N for $t=1,2$). This process provides a system of three key conditions: the intertemporal condition (Equation 6), and two intratemporal conditions (Equations 7 and 8). The intertemporal condition, or Euler equation, describes how households decide to consume imports between periods. The intratemporal conditions describe how households decide to consume imports and non-tradables within each period. For example, in period one, the household optimizes utility by choosing import consumption (C_1^M) and non-tradable consumption (C_1^N) based on Equation 7.

$$C_2^M=\frac{\beta C_1^M(1+\tau_1)(1+r_1)}{1+\tau_2} \tag{6}$$

[4]See Definition A4.

$$t = 1: \quad C_1^M = \frac{\gamma}{1-\gamma} \cdot \frac{C_1^N p_1}{1+\tau_1} \tag{7}$$

$$t = 2: \quad C_2^M = \frac{\gamma}{1-\gamma} \cdot \frac{C_2^N p_2}{1+\tau_2} \tag{8}$$

The Euler equation (Equation 6), tells us that consumption of imports tomorrow (time $t = 2$) depends on the level of consumption of imports today (time $t = 1$), the interest rate, and tariffs in both periods. $\beta \in (0, 1)$ is the subjective discount factor taken from the household's utility function[5] and describes the household's level of patience. For example, if $\beta = 0$, households have the lowest level of patience and only care about consumption today.

Equilibrium

With this framework in-hand, I can begin to think about equilibrium in this three-good TNT model. Equilibrium is the consumption decision (C_t^M, C_t^X, C_t^N) made by households that satisfies five items: the intertemporal budget constraint (Equation 5), the Euler equation (Equation 6), the interest rate parity condition, the household's utility[6], and the implication of free capital mobility,[7]

In equilibrium, the market for non-tradable goods clears, as all non-tradables are produced and consumed domestically. Formally, this means $C_t^N = Q_t^N$ for $t = 1, 2$. I can use this to eliminate C_1^N and C_2^N from the IBC (Equation 5), and obtain the economy's intertemporal resource constraint (IRC) (Equation 9):

$$TT_1 Q_1^X + \frac{L_1}{P_1^M} + \frac{TT_2 Q_2^X + \frac{L_2}{P_2^M} - p_1 Q_1^N}{(1+r^*) - p_1 Q_1^N} = C_1^M(1+\tau_1) + \frac{C_2^M(1+\tau_2)}{(1+r^*)} \tag{9}$$

I then use the Euler equation (Equation 6) to solve for C_2^M to find the equilibrium level of consumption of imports in period one. I call the left-hand side of equation 9 "$\underline{Y}$".

$$C_1^M = \frac{\underline{Y}}{(1+\beta)(1+\tau_1)} \tag{10}$$

With this framework of equilibrium in hand, I use it to examine the effects of internal shocks to the economy, specifically import tariffs.

[5]See Definition A.
[6]See Definition A.
[7]Under free capital mobility, a small open economy's interest rate today is the same as the world interest rate (see Definition A7), namely, $r_t = r^*$.

Analysis

The application of this 3-good model with regards too import tariffs can be seen through evaluation of the Euler equation (Equation 6). Consumption of imports in period two (C_2^M) positively depends on consumption of imports in period one C_1^M, the interest rate r_1, and today's tariffs τ_1, while it depends negatively on tariffs tomorrow τ_2. From this and the intratemporal conditions (Equations 7 and 8), we can glean that consumption of imports C_t^M depends positively on the consumption of non-tradables C_t^N and the relative price p_t for the respective period, while it depends negatively on tariffs tomorrow τ_2. Remember that the relative price in our model is the nominal price of non-tradables divided by the nominal price of imports, meaning an increase in tariffs would manifest as a lower relative price.

From Equation 10 I know that if $\beta = 1$ and $\tau_1 = 0$, consumption of imports in period one (C_1^M) are half the value of lifetime wealth ($\underline{Y}$). I use this to consider the effect on the trade balance and the current account. From Equation 10, I know that the terms of trade, the endowment of exports and non-tradables for the respective period have a positive relationship with import consumption in period one. I also know that the relative price in the respective period, as well as tariffs and the interest rate in period one, have negative relationships with import consumption in period one. This is shown in Equation 11:

$$\begin{aligned} C_1^M = C^M(TT_t, Q_t^X, Q_t^N, p_t, \tau_1, r^*) \\ (+)(+)(+)(-)(-)(-) \end{aligned} \tag{11}$$

The trade balance is the difference between tradables (composed of imports and exports), adjusted for terms of trade and the consumption of imports. I use this identity to determine the effects to the trade balance when there are movements in the independent variables (see Equations 12, 13, 14).

$$TB_1 = TT_1(Q_1^M + Q_1^X) - C_1^M \tag{12}$$

$$TB_1 = TT_1(Q_1^M + Q_1^X) - C^M(TT_t, Q_t^X, Q_t^N, p_t, \tau_1, r^*) \tag{13}$$

$$\begin{aligned} TB_1 = TB(TT_t, Q_t^X, Q_t^N, p_t, \tau_1, r_1) \\ (-)(+)(-)(+)(+)(+) \end{aligned} \tag{14}$$

As stated earlier, I believe an economy's trade balance and current account move closely together. The movements the model predicts regarding the trade balance listed above are therefore the lessons I will take about the determinants of the current account.

A Graphical Approach

We support the analysis above by modeling the effects of internal shocks to the current account in a series of graphs similar to those from Schmitt-Grohé et al. (2022). The shocks include a temporary increase, a permanent increase, and an anticipated future increase in import tariffs. I then take a closer look at the potential scenario that may come from the specific policy former President Trump proposes enacting, should he be re-elected.

A Temporary Increase in Tariffs

Figure 1 exhibits equilibrium in the 3-good model, specifically, the consumption path the utility-optimizing households choose (C_1^M,C_2^M), as determined by the point of tangency of the household's indifference curve that describes their utility[8], and downward-sloping line given by the economy-wide resource constraint (Equation 9). The slope of the IRC is the marginal rate of substitution between consumption of imports in periods one and two, given by the Euler equation (Equation 6).

Figure 1

Response of Intertemporal Import Consumption to Tariffs

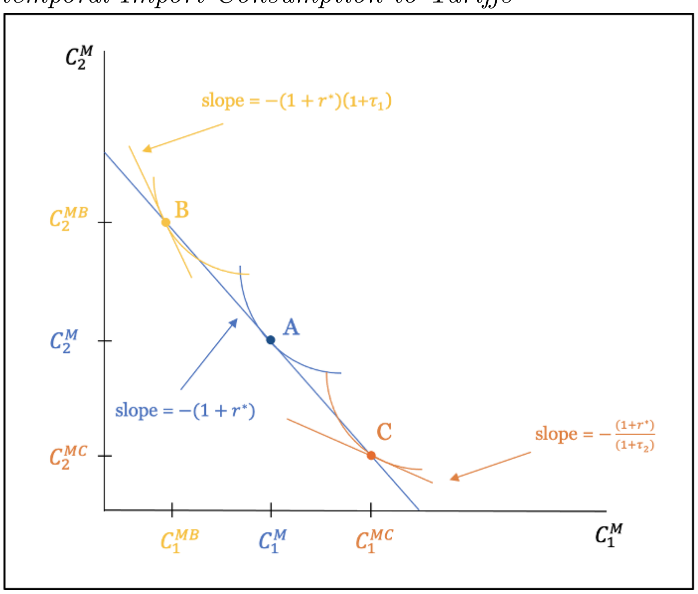

Assuming tariffs are initially $\tau_1 = \tau_2 = 0$, a temporary increase in tariffs would follow the case where $\tau_1 > 0$,$\tau_2 = 0$. The initial consumption path at Point "A" moves along the IRC to Point "B," as households shift consumption of imports from today ($t = 1$) toward tomorrow ($t = 2$). The slope of the point of tangency with the indifference curve changes

[8]See Definition A.

from $-(1+r^*)$ to $-(1+r^*)(1+\tau_1)$, showing a change in consumption from (C_1^M, C_2^M) to $(C_1^{M,A}, C_2^{M,A})$.

This increase in consumption of imports, when weighing the definition of the trade balance[9] and the current account[10], corresponds to a temporary improvement in the trade balance and the current account.

A Permanent Increase in Tariffs

A permanent increase in tariffs has no meaningful effect on the current account since it does not significantly affect the household's consumption of imports. I form this conclusion from what the Euler equation (Equation 6) signifies: the slope of the line of tangency at the IRC line and the household's new indifference curve does not change following the shock. Assuming initial tariffs are $\tau_1 = \tau_2 = 0$, and the shock of a permanent increase in tariffs is described $\tau_1 = \tau_2 > 0$, the slope of the tangent line is $-\frac{(1+r^*)(1+\tau_1)}{1+\tau_2}$. This reveals that a permanent increase does not have a significant effect on consumption, and therefore the current account.

An Anticipated Future Increase in Tariffs – Part A

An anticipated increase in tariffs, beginning with $\tau_1 = \tau_2 = 0$, follows $\tau_1 = 0, \tau_2 > 0$. The slope of the indifference curve changes from $-(1+r^*)$ to $-\frac{1+r^*}{1+\tau_2}$. This means a change in the consumption path in Figure 1 from Point "A" at (C_1^M, C_2^M) to Point "C" at $(C_1^{M,C}, C_2^{M,C})$. This change shows the household's shift in consumption away from tomorrow toward today. This fall in consumption of imports represents a worsening of the trade balance and the current account.

An Anticipated Future Increase in Tariffs – Part B

Now I observe what the model describes in the event reflective of Trump's proposed future tariff policy. In this case, the shock I model is an anticipated increase in future tariffs, but I begin not from $\tau_1 = \tau_2 = 0$ like in Part A, but from $\tau_1^{,} = \tau_2^{,} > 0$. The shock illustrated in this case, therefore follows $\tau_1^{,} > 0$, $\tau_2^{,} > 1$. Figure 2 illustrates this movement in the consumption path of imports.

The movement in Figure 2 from Point "C" to Point "D" reflects the change in the household's utility-maximized consumption path. The outcome of an anticipated additional increase in future tariffs adjusts the slope of the line tangent to the household's indifference curve, shifting intertemporal consumption preferences further from tomorrow toward consumption today. This shift further worsens the trade balance and the current account,

[9] See Definition A2.
[10] See Definition A3.

Figure 2

A Further Increase in Anticipated Future Tariffs

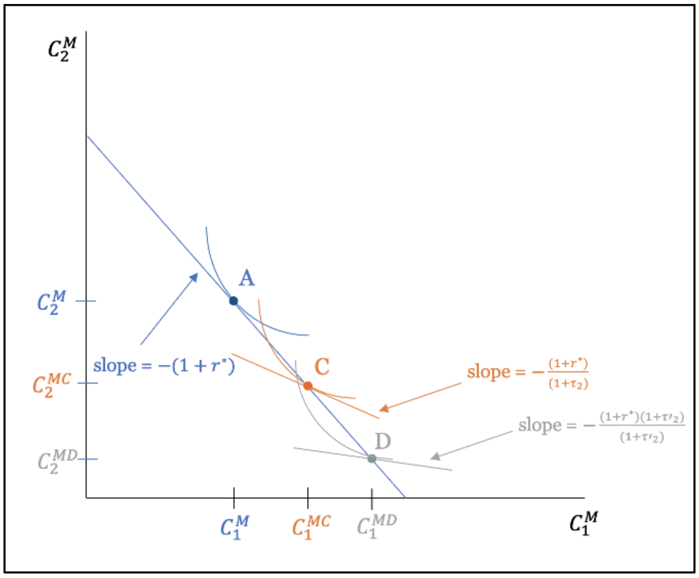

implying Trump's policy plan for a universal baseline tariff would be ineffective in addressing the US current account deficit.

An Empirical Assessment

From this model analysis, I draw conclusions about the impact of import tariffs on the trade balance and the current account. In general, tariffs have a positive relationship with the trade balance and the current account; an increase in tariffs increases the current account, thereby worsening the deficit.

This conclusion is supported by analysis of empirical evidence. According to the Bureau of Economic Analysis, the US current account stood at −\$212.1 billion in the second quarter of 2023 (BEA, 2023b). They list the figure at −\$96.9 billion when Trump enacted his first tariff in the first quarter of 2018 (BEA, 2023a). This indicates that since the Trump administration imposed the first steps of tariff policy, the current account has worsened by more than 100 percent. I conclude that basic empirical data suggests tariff policy is not a sufficient measure for addressing the deficit.

Limitations

In order to build our model and make analyses, I make three major simplifying assumptions. The first is integral to the structure of the model; a basic limitation of the TNT model is that it is representative of a two-period economy. The authors of one of the first working papers I review, Engel & Kletzer (1986), state that a two-period model cannot distinguish between short-run and long-run effects.

The most notable limitation to our analysis, however, is that I assume the economy in question is small. The impetus behind this study is to apply our framework to the United States economy, which is the largest in the world (by GDP). The way this is reflected in the model is through the interest rate r_1: large open economies, or a group of small open economies, often sway movements in the world interest rate. Small open economies, on the other hand, take the world interest rate as given.[11] In this model, I assume the economy in question takes r_1 as given to streamline our process.

The final simplifying assumption I make is that the household's utility is described by a function of log preferences. The reason I make this assumption is for the sake of completion of algebraic processes without the use of a computer. This statement of assumptions made and limitations to our study are equally as important as the conclusions I draw, as they highlight areas for improvement and potential next steps in further discussion in this area.

Concluding Remarks

In this paper, I take a thoughtful approach to evaluating the potential efficacy of import tariff policy as a tool for mitigating the growing current account deficit in the United States. Spurred by discussion around the economic policy put in place by the Trump administration, I begin by reviewing previous research done on the subject of tariffs, terms of trade, and the current account. I pinpoint a model of interest first used by Sebastian Edwards in the 1980's, and tweaked by Schmitt-Grohé et al. (2022).

I take the two-good model from Schmitt-Grohé et al. (2022) and adapt it into a three-good model in order to capture the direct theoretical effects of internal shocks to a small open economy caused by an increase in tariffs. My model points to a similar positive relationship between import tariffs and the current account, which aligns with our conclusions drawn from both a graphical analysis and the empirical data observed today.

In regard to Donald Trump's outlined plan for a "universal baseline tariff" to be imposed on all imports as a means of extending his protectionist policy and improving the US current account, economist and professor Menzie Chinn adds value to the discussion. Chinn (2017) views the current account as the difference between national saving (as a sum of government budget surplus and private saving) and investment, rather than our take of it as the sum of the trade balance and net investment income. Even using an alternative definition, he accurately predicted the effect Trump's import tariffs would have. Chinn states his expectation for the failure of the President's policy to reduce the deficit while

[11] See Definition A7.

maintaining growth. He supports his argument four years later with data accumulated over Trump's term to conclude just what I have in this paper.

What does this mean for the United States economy? It means I might rule out tariff policy as a powerful tool for addressing the country's amassing external deficit, and that economic researchers and policymakers must continue on this path of study. Theory outlined by Schmitt-Grohé et al. (2022) warns us about the unsustainability of a deficit and the perils one poses to not only the US economy, but the world economy.

References

BEA (2023a). *U.S. Current-Account Balance.* https://www.bea.gov/sites/default/files/2023-09/trans223-quarterly-current-account-balance.pdf. publisher: Bureau of Economic Analysis

BEA (2023b). *U.S. International Transactions, 2nd Quarter 2023 | U.S. Bureau of Economic Analysis (BEA).* https://www.bea.gov/news/2023/us-international-transactions-2nd-quarter-2023

Blanchard, O. J. (1985). Debt, Deficits, and Finite Horizons. *Journal of Political Economy*, 93(2), 223–247. https://doi.org/10.1086/261297

Chinn, M. (2017). *So You Want to Fix the Trade Deficit? | Econbrowser.* https://econbrowser.com/archives/2017/07/so-you-want-to-fix-the-trade-deficit

Edwards, S. (1987). *Tariffs, Terms of Trade, and the Real Exchange Rate in an Intertemporal Optimizing Model of the Current Account.* https://doi.org/10.3386/w2175

Edwards, S. (1988). *Temporary Terms of Trade Disturbances, The Real Exchange Rate and the Current Account.* https://doi.org/10.3386/w2629

Edwards, S. (2004). *Financial Openness, Sudden Stops and Current Account Reversals.* https://doi.org/10.3386/w10277

Edwards, S. & Ostry, J. D. (1987). *Anticipated Protectionist Policies, Real Exchange Rates and the Current Account.* https://doi.org/10.3386/w2214

Engel, C. & Kletzer, K. (1986). *Tariffs, Saving and the Current Account.* https://doi.org/10.3386/w1869

Laursen, S. & Metzler, L. A. (1950). Flexible Exchange Rates and the Theory of Employment. *The Review of Economics and Statistics*, 32(4), 281. https://doi.org/10.2307/1925577

Schmitt-Grohé, S., Uribe, M., & Woodford, M. (2022). *International Macroeconomics: A Modern Approach.* Princeton University Press.

Stein, J. (2023). Trump vows massive new tariffs if elected, risking global economic war. *Washington Post.* https://www.washingtonpost.com/business/2023/08/22/trump-trade-tariffs/

Uzawa, H. (1968). Market Allocation and Optimum Growth. *Australian Economic Papers*, 7(10), 17–27. https://doi.org/10.1111/j.1467-8454.1968.tb00149.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-8454.1968.tb00149.x

Yaari, M. E. (1965). Uncertain Lifetime, Life Insurance, and the Theory of the Consumer. *The Review of Economic Studies*, 32(2), 137–150. https://doi.org/10.2307/2296058

York, E. (2023). *Tracking the Economic Impact of Tariffs*. https://taxfoundation.org/research/all/federal/tariffs-trump-trade-war/

Appendix

Definitions

Terms of trade:

$$t = 1, 2: \quad TT_t = \frac{P_t^X}{P_t^M} \tag{A1}$$

Trade balance:

$$t = 1, 2: \quad TB_t = TT_t Q_t - C_t \tag{A2}$$

Current account:

$$t = 1, 2: \quad CA_t = TB_t + r_{t-1} B_{t-1} \tag{A3}$$

Relative price:

$$t = 1, 2 \quad p_t = \frac{P_t^N}{P_t^M} \tag{A4}$$

Household utility:

$$t = 1, 2 \quad U(C_t) = \ln(C_1) + \beta \ln(C_2) \tag{A5}$$

Consumption:[12]

$$t = 1, 2: \quad C_t = (C_t^M)^\gamma (C_t^N)^{1-\gamma} \tag{A6}$$

Free capital mobility:[13]

$$t = 1, 2: \quad r_t = r^* \tag{A7}$$

[12] $\gamma \in (0, 1)$ is a parameter describing the consumer's relative utility between imports and non-tradables.
[13] A small open economies' domestic interest rate, r_t for $t = 1, 2$, is equal to the world interest rate, r^*.

How do Fluctuations in Crude Oil Prices Differentially Influence the Stock Prices of Budget Versus Full-Service U.S. Airlines?

Tanya Nangpal

ECON 381: Introduction to Econometrics (Advisor: Amy Damon)

According to the International Air Transport Association (IATA), the global airline industry's fuel bill reached USD 271 billion in 2023, accounting for a significant 32% of its operating expenses. Additionally, the aviation industry represents 7.8% of final oil consumption worldwide (Planete-Energies, nd). Increasing fuel costs can limit the supply of flights or increase the overall price level for consumers (McKinsey, nd). Given budget airlines like Southwest and JetBlue focus on cost efficiency, they might be more vulnerable to fuel price changes than full-service airlines such as Delta, American, and United Airlines. Budget airlines' leaner operating models could expose them more to the direct impacts of rising oil prices, a scenario that has not been researched yet.

Consistent with studies examining the relationship between daily stock and oil price movements, this paper uses weekly prices of airline stocks and WTI oil prices while controlling for the performance of the market through the S&P 500 weekly close prices. The S&P 500 prices capture broader market trends that could influence individual stock prices. It also provides a benchmark to compare the performance of airline stocks against the overall market. Following the market efficiency theorem, which states that share prices reflect all publicly available information, we assume that airline-specific variables such as the effect of airline load factor, available seat miles, and revenue passenger miles are held by the stock price, as these are factors that analysts account for while determining bid and ask prices.

Prior research examined the impact of fuel price shocks on six major U.S. airlines, uncovering that shocks led to fluctuations in airline stock returns. Specifically, American Airlines, Delta Air Lines, United Airlines, and US Airways showed a significant negative relationship between stock returns and fuel price shocks. Hsu's 2017 study revealed that fuel price increases affected airline stock returns, unlike periods of fuel price decreases. However, Nandha et al. 2013 discovered that oil price volatility alone did not notably influence airline stock returns and risk. Yet, when combined with oil regimes such as laws, regulations, and

Tanya Nagpal is a sophomore with majors in Economics and Computer Science. Correspondence concerning this article should be addressed to tnangpal@macalester.edu.

agreements that govern the economic benefits from petroleum exploration and production, it significantly impacted the stock risk of various countries.

This paper aims to address a gap in current research on how oil price fluctuations affect budget versus traditional airlines differently. It explores the sensitivity of both airline types to fuel price volatility, considering their distinct strategies for managing fuel price risks and operational costs.

Our key findings reveal that contrary to the prevalent view, full-service airlines exhibit a stronger sensitivity to oil price fluctuations than budget airlines. The robustness of these findings is supported through diagnostic tests, using robust standard errors and alternative model specifications.

This paper is divided into six sections. First, we look at existing research to understand what has been established about oil price fluctuations, oil prices in relation to stock prices, and airlines and their stock prices in particular. Next, we focus on the economic theory of our research question, notably consumer price sensitivity and the different business models and operational strategies of the two types of airlines. The subsequent section dives deeper into our dataset to understand where data were sourced from and the different variables used in our empirical analysis. The empirical analysis section outlines the regressions and models used, including fixed effects, robustness checks, loggs, laggs and first differences. Following this, we present and discuss our findings, emphasizing their implications both from economic and statistical perspectives. The paper concludes by acknowledging its limitations and suggesting potential directions for future research.

Literature Review

Research on oil dynamics centers on two primary perspectives: market fundamentals, such as supply and demand dynamics, and speculative activities in oil futures markets. Hamilton 2009 links significant oil shocks, like the one in 2007-2008, to fundamental market changes where rising demand outstripped stagnant supply. On the other hand, some studies argue that speculation in oil futures significantly influences price volatility (Upadhyay, 2016), although recent trends suggest a diminishing impact of speculative activities as market fundamentals regain prominence (Hassani et al., 2017).

As jet fuel prices continue to fluctuate, airlines are increasingly focusing on fuel efficiency as a critical operational strategy. This emphasis is transforming the aircraft market as airlines opt for newer, more efficient models and retire older ones like the Boeing 747, hastened by stringent environmental regulations targeting carbon emissions (Boeing CDP Climate Response). Environmental regulations aimed at reducing carbon emissions significantly impact aircraft market values, as Vasigh et al. 2021 and Girardet & Spinler 2013 highlighted. During the 77th IATA Annual General Meeting in October 2021, member airlines (including full-service airlines such as United, American, Delta, Alaska, and Hawaiian) of IATA adopted a resolution to reach net-zero carbon emissions from their operations by the year 2050 (IATA, nd).

Two key factors significantly affect airline stock prices. First, exchange rate fluctuations— especially for airlines with international routes where ticket sales involve currency exchanges— influence profit margins (Tsai, 2008). Additionally, higher oil prices increase operational costs, leading to lower profits, reduced cash flow in the market, and negatively impacted stock performance (Alıcı, 2024).

Many studies have focused on the relationships between oil price shocks and stock market returns. Theoretically, such shocks influence stock returns by altering anticipated earnings, as Jones 2004 describes. Empirical investigations, including those by Park and Ratti 2008, have found a generally positive correlation between rising oil prices and stock market returns across 13 European nations and the United States. However, Apergis and Miller 2009 found that stock markets in developed economies showed negligible reactions to fluctuations in oil prices. Such broader studies on oil prices and stock markets show varied effects, but this analysis focuses on the airline industry, particularly distinguishing between budget and full-service airlines.

Airlines commonly hedge against rising oil prices by using forward contracts, which are agreements to purchase or sell a specified quantity of an asset at a predetermined price on a future date, ensuring stable fuel costs (Memon, 2019). Airlines secure fuel at competitive prices, even from distressed competitors, to stabilize costs during periods of high fuel prices (Carter et al., 2003). This practice not only preserves profit margins but also enhances the airline's valuation, as supported by Allayannis and Weston 2001. Furthermore, consistent hedging improves an airline's financial stability and attractiveness to investors, offering them greater opportunities for investment, as indicated by studies from Froot et al. 1993 and Géczy et al. 1997. Airline hedging decisions are confidential and important to their cost optimization strategies.

Economics of the Problem

Consumer price sensitivity drives the segmentation and pricing strategies within the airline industry. Budget airlines focus on cost-sensitive consumers by using static pricing models prioritizing affordability over convenience. These airlines attract flexible travelers willing to compromise on layovers and flight timings for lower fares. Conversely, full-service airlines cater to less price-sensitive passengers by offering premium services and more flexible travel options, charging higher prices for direct flights and optimal travel times.

Budget airlines often employ cost-reduction strategies, such as operating a single aircraft model to streamline maintenance and training costs. While cost-effective, this approach limits their operational flexibility compared to full-service carriers with diverse fleets capable of servicing a wider range of destinations. For example, Southwest Airlines exclusively uses the Boeing 737 series, optimizing cost efficiency but restricting service capabilities (Wikipedia, 2024).

Additionally, budget airlines tend to utilize secondary airports, which typically charge lower landing and ground fees than major hubs, further reducing operational costs. They also rely heavily on ancillary revenue streams, such as baggage fees and onboard services, to compensate for lower ticket prices. In contrast, full-service airlines derive revenue from various sources, including cargo services, loyalty programs, and subsidiary investments (Garbuno, 2020).

Data Description

Closing stock prices for selected budget and full-service airlines, alongside the S&P 500 index closing prices, are sourced from Yahoo Finance. Weekly prices of West Texas Intermediate (WTI) crude oil, a key benchmark for North American oil pricing, are obtained from the Federal Reserve Economic Data (FRED) database. Airlines are classified into two groups for analysis: Budget Airlines (including Southwest, JetBlue, Spirit, and Allegiant) and Full-Service Airlines (comprising United, American, Delta, Alaska, and Hawaiian). The dataset consists of weekly observations every Friday from May 2011 to November 2023, totaling 5,886 data points.

Table 1

Descriptive Statistics

Variables	Obs	Mean	Std. Dev.	Min	Max
Stock Price	5886	45.198	38.629	3.48	257.74
WTI Crude Oil Price	5886	70.297	22.628	8.252	120.932
S&P 500 Close Price	5886	2672.59	1016.166	1123.53	4766.18
Log Stock Price	5886	3.488	.836	1.247	5.552
Log WTI Crude Oil Price	5886	4.194	.358	2.11	4.795
Log S&P 500 Close Price	5886	7.817	.388	7.024	8.469
Full Service	5886	.556	.497	0	1
Lag WTI Crude Oil Price	5769	70.362	22.677	8.252	120.932
Lag S&P 500 Close Price	5769	2670.173	1014.126	1123.53	4725.79
First Diff Stock Price	5769	.02	3.055	-44.49	32
First Diff WTI Crude Oil Price	5769	-.046	2.782	-15.19	18.458
First Diff S&P 500 Close Price	5769	4.874	67.29	-406.1	301.17

'Airline' and 'Airline Code' identify the airline companies by name and a numeric code from 1 to 9, respectively. 'Stock Price' reflects weekly stock prices, and 'Full Service' is a binary variable indicating whether an airline is budget (0) or full-service (1). Prices of 'WTI Crude Oil' and 'S&P 500 Close' are recorded weekly. An 'Equally Weighted Average of Budget and Full-Service Airlines' is computed to analyze the collective behavior of different airline types. Transformations like logs and lags are applied to stock prices, WTI crude oil prices, and S&P 500 close prices. The changes in these variables from one week to the next are captured in the first difference variables to examine week-over-week fluctuations.

Table 2

Two-Sample T-test with Equal Variances

Variables	Obs1	Obs2	Mean1	Mean2	Diff	Std. Error	T Value	P Value
Stock Price by Full Service	2616	3270	54.379	37.853	16.526	.99	16.7	0

Notes: Obs1 and Mean1 correspond to Full Service = 0 (budget airlines), Obs2 and Mean2 correspond to Full Service = 1 (full-service airline).

The t-statistic of 16.7 indicates a very strong difference between the groups. The corresponding p-value is less than 0.0001, indicating that the difference is statistically significant at any conventional significance level. The test results strongly reject the null hypothesis that there is no difference between the stock prices of budget and full-service airlines. The positive mean difference indicates that budget airlines have significantly higher stock prices ($ 54.37) than full-service airlines ($ 37.85).

The graph illustrates average log-transformed airline stock fluctuations and crude oil prices over time. From 2011 to 2014, an inverse correlation is evident, with high oil prices coinciding with lower stock values. The 2015 to 2019 period is relatively stable, with prices moving in tandem. The 2020 COVID-19 impact shows a sharper decline in oil prices compared to airline stocks. Oil production, exports, and imports were affected, while airlines were impacted by travel restrictions. Subsequent recovery shows a divergence, suggesting a shift in the relationship between the two metrics.

Empirical Strategy

Before running any regression models, we set the panel to entities using airline codes and years-months from May 2011 to November 2023. Airline code fixed effects allow us to control for unobserved variables that could differ across airlines but not over time. Year-month or time fixed effects allow us to control for unobserved variables that could differ across time but not over entities (e.g. fuel policy changes, geopolitical events, or regulatory changes). This approach helps us control for unobserved heterogeneity that is constant over time within the entity.

Figure 1

Average Stock Prices & WTI Crude Oil Prices Over Time

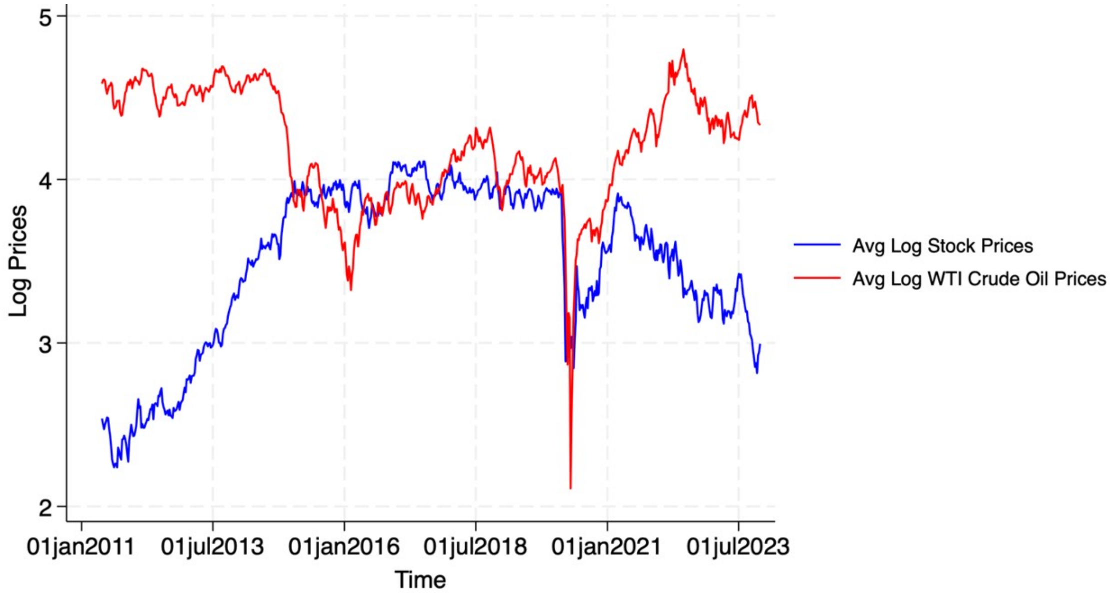

Our primary model uses a univariate regression model to examine the direct relationship between airline stock prices and WTI crude oil prices. This model shows the direct relationship between our independent and dependent variables before adjusting for potential confounding variables. We also control for the S&P 500 weekly closing price.

$$StockPrices_{it} = \beta_0 + \beta_1 WTICrudeOilPrices_t + \beta_2 S\&P500ClosingPrice_t + \varepsilon_{it} \quad (1)$$

where Stock Prices include the stock prices of all airlines (i), WTI Crude Oil Prices is the weekly closing price over time (t), the S&P 500 Closing Price at time (t), and an error term () holding all other potential confounders for airlines (i) at time (t).

Next, we use a log-log regression model to investigate the elasticity between airline stock prices and WTI crude oil prices. This model choice enables the interpretation of coefficients as percentage changes, allowing us to assess the impact of oil price fluctuations on airline stock prices in terms of elasticity. The dependent variable is the log-transformed stock prices of airlines, and the independent variable of interest is the log-transformed closing prices of WTI oil. We add the logged S&P 500 as a control variable. Additionally, we conduct diagnostic tests such as an F-test of joint significance to determine if the log of WTI crude oil prices and logged S&P 500 are jointly significant.

To analyze the percentage change in return (i.e. stock prices), we run a regression with lagged variables and a first difference regression model. Lagged variables help us determine how the independent variables' past values affect the dependent variable's current value. Using past values as predictors reduces the risk of endogeneity since past values are not likely to be caused by current values of the dependent variable. This first difference model directly estimates the effect of changes in independent variables (oil prices, market conditions) on changes in the dependent variable (stock prices). Financial data often exhibit heteroscedasticity, where the variance of the error term changes over time. Differencing can help mitigate this issue, leading to more efficient and unbiased estimators, as it focuses on the changes rather than the levels that might have different variances.

$$\begin{aligned} LagStockPrices_{it} - LagStockPrices_{it-1} = {} & \beta_0 + \beta_1 LagWTICrudeOilPrices_t \\ & - LagWTICrudeOilPrices_{t-1} \\ & + \beta_2 LagS\&P500ClosingPrice_t \\ & - LagS\&P500ClosingPrice_{t-1} + \varepsilon_{it} \end{aligned} \tag{2}$$

where (i) represents all airlines and (t) refers to different time periods. This regression looks at differences between time period (t) and the previous time period (t-1).

Additionally, to look at the relationship between an airline being budget or full service and the effect of oil prices on its stock price, we will run regressions with a binary variable (Full Service) indicating the type of airline interacting with WTI oil prices. By incorporating this dummy variable, the model can more precisely estimate the unique effect of oil prices on each type of airline. The approach aligns with previous discussions that suggest market participants do not respond uniformly to external shocks. Airlines have different elasticities of demand, cost structures, and capacities to pass costs onto consumers.

$$\begin{aligned} LogStockPrices_{it} = {} & \beta_0 + \beta_1 LogWTICrudeOilPrices_t + \beta_2 LogS\&P500ClosingPrice_t + \varepsilon_i \\ & \text{if Full Service} = 0 \end{aligned} \tag{3}$$

$$\begin{aligned} LogStockPrices_{it} = {} & \beta_0 + \beta_1 LogWTICrudeOilPrices_t + \beta_2 LogS\&P500ClosingPrice_t + \varepsilon_i \\ & \text{if Full Service} = 1 \end{aligned} \tag{4}$$

where (i) represents all airlines, (t) refers to different time periods, and () is the error term holding all other potential confounders for airlines (i) at time (t).

Regression Analysis

Table 3

Primary Regressions

Variables	(1) Stock Price	(2) Stock Price
WTI Crude Oil Price	0.0848	0.0451
	(0.0753)	(0.0750)
S&P 500 Close Price		0.0170***
		(0.00336)
Constant	39.23***	-3.296
	(5.298)	(9.718)
Observations	5,886	5,886
R-squared	0.866	0.866

Notes: Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1. Stock Price is the dependent variable, WTI Crude Oil Price is the independent variable, and S&P 500 Close Price is the control.

The WTI crude oil prices coefficient in Table 3 without market controls (Model 1) is 0.0848, indicating a non-significant relationship with airline stock prices. Upon introducing the S&P 500 as a control variable (Model 2), the coefficient slightly adjusts to 0.0451, remaining statistically insignificant. In contrast, the S&P 500 variable is significant ($p<0.01$) with a coefficient of 0.0170, underscoring the pronounced effect of overall market performance on airline stocks, as opposed to the direct impact of oil price fluctuations.

Table 4

Logged Variables Regression

Variables	(1) Log Stock Price	(2) Log Stock Price
Log WTI Crude Oil Price	0.183***	0.105*
	(0.0613)	(0.0595)
Log S&P 500 Close Price		1.461***
		(0.191)
Constant	2.722***	-8.373***
	(0.257)	(1.474)
Observations	5,886	5,886
R-squared	0.915	0.916

Notes: Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1. Log Stock Price is the dependent variable, Log WTI Crude Oil Price is the independent variable, and Log S&P 500 Close Price is the control.

In Table 4, Model 1, the elasticity of airline stock prices with respect to WTI crude

Table 5

F-test on Log WTI Crude Oil Price and Log S&P 500 Close Price

(1) Log WTI Crude Oil Price = 0
(2) Log S&P 500 Close Price = 0
$F(2, 5725) = 33.89$
Prob > F = 0.0000

oil prices is positive and statistically significant at the 1% level, with a coefficient of 0.183. This suggests that a 1% increase in crude oil prices is associated with a 0.183% increase in airline stock prices, highlighting a positive relationship between oil prices and airline stock performance in this model. Including the S&P 500 control variable in Model 2 modifies the log-transformed WTI crude oil prices coefficient to 0.105, which remains statistically significant at the 5% level but suggests a diminished elasticity compared to Model 1. This reduction indicates that part of the positive effect of oil prices on airline stocks may be mediated by overall market conditions, as captured by the S&P 500. The log-transformed S&P 500 closing price in Model 2 has a statistically significant coefficient of 1.461 at the 1% level, indicating that a 1% increase in the S&P 500 is associated with a substantial 1.461% increase in airline stock prices. The constants in both models suggest the baseline level of log-transformed stock prices when oil prices and, in Model 2, the S&P 500, are at their mean levels. The change from 2.722 in Model 1 to -8.373 in Model 2, both significant at the 1% level, emphasizes the impact of including market conditions on the analysis. Observing Table 5, the F-test on the logged WTI crude oil prices and the logged S&P 500 closing prices yields an F-statistic of 33.89 with a p-value of 0.0000. This statistically significant result ($p < 0.01$) indicates that both variables have a jointly significant effect on the dependent variable in the model.

In Table 6, Models 1 and 2 utilizing lagged variables indicate no significant predictive power from past WTI crude oil prices on current airline stock prices. Conversely, the S&P 500's lagged values yield a significant positive coefficient (0.0188, $p < 0.01$), suggesting that past market performance can predict current stock prices. Economically, this implies that a 1% increase in the S&P 500 is associated with a 0.0188% increase in airline stock prices, holding other factors constant. However, this interpretation is constrained by the model's inability to account for all possible factors influencing stock prices, such as macroeconomic shocks or industry-specific trends not captured by the S&P 500 index.

Models 3 and 4 present the first difference results, providing insights into the short-term elasticities. The negative coefficient (-0.0426, $p < 0.05$) for changes in WTI crude oil prices in Model 3 suggests an initial adverse impact on stock prices, which is reversed when controlling for market index fluctuations (0.0414, $p < 0.01$) in Model 4. The positive

Table 6

Lagged and First Difference Regressions

Variables	(1) Lag Stock Price	(2) Lag Stock Price	(3) First Diff Stock Price	(4) First Diff Stock Price
Lag WTI Crude Oil Price	0.0849	0.0233		
	(0.0711)	(0 .0703)		
Lag S&P 500 Close Price		0.0188***		
		(0.00348)		
First Diff WTI Crude Oil Price			-0.0426**	0.0414**
			(0.0198)	(0.0180)
First Diff S&P 500 Close Price				0.0177***
				(0.00105)
Constant	39.21***	-6.577	0.0179	-0.0646*
	(5.005)	(10.06)	(0.0385)	(0.0364)
Observations	5,769	5,769	5,769	5,769
R-squared	0.866	0.866	0.108	0.226

Notes: Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1. Lag stock price is the dependent variable, Lag WTI crude oil price is the independent variable, and Lag S&P 500 Close Price is the control. First Difference Stock price is the dependent variable, First Difference WTI Crude Oil Price is the independent variable, and First Difference S&P 500 Close Price is the control.

and significant coefficient for First Diff S&P 500 Close Price (0.0177, $p < 0.01$) in Model 4 underscores the market's role in driving airline stock performance. The change in R-squared from Models 1 and 2 to 3 and 4 indicates that the first differences capture the short-term volatility in the relationship between the variables, perhaps more accurately reflecting the temporary nature of stock market reactions to oil prices and market changes.

Table 7 shows us results that contrast our initial hypothesis. For budget airlines (Column 1), the coefficient for the logged WTI crude oil price is -0.595, which implies that a 1% increase in WTI crude oil prices is associated with a 0.595% decrease in airline stock prices. Similarly, for full-service airlines (Column 2), the coefficient is slightly larger at -0.670, indicating a 1% increase in oil prices is associated with a 0.670% decrease in their stock prices. This suggests that full-service airlines are more sensitive to oil price changes than budget airlines. The previous aggregate model's (in Table 4) positive coefficient may result from averaging out these individual negative impacts.

The coefficients for the logged S&P 500 closing price are positive in both models (0.401 for budget airlines and 0.473 for full-service airlines), indicating that a 1% increase in the S&P 500 is associated with a 0.401% and 0.473% increase in stock prices for budget

Table 7

Regressions for Budget (1) and Full Service (2) Airlines

Variables	(1) Log Stock Price	(2) Log Stock Price
Log WTI Crude Oil Price	-0.595***	-0.670***
	(0.0608)	(0.0440)
Log S&P 500 Close Price	0.401***	0.473***
	(0.0507)	(0.0298)
Constant	2.915***	2.548***
	(0.505)	(0.314)
Observations	2,616	3,270
R-squared	0.086	0.221

Notes: Robust standard errors in parentheses: *** p<0.01, ** p<0.05, * p<0.1. Log Stock Price is the dependent variable, Log WTI Crude Oil Price is the independent variable, and Log S&P 500 Close Price is the control. Model 1: Full Service = 0, Model 2: Full Service = 1.

and full-service airlines, respectively. This positive relationship is intuitive as an increasing S&P 500 generally reflects improving market conditions, which could benefit airline stocks.

R-squared values show that the model explains 8.6% of the variability in stock prices for budget airlines and 22.1% for full-service airlines, which suggests that other factors not included in the model also play a significant role in determining stock prices, and that the model fits full-service airlines better.

Discussion and Limitations

A key assumption underlying our analysis is market efficiency. We assume that all publicly available information is reflected in airline stock prices. In reality, information asymmetry, delayed market reactions, or other forms of market inefficiency may lead to deviations from this assumption, potentially skewing the results.

Contrary to the initial hypothesis, the results suggest that full-service airlines are more affected by oil price changes than budget airlines. This counterintuitive finding might be explained by the fact that full-service airlines are more likely to be locked into longer-term fuel contracts, possess more effective fuel hedging strategies, or have a greater capacity to pass fuel costs onto consumers. Additionally, their commitment to achieving net-zero carbon emissions by 2050 might prompt a more acute short-term response to oil price volatility. However, the precise mechanisms behind this finding remain unclear due to data limitations.

Another significant limitation is our lack of access to airlines' proprietary hedging strategies. Understanding these strategies would allow for a more comprehensive analysis of how airlines manage the risks associated with oil price volatility. Hedging practices could provide a buffer against oil price changes, thus impacting the observed relationship between oil prices and stock values.

The external validity of the findings may extend to global airline markets similar to the U.S., particularly those in regions like the European Union, Canada, and Australia, which share comparable market structures, regulatory environments, and economic dynamics. This applicability, however, depends on factors such as the prevalence of fuel hedging, demand elasticity, and regulatory intensity. Properly recognizing these factors can refine predictions about airline stock responses to oil price fluctuations in other markets, ensuring that this paper's results are adapted to specific conditions rather than broadly generalized.

Conclusion

This paper examined how crude oil prices differentially influence the stock prices of budget (Southwest, JetBlue, Spirit, and Allegiant) and full-service Airlines (United, American, Delta, Alaska, and Hawaiian). Fixed effects were used to control for airline-specific and time-specific unobserved heterogeneity. Our initial hypothesis of budget airlines being more sensitive to oil price fluctuations was proven incorrect. Our analysis indicated that full-service airlines are affected more by such movement in oil prices.

A key limitation of this approach is the assumption of perfect market efficiency. This suggests that all publicly available information is captured by the airline stock price. However, this is likely not the case. One key piece of information not reflected in the stock prices and this paper is airline-specific hedging strategies and decisions. While this information is confidential, it is a large concern that limits how much we can generalize our findings.

Further research might focus on different or larger geographical contexts to examine whether findings are consistent across regions. Moreover, an investigation into airlines' hedging behaviors and how these strategies impact their financial vulnerability to oil price volatility would help address the limitations of our research.

References

Allayannis, G. and Weston, J. P. (2001). The use of foreign currency derivatives and firm market value. 14(1).

Alıcı, A. (2024). Analysis of macroeconomic factors affecting airline stock prices. (25):93–137. Number: 25.

Apergis, N. and Miller, S. M. (2009). Do structural oil-market shocks affect stock prices? 31(4):569–575.

Carter, D. A., Rogers, D. A., and Simkins, B. J. (2003). Does fuel hedging make economic sense? the case of the US airline industry.

Froot, K. A., Scharfstein, D. S., and Stein, J. C. (1993). Risk management: Coordinating corporate investment and financing policies. 48(5):1629–1658. Publisher: [American Finance Association, Wiley].

Garbuno, D. M. (2020). The difference between full service & low cost carriers. Section: Analysis.

Girardet, D. and Spinler, S. (2013). Does the aviation emission trading system influence the financial evaluation of new airplanes? an assessment of present values and purchase options. 20:30–39.

Géczy, C., Minton, B. A., and Schrand, C. (1997). Why firms use currency derivatives. 52(4):1323–1354. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.1997.tb01112.x.

Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007–08. 2009:215–261. Publisher: Brookings Institution Press.

Hassani, H., Moudassir, M., Christodoulides, P., and Odulaja, A. (2017). A statistical study of the short- and long-term drivers of crude oil prices. 41(2):93–114. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/opec.12095.

Hsu, C. (2017). How fuel price shocks affect airline stock returns: An empirical study of major us carriers.

IATA (n.d.). Our commitment to fly net zero by 2050.

Jones, D. W., Leiby, P. N., and Paik, I. K. (2004). Oil price shocks and the macroeconomy: What has been learned since 1996. 25(2):1–32. Publisher: International Association for Energy Economics.

McKinsey (n.d.). How airlines could adapt to rising fuel prices | McKinsey.

Memon, O. (2019). What is fuel hedging and why do airlines do it? Section: Analysis.

Nandha, M., Brooks, R., and Faff, R. (2013). OIL, OIL VOLATILITY AND AIRLINE STOCKS: A GLOBAL ANALYSIS. 12(2).

Park, J. and Ratti, R. A. (2008). Oil price shocks and stock markets in the u.s. and 13 european countries. 30(5):2587–2608.

Planete-Energies (n.d.). Fuels in aviation and shipping.

Tsai, B. (2008). Financial risk exposures in the airline industry - case of south african airlines.

Upadhyay, R. (2016). Ignoring fundamentals: Speculation has been driving oil prices.

Vasigh, B., Azadian, F., and Moghaddam, K. (2021). Methodologies and techniques for determining the value of an aircraft. 2675(1):332–341. Publisher: SAGE Publications Inc.

Wikipedia (2024). Southwest airlines fleet. Page Version ID: 1216332360.

Euro* Trash?: An Econometric Examination of Evaluating Discrimination in the NBA on the Basis of Nationality

Samina Stack

ECON 381: Introduction to Econometrics (Advisor: Felix Friedt)

For many years, foreign-born National Basketball Association (NBA) players, especially Europeans, were stereotyped as “out of shape, soft, indifferent, or a combo platter of all three” (Roberson, 2023) by fans and scouts alike. However, stereotyping in the media does not necessarily imply discrimination within the league itself. In recent years, the number of foreign players in the league has been increasing with 120 on opening night of the 2022-2023 season (NBA, 2022), which may imply increasing willingness by teams to hire foreign players and thus less discrimination. How can we test whether foreign players actually do face discrimination?

Nick Hornby wrote that “[o]ne of the great things about sport is its cruel clarity [...] in sport, you get found out” (Hornby, 1994). Sports are unique in that individual skill is highly measurable. Nikola Jokic, a Serbian NBA player, won back to back Most Valuable Player (MVP) awards in 2021 and 2022 not because he was the most fun player to watch but because he played with ruthless efficiency and led the league in most prized statistical categories (Nerkar, 2023). The ability to quantify worker output as nicely as in basketball creates ideal conditions to conduct objective quantitative analysis on such a hard to define subject as discrimination.

In this paper, I will examine whether there is a noticeable difference in player salary between foreign and American NBA players of the same skill level. Discrimination can be defined as “differences in pay or wage rates for equally productive groups” (Cain, 1984). Thus, if I find noticeable differences in salaries among the average foreign and domestic NBA player of the same skill level, there is strong evidence to suggest discrimination against foreign players.

Samina Stack is a junior with a major in Economics and minors in Mathematics and Statistics. Correspondence concerning this article should be addressed to sstack@macalester.edu.

Literature Review

Much of the economic literature on discrimination examines gender or racial discrimination. Oaxaca (1973) and Blinder (1973) use residual approaches to estimate the extent of discrimination against female workers and the white-black and male-female wage differentials respectively. Now termed the Blinder-Oaxaca decomposition, the methodologies in Oaxaca and Blinder decompose the differences between wage means into the part due to race or gender and the part not affected by race or gender. Both Oaxaca and Blinder find discrimination effects to explain sizable portions of wage differentials. Previously, Becker (1971) proposed examining racial discrimination within the context of a market. His utility-based approach emphasizes interaction between individuals such that discomfort with black workers can depress their wages and subsidize discrimination. However, Becker also proposes that the costs to employers of discrimination, such as having to pay more for white employees when similarly skilled black workers are available for a lower wage, incentivizes nondiscriminators to hire black workers. With sufficient competition, Becker hypothesizes discrimination should be reduced but ultimately discrimination reflects employer tastes.

Within the context of sports, racial discrimination is well studied. Szymanski (2000) examines discrimination within English soccer using a Becker-inspired market test. Assuming an efficient talent market exists, he asserts club's wage bills should indicate productivity and league performance. His research finds statistically significant evidence of discrimination using these assumptions. Johnson and Gius and Johnson (1998), Groothuis and Hill (2013), Johnson and Minuci (2020) examine racial discrimination within the National Basketball Association (NBA). Using salary data from the 1997 season and performance data from the 1996 season, Johnson's and Gius's log-linear wage regression and Chow test, which tests whether significant differences between the two sets of regression parameters exist, finds no evidence of wage discrimination based on race in the NBA. Groothius and Hill examine exit and pay discrimination, finding a premium in career earnings paid to White players and survival rates of highly-skilled players leading to overrepresentation at high experience levels though overall results are not robust. Johnson and Minuci use data from 2011 to 2017 and find Black athletes to be paid significantly less than White athletes. They use free agency data to represent how players are valued for their marginal revenue products.

Recently, discrimination against foreign-born basketball players has received greater research attention. Yang and Lin (2012) using data from the 2000 to 2008 NBA seasons and a two-stage double fixed-effect model find that foreign players receive lower salaries than similarly skilled US-born players but foreign players from large economies receive preferential labor market treatment as the viewership from their home country can generate significant revenue for teams. Hoffer and Freidel (2014) evaluate salary discrimination against foreign NBA players using Becker's theory of wage discriminators eventually being priced out of the market and 2011 season data. They find foreign-born players to receive a wage premium, earning more than similar domestic players, attributable to the extra revenue they can bring their franchises, though only 89 foreign-born players are in their data. Looking at both the NBA and Spanish Liga ACB, Berri et al. (2015) find opposing evidence to suggest US-born players receive preferential treatment in both leagues by considering playing time allocation instead of salary due to limited data.

My analysis will build off of the existing research by examining a broader range of data spanning a decade to capture more variation in the data. Given that foreign players comprise a much smaller portion of players than domestic players do, I should get a big enough sample to draw reasonable conclusions. I will follow Becker's methodology of using a market test to identify any valuation inefficiencies assuming a competitive market for talent exists. If this efficient market exists, then any inefficiencies in the form of discrimination would be identifiable.

Theory

My research examines whether NBA teams discriminate against foreign-born players. We can think of each NBA team as a firm in the NBA market. The objective of each team is to win. As it is the athletes themselves who play games a team wants to accumulate the best roster it can to maximize its chances of winning. Thus, in an ideal world, teams can select the best available players for their team. In actuality there are some constraints in place such as the total amount of money a team can spend on its roster and the number of roster spots on a team. Thus to maximize winning, a team should look to construct the most talented roster it can for a given amount of money and number of roster spots. If two equally talented players are priced differently and teams are cost-minimizers, then they should select the lower cost option. However, we will investigate whether the value to winning alone does not determine roster construction, but a player's country of origin influences selection decisions. Specifically, we analyze whether teams prefer domestic or US-born players to foreign-born ones.

To empirically provide evidence of any discriminatory practices, we first need to rule out other possible explanations for wage differences between the two groups.

One possibility we need to consider, is that foreign-born players are simply willing to work for a lower wage than are domestic players. If we assume each team is allowed a fixed number of roster spots, then demand for players is perfectly inelastic at some fixed value $\overline{D}$. Thus, any variation in wage should be attributable to differences in supply of labor. If for players of the same skill level, foreign-born players earn less than their domestic counterparts, we can model two different supply curves, one for each type of player. We can think of it as foreign-born players being willing to work for a lower wage because of a lack of opportunities outside the US of the same caliber. If the NBA is regarded as the best league in the world, then players might be willing to accept lower pay than their ability would indicate to get the opportunity to compete in such a prestigious league. Thus, as we can see from Figure 1, under these conditions we expect for the same level of popularity, skill, and team income, foreign-born players will earn less than domestic players. Under these conditions, foreign-born players make themselves less expensive for teams rather than teams discriminating against foreign-born players to make them undervalued in the market.

Another way we can conceptualize a possible non-discriminating market inefficiency like this is in terms of ability to predict the skill of both groups of players. Suppose that teams have a test or calculation they use to determine how much to pay a player. Consider the equation

$$w = \alpha \cdot skill_m + (1-\alpha)\overline{skill}_p, \quad \text{where } \alpha \in (0,1) \tag{1}$$

where α represents the reliability of the test and $\overline{skill}_p$ is the average skill level for domestic or foreign-born players (Lundberg & Startz, 1983). Here wage represents the predicted skill level of a given player. If the skill test is more reliable for domestic players (US) than foreign-born players (FB), then predicted skill levels should be more accurate for those domestic players. According to the equation, this means that predicted skill of foreign-born players will be closer to the average skill observed among foreign-born players in the NBA. Looking at Figure **??**, we see that, under these conditions, for lower skill levels, foreign-born players are valued higher than domestic players in terms of predicted skill, but for higher skill levels domestic players are valued more highly. This implies that there may be competing relationships between foreign-born status and valuation for different levels of skill. If predicted skill can be approximated by salary, this means that lower-skilled foreign-born players and higher-skilled domestic players receive salary premiums under these assumptions of information asymmetry that make the test less reliable.

However, if the group average were subject to employer bias, a difference in wages would not just be information asymmetry. Consider the predicted group average written as

$$\overline{skill}_p = \overline{skill}_r(1-b) \quad \text{where } b \in [0,1], \tag{2}$$

where $\overline{skill}_r$ is the real group average and b represents employer bias. If employers are unbiased, then the value for average skill is the same as the actual average skill observed. If they are fully biased, then the predicted average becomes zero, and with a noisy enough test, teams choose not to hire foreign-born players.

Thus, to make any claim of discrimination, we need to ideally rule out the possibility of information asymmetries or lower NBA reservation wages among foreign-born basketball players. If teams discriminate against foreign-born players, then we should be able to empirically find this by modeling player salary as a function of whether they were born in a foreign country or the United States, holding marginal revenue product (as measured by skill), reservation wage, and availability of information constant.

Empirical Theory

To estimate the relationship between foreign-born status and salary in the NBA a few considerations need to be made. Firstly, as discrimination is not a measurable variable we will approximate it using the coefficient on foreign-born status. If a foreign-born player with all the same characteristics as a domestic player except for country of origin is on average paid statistically significantly less than that domestic player, and the conditions for willingness of foreign players to accept lower wages or information asymmetry do not hold, we can infer discrimination on the basis of nationality.

There are many factors that can influence salary. To simplify our analysis, we make the assumption that teams are win-maximizers (as opposed to profit-maximizers) such that the primary factor is player skill, which should approximate contribution to winning. As we can see from Figure 3, other factors correlated with salary may be experience as though younger players can play for more years than older players, they require time to adjust to the level of competition, and draft round as we expect more talented prospects are selected in earlier rounds of the draft. However, if nationality impacts salary then we expect it to also correlate with draft round as a pro-domestic bias should result in more domestic players being drafted earlier. Our general model specification can be written:

$$\ln \widehat{Salary}_{pst} = \hat{\beta}_0 + \hat{\beta}_1 FBorn_p + \hat{\beta}_2 X_{ps} \times FBorn_p + \hat{\beta}_3 C_p + \hat{\beta}_4 Allstar_{ps} + \alpha_s + \alpha_t \quad (3)$$

where X represents a vector of lagged performance variables including points per game, assists per game, true shooting percentage, and rebounding percentage, all interacted with foreign-born status. C is a vector of player characteristic variables including height, experience, and experience squared. *Allstar* is a dummy variable representing whether or not a player was voted an all-star in a given season. Finally, α_s, α_t are matrices of season and team fixed effects. We implement these fixed effects as perception of foreign-born players may change over time and certain team playing styles could lead us to misattribute a discrepancy as discrimination.

Data Description

To conduct the investigation, we use panel data of seasonal performance and salary on individual NBA players from the 2000 to 2019 seasons (NBA, 2022). The data contain close to 5,000 cases containing information on player position, team, salary, age, country of origin, and performance statistics. These data were compiled using ESPN NBA salary data, NBA players biographic data compiled by Justinas Cirtautas, and data from Basketball Reference (ESPN, 2000; Cirtautas, 2023).

We use salary as the outcome variable. Assuming rational actors and a competitive market, if the average domestic player is paid more than the average foreign-born player of the same skill level, then we can infer possible wage discrimination by the player's country of origin. Looking at Figure 4, we can see that in this data, median salary in 2019 US dollars for both US-born and foreign-born players look very similar, with slightly more variation in the foreign-born interquartile range.

It is difficult to distinguish the difference within the interquartile ranges of foreign-born and US-born players though due to the outliers present in both that graphically compress the boxes. Comparing the outliers between the two, which we may interpret as highly-skilled players or superstars, far more US-born players achieve this status than foreign-born and there are many US-born superstars within the data that received a higher seasonal wage than the highest paid foreign-born player. This is the only obvious difference in compensation indicated by the visualization, suggesting that there may be more discrimination between superstar status players based on country of origin than between average players.

We can look at this relationship over time to see if these outliers are perhaps influenced by one specific year and to see if salaries for both foreign and domestic players generally align. Figure 5 displays this relationship. We can see that the mean salary for foreign and domestic players has moved differently over time. More recently, after about 2015, the two curves converge and closely align. Previously though, they would at times move in opposite directions. The graph seems to indicate that foreign-born and domestic players have been valued differently throughout the seasons in the data, with foreign-born sometimes being valued lower or higher than domestic players.

Table 1 shows us some summary statistics regarding variables of interest across foreign-born and domestic players. From this we can see that the average salary in the data does not vary substantially between the two groups. Player height may influence salary as height is a useful trait in basketball and cannot be acquired in a controlled manner so taller players may be valued higher than shorter players of the same skill level. Here, we again see very little difference, which might partially explain the similarity in salaries. Domestic players tend to score more points per game, which we theorize is positively associated with salary due to the direct impact on winning that scoring points has. This could partially explain a wage premium paid to domestic players if we find this in our models. Lastly, experience in the NBA also remains fairly consistent across the two groups. Based off of the findings in the summary table, we cannot locate any major differences between foreign-born and domestic players in key variables. The largest difference is that in sample sizes as there are far more domestic than foreign-born players in the data. This limited sample of foreign players could potentially skew results as it may not be as representative of a sample.

Results

Table 2 displays the coefficient estimates for three models. The first is a naive model that regresses salary only on foreign-born status to set a baseline estimate of potential discrimination. The second interacts a vector of lagged skill measures with the foreign-born dummy variable, and the third adds season and team fixed effects to the second model specifications to control for time-specific variables such as the global popularity of the NBA and different playing styles employed by different teams whether that be due to coaching preferences or roster makeup.

The coefficient on our primary regressor of interest, foreign-born status varies in magnitude, but not sign, across all three models. The naive model estimates that on average we expect domestic players to earn about 24.9 percent more than their domestic counterparts and this estimate is statistically significant at the one percent level. This would indicate that foreign-born players, on average, receive a wage premium relative to their domestic counterparts, implying possible discrimination against domestic players, going against the intuition laid out in our theory. The next two models estimate much larger coefficients on the foreign-born dummy but neither are particularly significant. Both estimate foreign-born players to have salaries more than 100 percent larger than those of domestic players, holding all else constant, with the third estimate being significant at the 10 percent level, which for the purposes of this paper is insignificant. We can note that as more specifications are added to the model, the coefficient on foreign-born status only increases in magnitude, indicating that these estimates may not be very accurate. If these unexpectedly high results are accurate, then NBA teams pay a significant wage premium to foreign-born players, creating a large market inefficiency that should have been corrected over the 20 years of data in our sample.

The direction of the coefficients on the significant controls agree with our intuition for the most part and are consistent across the last two models. The models predict that scoring more points and blocking more shots per game, and having more experience in the NBA is associated with higher salaries, holding all else constant, among foreign-born players. This agrees with our intuition as the team with the most points in a game is the winner and scoring increases a player's team's points while blocking shots prevents the other team from scoring and players likely adjust to the level of the NBA with more time in the league. Less intuitively, higher steals per game are associated with lower salaries and more turnovers per game are associated with higher salaries for foreign-born players, holding all else constant.

To verify the initial model results we run three more models, the results of which are shown in Table 3. The first of these keeps everything in the primary specification the same except for the outcome variable, which become minutes played per game. Here we see a negative association between being foreign-born and the outcome. In fact, according to the model, a foreign-born player is, on average, associated with playing about 17 fewer minutes per game than a domestic player of the same skill level. This estimate conforms much more closely with our theory but is only significant at the 10 percent level, which we consider insignificant.

The second model adds a variable, country GDP (measured in 2019 US dollars), to the primary specification. Previous research (Yang & Lin, 2012) has found that foreign-born players from larger economies may receive higher salaries than foreign-born players from smaller economies due to marketability and viewership in those larger economies. As such, we add GDP for each player's country of origin as a measure of the economic size of the country. The model estimates the relationship between country GDP and logged salary to be statistically significant at the one percent level, though the estimate of this relationship is so small, it is almost zero. That said, in this specification of the model the coefficient on foreign-born status is substantially smaller in magnitude than in the primary model, but statistically indistinguishable from zero.

The third robustness check decreases the sample to players drafted between the years 1999 and 2016 to limit our sample to players whose entire career or first three years we have. Despite decreasing our sample considerably , the predictions of this model are consistent with our initial findings. This specification again indicates foreign-born players to receive more than 100 percent more than domestic players of the same skill level.

The inconsistency in our robustness check model predictions indicate further research is required to properly test for discrimination on the basis of country of origin in the NBA.

Conclusion

In this paper, we analyzed possible discrimination in the NBA based on country of origin. While we expected to find evidence of discrimination against foreign-born players, our primary models indicated just the opposite. However, while these primary findings were highly statistically significant, the extremity of the estimates, incompatibility with theory, and inconsistent robustness tests indicate further research is needed to make any strong claims regarding possible discrimination.

There are a number of factors that contributed to these unintuitive and inconsistent findings. Firstly, one of the most important factors to consider when attempting to empirically test for discrimination is how to accurately measure worker marginal revenue products (MRPs). If player MRPs are misapproximated, we fail to properly control for individual skill. Additionally, we failed to fully account for the global popularity of the NBA. While our season fixed effects controlled for variations in popularity over time, we failed to account for popularity within a given country or by player. More popular players may be offered higher salaries because they attract more fans to purchase tickets and merchandise. We also only lagged performance variables by one season. As players sign multiyear contracts, during the final season in a contract, a player may play much better than in previous years to secure a lucrative contract. This also introduces some reverse causality to the model as in this scenario, current and potential salary influence performance. We also fail to capture movement within the job market as players are traded to or move in free agency to different teams, possibly midway through the season. The adjustment time and different play styles could impact performance, and thus salary, in ways not captured by the team fixed effects as each player is only associated with one team per season.

Future research on this topic should examine different groups of NBA players based on performance to assess whether any discrimination is more localized to those particular subgroups. Additionally, only looking at more recent data to factor in the impact of social media on player popularity and marketability might capture the relationships in the data more effectively.

References

Becker, G. S. (1971). *The economics of discrimination.* Chicago, University of Chicago Press.

Berri, D. J., Deutscher, C., and Galletti, A. (2015). Born in the USA: National Origin Effects on Time Allocation in US and Spanish Professional Basketball. *National Institute Economic Review*, 232(1):R41–R50. Publisher: SAGE Publications Ltd.

Blinder, A. S. (1973). Wage Discrimination: Reduced Form and Structural Estimates. *The Journal of Human Resources*, 8(4):436–455. Publisher: [University of Wisconsin Press, Board of Regents of the University of Wisconsin System].

Cain, G. G. (1984). The economics of discrimination: Part.

Cirtautas, J. (2023). NBA Players.

ESPN (2000). NBA Player Salaries - National Basketball Association - ESPN.

Gius, M. and Johnson, D. (1998). An empirical investigation of wage discrimination in professional basketball.

Groothuis, P. A. and Hill, J. R. (2013). Pay Discrimination, Exit Discrimination or Both? Another Look at an Old Issue Using NBA Data. *Journal of Sports Economics*, 14(2):171–185. Publisher: SAGE Publications.

Hoffer, A. J. and Freidel, R. (2014). Does salary discrimination persist for foreign athletes in the NBA?

Hornby, N. (1994). *Fever pitch.* London : Gollancz.

Johnson, C. and Minuci, E. (2020). Wage discrimination in the NBA: Evidence using free agent signings. *Southern Economic Journal*, 87(2):517–539. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/soej.12461.

NBA (2022). NBA Rosters Feature 120 International Players from 40 Countries.

Nerkar, S. (2023). Nikola Jokić's MVP Case Is Hidden In Plain Sight (Again).

Oaxaca, R. (1973). Male-Female Wage Differentials in Urban Labor Markets. *International Economic Review*, 14(3):693–709. Publisher: [Economics Department of the University of Pennsylvania, Wiley, Institute of Social and Economic Research, Osaka University].

Roberson, M. (2023). The NBA Officially Belongs to International Players. Section: tags.

Szymanski, S. (2000). A Market Test for Discrimination in the English Professional Soccer Leagues. *Journal of Political Economy*, 108(3):590–603. Publisher: The University of Chicago Press.

Yang, C.-H. and Lin, H.-Y. (2012). Is There Salary Discrimination by Nationality in the NBA?: Foreign Talent or Foreign Market. *Journal of Sports Economics*, 13(1):53–75. Publisher: SAGE Publications.

Tables and Figures

Figure 1

Demand and supply diagram explaining theoretical difference in wages.

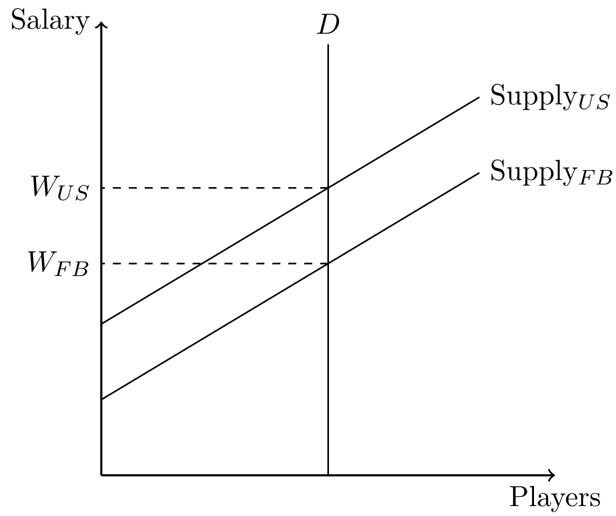

Figure 2

More reliable test for domestic worker skill and less reliable for foreign-born leads to predicted skill of domestic having more variation while predicted foreign-born is closer to the average.

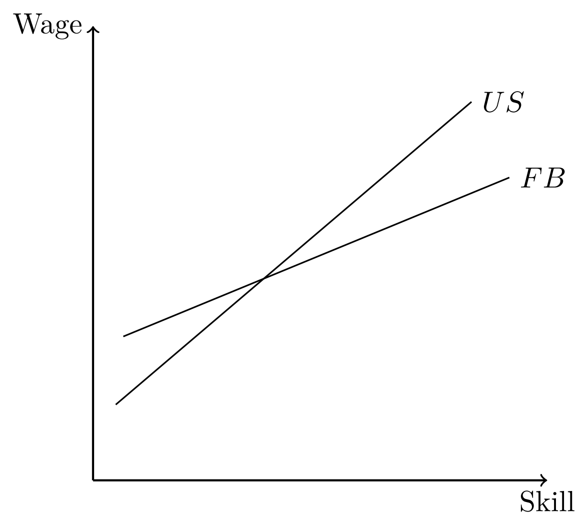

Figure 3

Causal diagram showing theorized relationships relevant to the nationality-salary relationship.

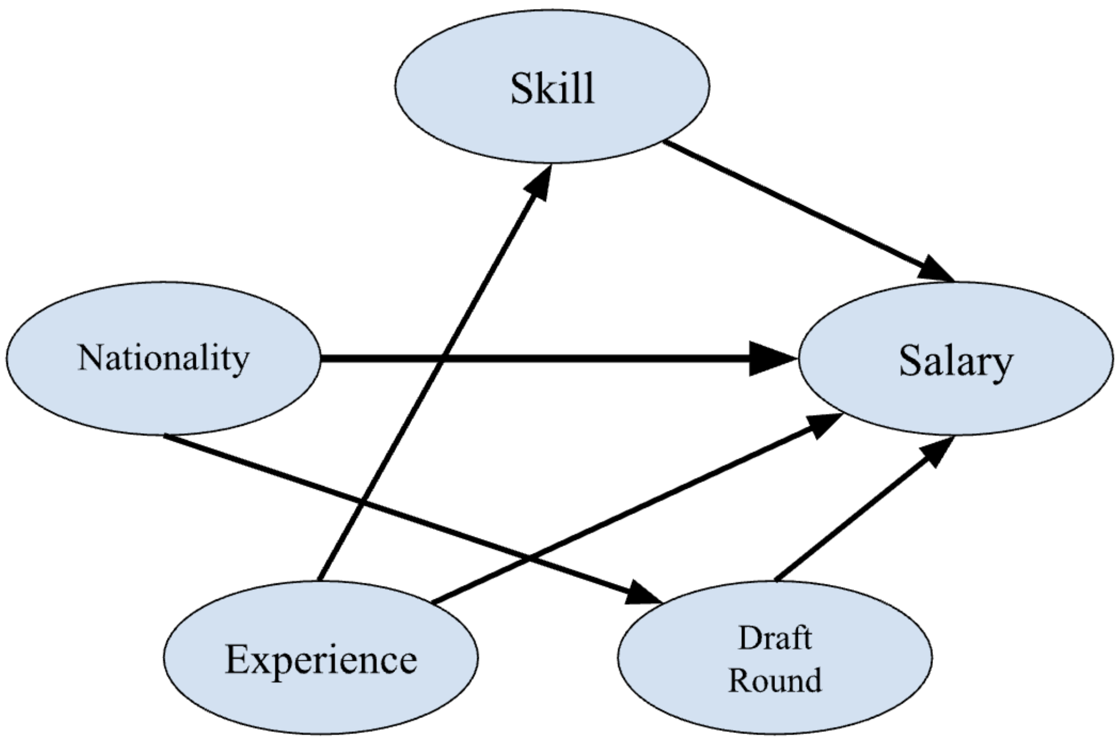

Figure 4

Boxplot comparing salary between domestic and foreign-born players.

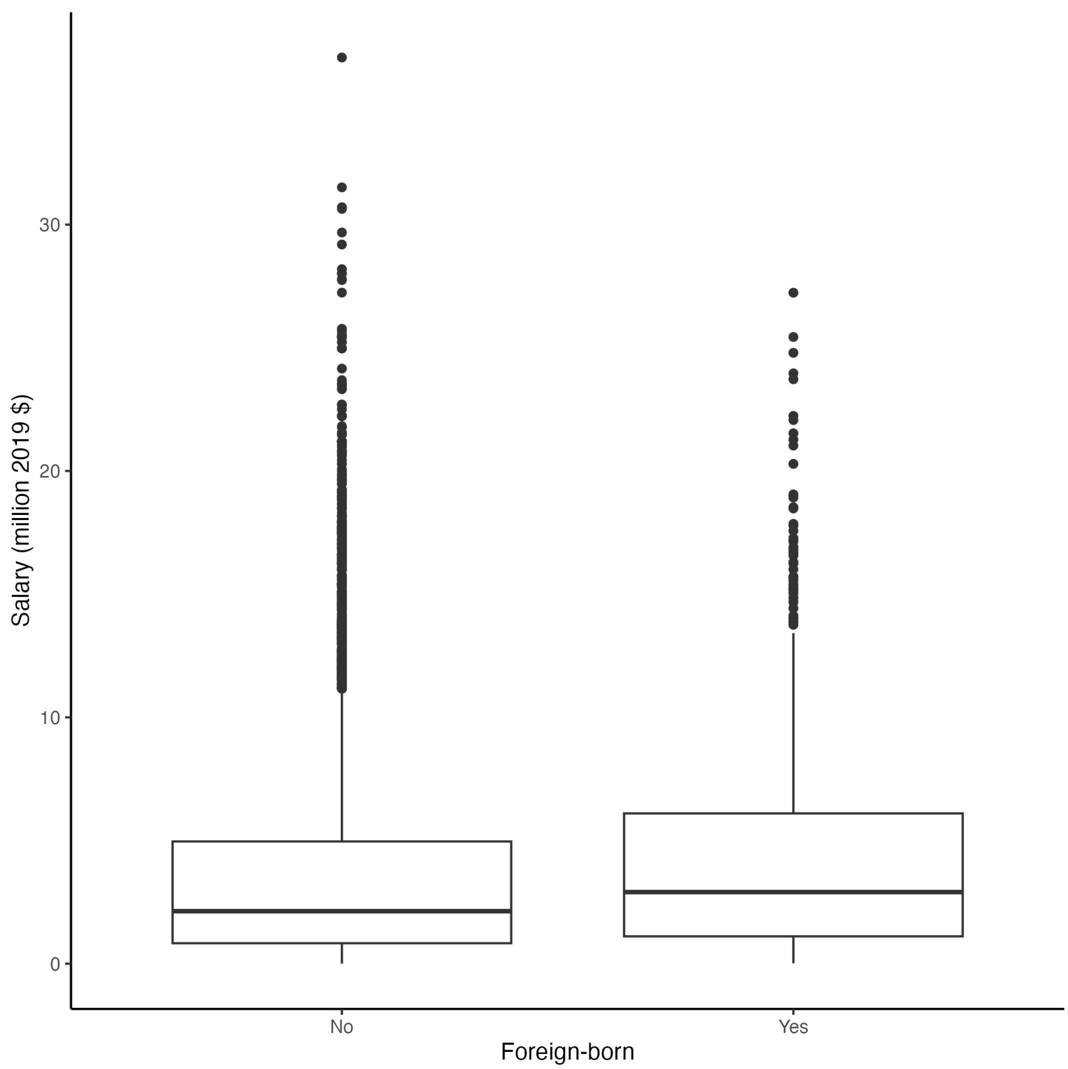

Figure 5

Line plot showing movement of mean salary for foreign-born and domestic players over time.

Table 1

Summary statistics for key variables

Summary statistics:
Nationality Dummy (1 if FB):
Domestic

	N	Mean	SD	Min	Max
Salary ($ 2019)	3193	14.745	1.144	10.039	17.421
Player height	3193	200.471	8.510	175.26	218.44
Points per game	3193	9.4	5.966	0	32
Experience	3193	5.784	3.283	1	18
Foreign-born					
Salary ($ 2019)	722	14.967	1.042	11.04	17.12
Player height	722	206.222	8.405	182.88	228.6
Points per game	722	8.72	5.191	.3	25.2
Experience	722	6.411	3.382	2	19

Table 2
Primary specification output table

VARIABLES	(1) Naive Model	(2) Interaction	(3) Season & Team FE
Nationality Dummy (1 if FB)	0.267***	1.267***	1.145***
	(0.0420)	(0.375)	(0.368)
Lagged Points per game		0.0846***	0.0800***
		(0.00546)	(0.00539)
Logged lagged assists per game		-0.0118	-0.0143
		(0.0230)	(0.0227)
Lagged usage percent		-1.452***	-1.155**
		(0.460)	(0.452)
Lagged steals per game		-0.0671	-0.0316
		(0.0447)	(0.0441)
Lagged blocks per game		0.263***	0.278***
		(0.0292)	(0.0290)
Lagged Effective Field Goal %		0.0444	0.0490
		(0.192)	(0.191)
Lagged turnovers per game		0.126***	0.131***
		(0.0335)	(0.0331)
Experience		0.306***	0.302***
		(0.0128)	(0.0128)
Experience squared		-0.0146***	-0.0140***
		(0.000745)	(0.000745)
Draft round (=4 if undrafted)		-0.248***	-0.238***
		(0.0293)	(0.0290)
Allstar status		-0.00849	0.0116
		(0.0536)	(0.0530)
Constant	14.69***	13.05***	12.56***
	(0.0166)	(0.141)	(0.555)
Observations	6,189	4,073	4,073
R-squared	0.006	0.512	0.545

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

Table 3

Robustness checks output table

VARIABLES	(1) Games Played	(2) Country GDP	(3) Careers
Nationality Dummy (1 if FB)	-17.464*	-0.081	1.087***
	(10.010)	(0.514)	(0.392)
Lagged Points per game	1.705***	0.087***	0.0865***
	(0.167)	(0.007)	(0.00655)
Logged lagged assists per game	4.699***	0.005	0.00160
	(0.712)	(0.028)	(0.0279)
Lagged usage percent	-83.580***	-1.170**	-1.132**
	(13.270)	(0.519)	(0.520)
Lagged steals per game	1.114	-0.067	-0.0632
	(1.374)	(0.054)	(0.0538)
Lagged blocks per game	2.903***	0.266***	0.265***
	(1.016)	(0.040)	(0.0398)
Lagged Effective Field Goal %	19.439***	0.329	0.318
	(5.793)	(0.227)	(0.227)
Lagged turnovers per game	-3.074***	0.087**	0.0912**
	(1.056)	(0.041)	(0.0414)
Experience	-0.871	0.270***	0.273***
	(0.547)	(0.021)	(0.0214)
Experience squared	-0.005	-0.012***	-0.0122***
	(0.040)	(0.002)	(0.00158)
Draft round (=4 if undrafted)	-3.210***	-0.328***	-0.321***
	(0.894)	(0.035)	(0.0350)
Allstar status	-2.825*	-0.012	-0.0156
	(1.686)	(0.066)	(0.0661)
Country GDP ($ 2019)		-0.000***	
		(0.000)	
Constant	66.483***	13.388***	12.80***
	(5.667)	(0.279)	(0.222)
Observations	2,632	2,626	2,632
R-squared	0.306	0.588	0.586

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