THE EAGLETARIAN

BOSTON COLLEGE ECONOMICS ASSOCIATION

The **Boston College Economics Association** provides students a forum to discuss and explore economics related issues with classmates, professors, and professionals. We accomplish this through small events designed to allow students the opportunities to meet BC faculty as well as larger lectures with professionals in the field. In addition, we publish our annual economics research journal, *The Eagletarian.*

BC ECONOMICS ASSOCIATION

The 2025 Eagletarian

Boston College's Undergraduate Economic Research Journal

Boston College Economics Association 2023-2024 Executive Board:

Dear Reader,

The Boston College Economics Association (BCEA) is proud to share the 2025 edition of The Eagletarian. We received a number of excellent papers from across the Boston College student body and believe this publication provides a glimpse of the talent and thoughtfulness of our peers. After several rounds of careful reading and discussion, the editorial board selected the six essays featured here. These pieces reflect not only the dedication of their authors but also the editorial board's commitment to presenting what we felt were the most insightful and compelling works.

These six fantastic essays will appeal to a wide audience, covering topics ranging from an analysis of the cost efficiency of the death penalty to exploring intergenerational income mobility. The content covers policy suggestions, deep regression analysis, and insightful literature reviews, allowing for a rich reading experience for a broad community of readers.

We would also like to express our gratitude to several people whose support was invaluable in putting together this journal. This publication would not have been possible without Professor Geoffrey Sanzenbacher, Professor Kenneth Felter, and our amazing faculty advisor, Professor Matthew Rutledge. We extend our sincere thanks to the faculty of the Boston College Economics Department for their exceptional teaching and for encouraging their students, our peers, to share their work with the journal and the broader Boston College community. Lastly, we would like to thank BCEA Co-Presidents Sammy Wood and Lisa Su for their leadership and support this year.

Sincerely,
Annie Li ‘25, Anthony Yang ‘25, and Nico DeZerega ‘26

TABLE OF CONTENTS

The Parent Trap: Investigating the Relationship Between Children and Retirement Age

Brendan McLaughlin

ECON3242: Economics of Aging

Professor Rutledge

May 2, 2024

Abstract

The timing of one's retirement can be a complex decision, as individuals must consider their immediate needs alongside long-term goals. One significant factor to consider is the impact of raising children. For example, individuals may retire earlier to spend more time with their children, or later to continue financially supporting their family. Using fully retired individuals from the University of Michigan's *Health and Retirement Study* (HRS), I demonstrate a positive correlation between having additional children and one's retirement age. This relationship highlights the broader challenges of retirement security, as extended working years may be necessary to support larger families. These findings contribute to understanding how family size influences retirement decisions, as well as inform policymakers of potential action to assist families' retirement security.

Introduction

Children are an expensive investment. To raise a child effectively takes time, effort, money, and planning. Financially, they require significant costs to bring into this world, provide basic needs, and raise adequately. These costs include hospital bills for childbirth, ongoing expenses for food, clothing, diapers, medical care, and more. Children also do not require one-time down payment; they are a recurring cost for families until they are able to live on their own and provide for themselves. Parents may face reduced income due to one parent cutting back work hours or leaving the workforce entirely. Even if both parents stay in the workforce, they must carve large chunks of time out of their work weeks in order to successfully raise their child. Especially when it comes to newborns, care is needed around the clock. Mothers may have to go on maternity leave, paid or not. This leave may lead to a temporary or even permanent loss of future wages. Fathers will take care of their wife throughout the whole process. There is also lots of planning that goes into raising a child. Parents must work out new living arrangements for a child, whether that means a room for their child or upgrading to a larger, more expensive home. Parents must research the childcare options in their area such as daycare centers, nannies, or family daycare providers, and their associated costs. They also will need to review their health insurance coverage to understand what is covered during pregnancy, childbirth, and pediatric care. Overall, bringing a child into the home brings a complete change in lifestyle in all aspects.

Despite the costs of raising a child, the return one can receive is extraordinary. One has the ability to create a child with the person they love. The joys of nurturing and caring for a child can be immensely rewarding. The emotional bonds formed with one's children provide a deep sense of fulfillment and happiness. From a long-term health standpoint, one's children are the number one provider of long-term care other than one's spouse. While the decision to have children is deeply personal and comes with its challenges, many parents find that the rewards and joys of parenthood far outweigh the costs.

Part of the difficulty of planning around children is ensuring financial stability. It is difficult to both raise children adequately and effectively save to retire at one's desired age. With the dropping fertility rate, fewer adults in the US are having kids today, and this could be in part due to the financial burden children bring. While saving effectively and raising kids is difficult, it is still possible to do both. One way that people may make up for that investment is by staying longer in the workforce. This way, they can make up for lost savings potential from investing in their children. I look at retired individuals from the HRS, specifically from the years 1998 to 2020, and view the age in which they retire. Using multiple regression analyses, I demonstrate that each additional child that a respondent has is associated with an increase in their retirement age of around six months. This was a significant result, and supports my hypothesis. The regression also shows other significant factors that can increase one's retirement age, such as average earnings or one's self reported health.

In the next section, I review related economic research on the history of fertility rate in the US and why children add to the challenge of retirement saving. I outline how I use a sample of adults who have had at least one child from the HRS dataset, and pinpoint the age in which they first declared being fully retired. In the following section, I describe the results of my statistical analysis and regression of the number of children at one's retirement age, as well as the caveats of my analysis. Finally, I conclude that adults who have children tend to - on average - retire later in their lives, and these effects are greater for newer generations of retired parents. I propose ways to incentivize adults to both have children and contribute to their retirement at the same time.

Literature Review

In examining the relationship between children and economics, fertility rates come at the forefront of the discussion. Since the 1960s, there has been a steady decline in the fertility rate in the United States. Boldrin et al. (2015) argue that this decline could coincide with the increase in popularity of government retirement programs. The analysis done by Boldrin exhibits a strong negative correlation between the TFR (Total Fertility Rate) of the US and the size of the Social Security system, as well as the public pension system. This is true for other countries as well. This study's model of fertility rate is based on the "old age security" hypothesis: The idea that having more children is more beneficial to one's retirement utility than saving money, because one's children will be able to provide and care for them. With the introduction of Social Security and public pension systems, people feel less of an incentive to have children as they are more able to seek assistance from the government in old age rather than needing to rely on their children. Using the old age security model, Boldrin accounts for 40% to 60% of the total drop off in TFR to date. Understanding Social Security's impact on both having children and retirement saving is vital to this paper's analysis.

In a more recent paper, the falling US fertility rate is put in the context of women's increasing presence outside of the home (Kearney & Levine, 2022). Kearney and Levine hypothesize that women who are in child-bearing years now have different priorities than women of the previous generation. They posit this could be because women who were born in the 1980s or 1990s grew up with an emphasis from their parents on exploring a future outside of giving birth to children and staying within the home. Contemporary cultural norms emphasize a decline in the stay-at-home mom archetype and an increasing percentage of females entering the labor force. More women in the workforce and fewer children being born have eased financial pressures for couples, allowing them to save more for retirement. The changing priorities of US adults discussed in this paper point to the declining desire to have children and increasing desire to invest in oneself. These different norms have influenced how and when individuals choose whether to have children, in addition to how many they have.

One reason Social Security's introduction has led to a greater emphasis on working more is that Social Security benefits are in part calculated by Average Indexed Monthly Earnings (AIME) which consist of an individual's thirty-five highest earning years of their career. Therefore, working an additional year later in life enables program participants to substitute a year of low earnings with a year of the higher earnings typically expected later in life. In a working paper by Rutledge and Lindner (2016), they propose that women stand to gain more from an additional year of working due to potential lost earnings from child-rearing. Specifically, they use the HRS dataset to show that working an extra year increases monthly Social Security retirement benefits by an average of 8.6 percent for women and 7.6 percent for men. Coinciding with Boldrin's analysis of Social Security's effect on the TFR in the US, the program incentivizes workers to maximize thirty-five well-earning years in order to grow their retirement benefit to the optimum level. This is especially true for women who are more likely to have low or zero-earning years from child-rearing. Rutledge and Lindner highlight the implicit costs of having children in the form of lost Social Security benefits from low-earning years. These costs can be offset by delaying retirement and working longer in order to make up for low-earning years. They can also be offset by having fewer children and reducing the number of low-earning years altogether. Given the evolving cultural attitudes toward work and family, both of these strategies could serve as viable solutions that parents take to optimize financial flexibility in retirement.

Children are associated with a number of implicit and explicit costs. These costs are difficult to quantify, due to their inherent variability and wide-ranging impact upon resource allocation within the household. Child expenditures do not only include financial spending, but also the opportunity costs of lost wages, lost leisure time, shifts in routine, and countless others. Despite the challenges of quantifying these costs, Bradbury's paper "Time and the Cost of Children" (2008) clarifies that "Parents reduce their leisure and personal hours considerably when they are raising their children … This change in time-use arises from a combination of the time and the expenditure costs of children." (Bradbury, 2008). Bradbury's use of an "adult goods" economic model that takes into account both time and financial expenses associated with raising children provides a foundation for his findings. He concludes that a child's most cost-intensive years come within the first five years of their life. Bradbury posits parents could benefit by proactively spreading out child-rearing costs across working years, providing policymakers with data to focus on cost amelioration in the first five years. Being aware of this information lends insight to prospective parents who are aiming for both financial security and the joys of parenthood.

In addition to their costliness, children are a long-term investment as well. Even having just one child can majorly change expenditure patterns for working adults. Biggs (2019) discusses that expenditures between parental households and non-parental households begin similarly, but diverge over time. On average, the expenditures of parents peak in their 50s and then decline. Non-parents' expenditures tend to rise continually until retirement age. This decrease in spending for parents in their 50s is hypothesized as concordant with the time point when the household's children become economically independent from the parents. Because the process of raising children significantly affects wealth and consumption, it therefore affects the household savings/consumption rate and retirement security. Adults without children have greater flexibility in managing their spending and savings habits. Conversely, parents face additional expenses that must be addressed before they can allocate funds to discretionary spending and retirement savings.

Expenditure and savings patterns are explored more by the Center for Retirement Research (CRR) of Boston College. Biggs, Chen, and Munnell (2021) explore how a household's savings rate changes after children leave the home and become financially independent. Once children leave the household, parents allocate their extra disposable income to various endeavors, such as boosting savings, repaying mortgage or settling other debt. At this point in their lives, most parents will work less hours, thereby decreasing their income, instead opting for more leisure as they approach retirement. There are many different financial routes parents may choose from, but it is apparent that when children leave the house there are strategies that lead to saving more if they were unable to save while raising children.

Given the notion that a household deciding to have children results in major costs during the time children are living at home as dependents, there are strategies that parents can use to ensure they retire with adequate funds. The CRR also explored how some parents behave suboptimally in this scenario. In a Policy Brief discussing retirement security, Munnell, Hou and Sanzenbacher (2017) found that some parents increase their consumption when they have children, but fail to bring that consumption back down after the child becomes independent and leaves the home. The optimal strategy would be to keep consumption level flat, or to bring consumption back down after increasing it while raising children. Overall, the brief finds that each child is associated with a 4% decrease in overall wealth. While this may seem significant, it is important to note that for most parents it is certainly possible to both raise children and save adequately for retirement.

Lastly, using the HRS dataset, Scholz & Seshadri (2007) examined the effects of children on wealth in the household using complex models concerning consumption levels during worker's tenure in the labor force. Their study compared 1992 net worth to lifetime earnings, measuring the impact that having young kids in 1992 would have on total lifetime earnings. Their results concluded the ratio of 1992 wealth to lifetime wealth was highest with 1-2 children, and decreased with each additional child. This suggests that families with fewer children tend to accumulate more wealth relative to their lifetime earnings compared to those with more children. It is clear that having and raising children is not solely a binary variable, but the number of children a family raises is just as important. More children in a household means more expenditures on food, shelter, clothes, and other child care.

In conclusion, raising children changes opportunity costs, retirement savings behavior, and overall wealth. These papers and several others examine the multitude of factors that describe the intersection of raising children and the age in which one retires. While previous research focuses on the financial costs of children on families and their retirement plans, my analysis focuses on how it impacts the age that an individual retires. Exploring these impacts while controlling for other factors such as wealth, race, education, health, and others provides further insights into targeting policy and education towards individuals' retirement savings. It would be beneficial to shed light on these costs and grasp a better understanding of how these numbers should factor into a household's long term retirement plan.

Data and Methodology

The data I use in my analysis comes from the *Health and Retirement Study* (HRS) provided by The University of Michigan, a longitudinal survey of U.S. households with at least one adult age 50 or older. Every two years, respondents are surveyed about demographics, wealth, health, Social Security, employment history, and many other factors. The cohorts ranging from 1998-2020 are used to provide a comprehensive analysis of these trends over many years. The HRS over-samples Hispanics, Blacks, and residents of Florida, and provides weighting variables to make it representative of the community-based population.

I look at the intensive margin associated with the number of children parents have, and attempt to isolate the impact that each additional child has on retirement age. I restrict my sample to only households that have children, discarding non-parent observations. Non-parents may have different lifestyle choices, financial priorities and retirement goals that will affect my analysis. My goal is to create a more homogeneous sample that allows for clearer comparisons and conclusions.

The dependent variable is “Age Retired” and measures the age in which a respondent answers that they are fully retired. The HRS does not have an exact variable for this in their system, so it was constructed using two variables: one which asks respondents their age, and one which asks respondents about their retirement status. I create this variable for age retired by pinpointing the first wave in which respondents self report being fully retired. I drop individuals who reported being retired in their first wave, as it cannot be determined if their age in their first wave is when they first retired, or if they were retired for some time before their first wave.

The independent variable is “Children” and measures the number of children ever born to a respondent. It is usually asked once, and does not include step-children, adoptions, or miscarriages. For confidentiality protection, the HRS dataset limits the value of this measure to a maximum of 11.

Several control variables are included to increase the precision of my estimates. Omitting these variables could contribute to bias in my analysis, and accounting for them will improve the interpretability of my regressions. I include variables that are associated with wealth and status, such as average earnings (inflation-adjusted to 2020 dollars using year-specific Consumer Price Index values) and years of education. Those who are higher educated or average higher yearly salaries are better able to pay for the costs associated with children, as well as save for retirement. Other demographic variables are included, such as race and gender. By including demographic control variables, I can better ensure that any observed relationships between the number of children and retirement age are not simply artifacts of racial or gender differences within the sample. I also include health related control variables, such as self-reported health and whether a respondent has a heart problem. Those with poor health or other complications may retire earlier due to the possibility of lower life expectancy. On the contrary, healthier individuals do not have the same pressure to retire sooner if they expect to live longer lives. Additionally, I include religion as a control variable, as certain religious beliefs or practices can influence attitudes toward both childbearing and work. Finally, marriage status is included, as individuals who are or have been married before are more likely to have children compared to those who are single. Moreover, marriage can significantly influence financial stability, which may, in turn, affect decisions about family size and the timing of retirement.

In my analyses, I create a number of graphs to identify trends, outliers, and relationships between key variables over time. I regress the number of children an individual has and these other control variables on age that an individual retires. In addition, I examine the difference between retirees from early and late waves in two separate regressions to evaluate how the impact of having children on retirement age varies across cohorts.

Results

Table 1 exhibits the sample's variables, number of observations, their means and medians, standard deviations, and minimum/maximum values. The sample used consists of about 10,600 retired individuals. Over half of my sample consists of females, and a majority of my sample are married individuals. A large part of the sample is also Protestant, but consists of a multitude of religions. The typical number of years of education a respondent has is around twelve. The average yearly income of respondents in this sample, inflation-adjusted to reflect 2020 purchasing power, is $76,872. It is important to note that this is much larger than the median, which is $54,765. This indicates that there are some large outliers that skew the average yearly earnings to a much higher number. To help mitigate the influence of large outliers and to normalize the distribution of earnings, the log form of this variable is used in my regressions. Finally, the average self-reported health is "Good" and over 25% of the sample reports having heart problems.

In Figure 1, I plot the average retirement age for each cohort of respondents from the HRS study. The overall average retirement age among the entire sample is 64.75. Other than two dips in average retirement age, it tends to trend upwards across each cohort. Newer generations of retirees tend to have more time spent in the workforce. This could be for a variety of reasons, such as higher life expectancies of recent years allowing for people to earn more before retirement, or the culture of US adults becoming more work oriented. Since retirement age tends to increase across waves while the number of children trends downward, I introduce wave dummy variables in my regressions to account for potential time-period effects. These dummies help isolate the relationship between children and retirement by controlling for systematic differences between survey waves.

I also plot the average number of children for each cohort from the study in Figure 2. The overall average number of children for all respondents is 3.01. The average trends down for each wave. This coincides with the drop in fertility rates seen in the US, which can be for a number of reasons as reviewed in previous literature. These reasons could include the changing priorities of women today, the utility of children being viewed differently, or many others. A steady decline in the number of children born per wave shows that the incentives for adults to have children are also decreasing over time. While the joys of parenthood stay unchanged for each generation, the costs of having children - and the opportunity costs of not having children - are becoming less appealing.

Also analyzing the average retirement age for each additional child in Figure 3, one can see a general upward trend. Notably, individuals with only one child retire at a much earlier age compared to the rest of the sample. However, the upward trend is less reliable among respondents with more than five children, likely due to smaller sample sizes and greater variability, as reflected in the larger error bars. Lastly, Figure 4 presents the sample distribution of respondents by number of children. It is important to note that in my sample, the most common number of children is 2, and the sample size per number of children decreases thereafter. Despite rising retirement age and decreasing number of children for newer generations of retirees, this result shows that there is a positive relationship between them.

In Regression 1, I estimate the associated increase in retirement age for each additional child. I also include a number of control variables that may help to mitigate bias throughout my sample. All else being equal in this multiple variable regression, having an additional child is associated with an increase of 0.629 - or over seven months - in the age in which one retires. This coefficient is statistically significant at the 1% level, both with and without the inclusion of control variables. Most control variables are significant as well, indicating their effectiveness. The R-squared value suggests that this model explains about 16% of the variation in retirement age. Overall, these regression results show that parents who choose to have an additional child, on average, see an uptick in the age in which they retire in order to make up for that investment.

With regressions 2 and 3, I compare the associated impacts from the number of children on retirement age between those who responded being fully retired between waves 5-9 (1998-2008) and those who fully retired in waves 10-15 (2010-2020). For respondents in the earlier waves, the associated impact of an additional child on retirement age is 0.495, all else equal. On the contrary, for respondents in the later waves, the associated impact of an additional child on retirement age is 0.729, all else equal. Both of these coefficients are statistically significant at the 1% level, and the estimated difference between them equates to an approximate three-month differential. The effect that children have on retirement age is greater on the newer generations of retirees, compared to the older generations. This is consistent with previous findings of rising costs to have children, decreasing incentives to have children, and the declining fertility rate in the US.

Caveats

Some of the limitations in my variables have to do with how they are defined. For example, the variable "Children" leaves out step-children, adoptions, and miscarriages. While these are not as common, they can still be significant investments for some households. Leaving these responses out of the HRS survey could skew the hypothesized impact on one's retirement age since respondents who only have step-children or adoptions are not included, and the regression cannot account for those who have their own children as well as other step-children or adopted children.

Another limitation could come from individuals unintentionally misreporting certain variables, such as their retirement status. The HRS asks respondents to describe whether they are not retired, fully retired, or partly retired. These categories are ambiguous and it is possible that certain individuals interpret their retirement status differently than how the HRS would. This could lead to respondents misidentifying themselves to inaccurate categories that can misrepresent the results. This is true for other categorical variables in the HRS, such as self-reported health.

In addition, the HRS interviews its respondents every two years. This means when an individual reports being fully retired for the first time in a new wave, their actual retirement decision could have happened at any point within the two years between their last survey response and reporting being fully retired. Therefore, a respondent's reported retirement age could be up to two years past the actual age in which they reported being fully retired. This could misrepresent the retirement age of some and increase the overall average retirement age among respondents.

Also, a variable measuring Social Security wealth was omitted from my regression. Average yearly income is used as a control variable, and Social Security is factored into the income of respondents. Furthermore, people don't always claim their Social Security benefits at retirement; they do so at various ages. For this regression, which concentrates on a respondent's answers in the year they retire, and their average total income, including Social Security wealth would not provide accurate contributions to my analysis.

Finally, while my regression models control for key demographic and socioeconomic variables, it omits potentially important factors such as occupation type, spousal retirement status, and health shocks, which may influence both earnings and retirement age. The exclusion of these variables could result in omitted variable bias, potentially overstating or understating the true effect of earnings on retirement decisions.

Conclusion

Having children is an important decision to make for every family. They take a considerable amount of time, money, and effort to raise, and can come with a number of risks. All of these factors add to the already difficult task of saving for retirement years. It is possible for parents to navigate these extra expenses and still reach their retirement goals. Being able to plan for the investment of a child and understanding its impacts on retirement plans can assist lots of families in this decision making process. However, it is easier said than done. My analysis shows that all else equal, having an additional child is associated with a slight increase in full retirement age. This result is indicative of my hypothesis that - in myopic fashion - parents tend to pay for their lack of savings by delaying retirement. While my findings are significant, they are also not strong enough to claim as causal. An important limitation of my analysis is the potential endogeneity between the number of children and retirement age. Individuals with more children may have other characteristics, such as a stronger work ethic or different financial habits, that lead them to retire later, regardless of family size. Thus, while my results offer insight into understanding this relationship, they may not fully capture the complexities of this dynamic.

Understanding this relationship can inform couples and policymakers about the potential impact of larger families on retirement patterns. In the government's efforts to try to increase and maintain sufficient fertility rates, it may consider policies or programs that both encourage child-rearing and also incentivize saving for retirement. It would be easy if institutions could just simply stress the importance of saving early and often. Nevertheless, saving optimally is difficult for the large majority of families, so some intervention may help. Since there is evidence that the introduction of Social Security has contributed to falling fertility rates, it can be edited to incentivize having children. One idea is to provide financial credits toward an individual's future Social Security benefits for each child they raise. For example, there could be small incremental increases in one's Social Security benefit. Another idea could be a contribution in the form of a "caregiver credit." This could be some type of bond or small investment contribution that becomes liquid when one reaches their full retirement age. This way, the government can lessen fears adults may have of the costs of having children, and help alleviate the financial burden associated with it. The concern with financial incentives would be that prospective parents see them as a replacement for retirement saving, rather than a supplement. Grasping the relationship between the decision to have children and its impact on retirement patterns is crucial for informing both couples and policymakers. Implementing policies that incentivize both child-rearing and retirement savings can help alleviate the financial burden associated with raising children and ensure a stable future for many families across the US.

References

Biggs, Andrew G., Chen, Anqi, and Munnell, Alicia H. "How do Households Adjust Their Earnings, Saving, and Consumption After Children Leave?". CRR WP 2021-20 November 2021. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Biggs, Andrew G., "How do Children Affect the Need to save for Retirement?". AEI Economics Working Paper 2019-23 December 2019. American Enterprise Institute.

Boldrin, Michele, De Nardi, Mariacristina, and Jones, Larry E. "Fertility and Social Security". Journal of Demographic Economics, 81, 2015, 261–299. Cambridge University Press.

Bradbury, Bruce. "Time and the Cost of Children." The Review of Income and Wealth, vol. 54, no. 3, 4 Aug. 2008, pp. 305–323.

Health and Retirement Study, (RAND HRS Longitudinal File 2020 (V1)) public use dataset. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740). Ann Arbor, MI, (March 2023).

Kearney, Melissa S. and Levine, Phillip B. "The Causes and Consequences of Declining U.S. Fertility." Economic Policy in a More Uncertain World January 2023. Aspen Economic Strategy Institute.

Munnell, Alicia H., Hou, Wenliang, and Sanzenbacher, Geoffrey T. "The Impact of Raising Children on Retirement Security". CRR IB 17-16 September 2017. Chestnut Hill, MA: Center for Retirement Research at Boston College.

RAND HRS Longitudinal File 2020 (V1). Produced by the RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration. Santa Monica, CA (March 2023).

Rutledge, Matthew S. and Lindner, John E. "Do Late-Career Wages Boost Social Security More for Women than Men". CRR WP 2016-13 November 2016. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Scholz, John K and Seshardi, Ananth. "Children and Household Wealth". WP 2007-158 October 2007. University of Michigan Retirement Research Center.

Appendix

Table 1

Variable	Obs.	Mean	Median	Std. dev.	Min	Max
Age Retired	10,694	64.74911	65	6.203475	31	88
Children	10,644	3.012777	3	1.692744	1	11
Mean Earnings	10,694	76872.31	54765.07	123254.8	0	8033232
Years of Education	10,664	12.39704	12	3.14006	0	17
Self-reported Health	10,686	3.005896	3	1.12431	1	5
Heart Problems	10,683	0.2621232	0	0.5226859	0	6
Female	10,694	0.5970638	1	0.4905111	0	1
White	10,671	0.7168025	1	0.4505726	0	1
Black	10,671	0.2050417	0	0.403751	0	1
Other Race	10,656	0.0781557	0	0.2684291	0	1
Protestant	10,656	0.6266892	1	0.4837063	0	1
Catholic	10,656	0.2637012	0	0.4406598	0	1
Jewish	10,656	0.013795	0	0.1166448	0	1
No Religion	10,656	0.0789227	0	0.2696307	0	1
Other Religion	10,656	0.0168919	0	0.1288725	0	1
Married	10,677	0.6794043	1	0.4667274	0	1
Divorced	10,677	0.1580968	0	0.3648489	0	1
Widowed	10,677	0.1370235	0	0.3438883	0	1
Never Married	10,677	0.0254753	0	0.1575711	0	1

Figure 1

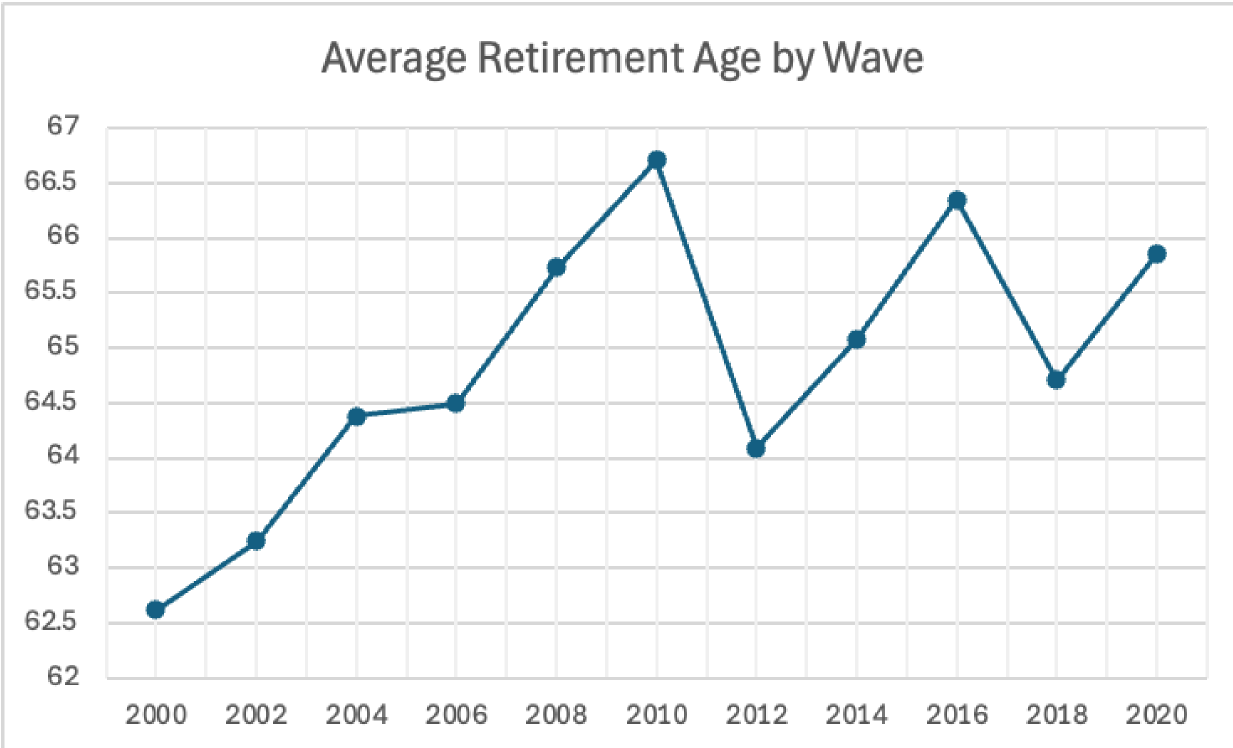

Figure 2

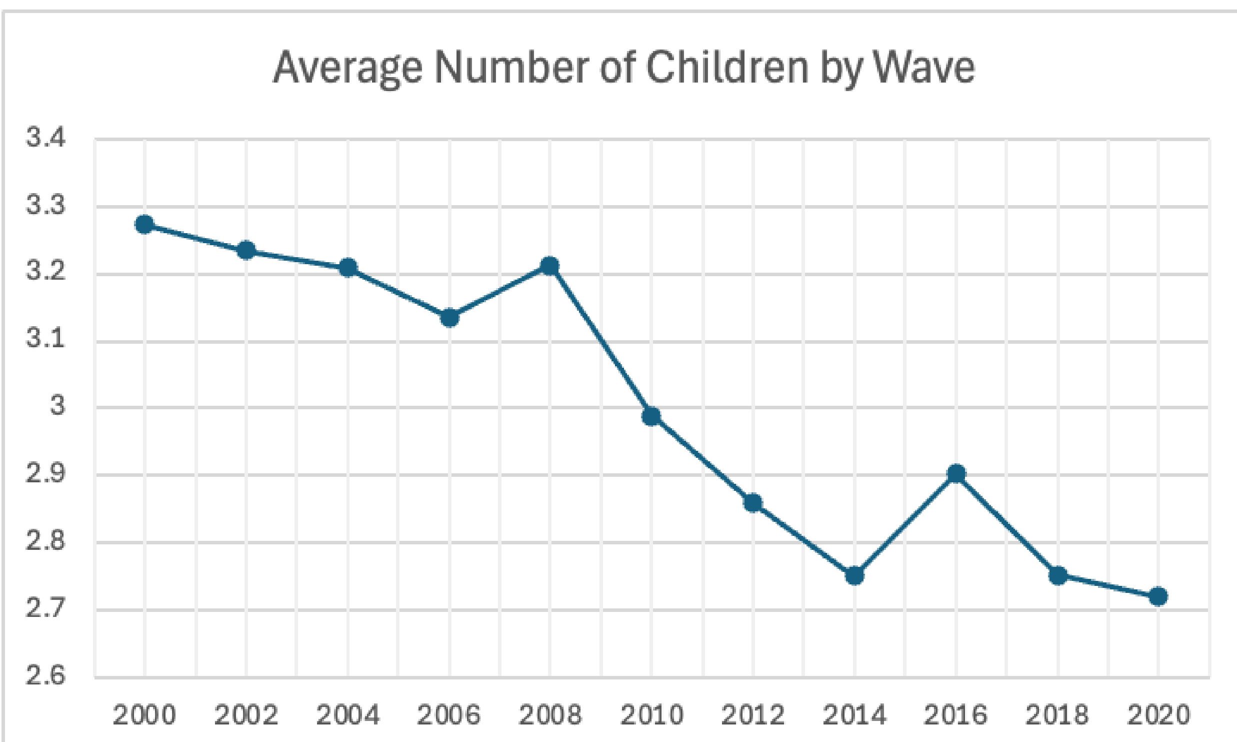

Figure 3

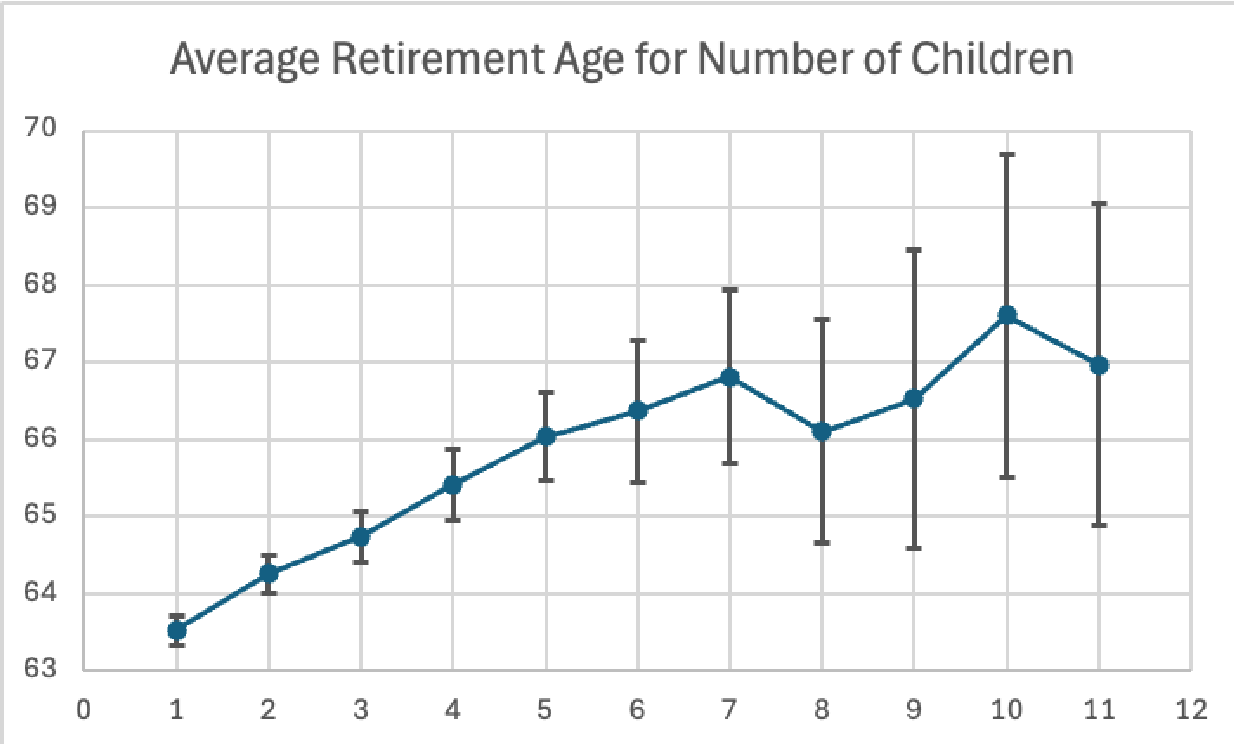

Figure 4

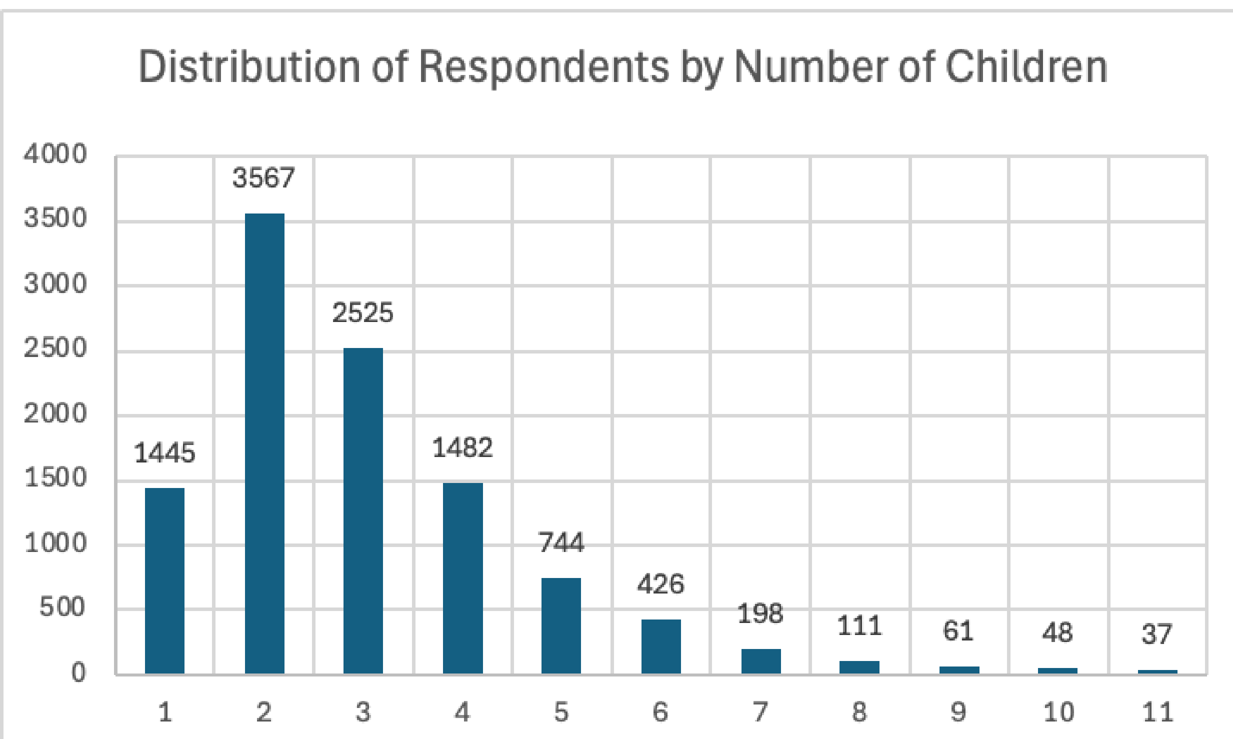

Regression 1

REGRESSION 1		
VARIABLES	Age Retired Coefficient	Standard Error
Children	0.629***	0.0353
log(Average Earnings)	0.760***	0.0868
Years of Education	-0.13***	0.022
Self-reported Health	-0.626***	0.0553
Heart problems	0.706***	0.1061
Female	-1.141***	0.1177
Black	-1.796***	0.1536
Other Race	-2.152***	0.2207
Catholic	-0.166	0.1358
Jewish	1.921***	0.4793
No Religion	-1.887***	0.2138
Other Religion	-1.279***	0.438
Divorced	0.534***	0.1722
Widowed	3.855***	0.18771
Never Married	-0.041	0.3824
Constant	56.21***	0.9735
Observations	10,513	
R-squared	0.1617	
Wave Dummies?	Yes	
*** p<0.01	** p<0.05	* p<0.1

Regression 2 & 3

WAVE SPLIT	REGRESSION 2 (WAVES 5-9, 1998-2008)		REGRESSION 3 (WAVES 10-15, 2010-2020)	
VARIABLES	Age Retired Coefficient	Standard Error	Age Retired Coefficient	Standard Error
Children	0.495***	0.0412	0.729***	0.0569
log(Average Earnings)	-0.358**	0.1204	1.292***.	0.1223
Years of Education	-0.0549**	0.0275	-0.141***	0.0333
Self-reported Health	-0.569***	0.0664	-0.693***	0.0864
Heart problems	0.347***	0.132	0.944***	0.1601
Female	-1.082***	0.1448	-1.228***	0.1796
Black	-1.201***	0.1968	-2.248***	0.2268
Other Race	-1.671***	0.3387	-2.367	0.2948
Catholic	-0.164	0.1627	-0.134	0.212
Jewish	1.712***	0.6137	1.944***	0.7064
No Religion	-1.209***	0.3061	-2.147***	0.2945
Other Religion	-1.12	0.756	-1.329**	0.5579
Divorced	0.262	0.2186	0.605**	0.2565
Widowed	2.477***	0.2207	4.701***	0.2685
Never Married	0.985	0.625	0.091	0.4991
Constant	67.616***	1.309	54.445***	1.426
Observations	4,991		5,522	
R-squared	0.1304		0.1830	
Wave Dummies?	Yes		Yes	
	*** p<0.01	** p<0.05	* p<0.1	

Thank You, Mom and Dad:
An Analysis of Parents' Education and Children's College Enrollment

By Zhanshen Weng & Bojun Zhang

I. Introduction

Educational attainment has long been considered a critical pathway to economic opportunity and social mobility. However, disparities in college attendance rates persist across various demographic and socio-economic groups, highlighting deep-seated inequalities in access to higher education. Among the many factors that influence college attendance, parental education stands out as a significant predictor, often intertwined with household income, family structure, and the academic performance of children. Understanding the independent impact of parental education on college attendance, while controlling for other key factors such as parental income and students' high school GPA, is crucial for addressing the root causes of educational inequality.

This paper investigates how parents' educational attainment affects their children's likelihood of enrolling in college, using data from the High School Longitudinal Study of 2009 (HSLS:09). By employing a probit regression model, we analyze the relationship between parental education and college enrollment, accounting for variables that represent family income, household size, demographic characteristics, and student ability. Specifically, our analysis focuses on whether having a mother or father with education above a bachelor's degree increases the probability of college attendance when holding other factors constant.

The relevance of this question extends beyond academic interest. Educational inequality perpetuates broader societal inequities, influencing income distribution, social mobility, and labor market outcomes. Policymakers and educators need empirical evidence to design interventions that address barriers to college access. By isolating the role of parental education, this study contributes to a more nuanced understanding of how family background shapes educational opportunities and outcomes.

The discussion proceeds as follows. The first section discusses the importance of education and existing research on how children's educational attainment is affected by family conditions, especially parents' education. The second section presents a description of the HSLS dataset used, the econometric methods used, and a summary of the stages of the analysis. The third section breaks down the three stages of our analysis, going from showcasing the education level of parents of those enrolled or not enrolled in college, to the impact of any family member's highest education level, and the impact of parents' education level on children's college enrollment. The final section concludes that, controlling for demographic factors and children's academic abilities, parents' education level has a significant impact on children's college enrollment rate.

II. Background

According to Bourdieu's theory of social reproduction, social inequalities are based on the effective transmission of family-based endowments to the offspring, including physical, human, and, especially, cultural capital. Cultural capital is the accumulated cultural knowledge that confers power and status. Education is a major form of cultural capital that plays an important role in creating inequalities (Bourdieu, 1977, 1984).

There are three forms of cultural capital: embodied cultural capital (ingrained cultural knowledge and skills, such as refined language use), objectified cultural capital (cultural objects and goods, such as artworks and books), and institutionalized cultural capital (formal recognition, such as academic credentials or awards). Sullivan (2001) illustrates how these forms of cultural capital interact within the education system to create a cycle of cultural reproduction and empirically confirms such a cycle using a sample of 465 pupils across four different schools in England. First, parental cultural capital is inherited by children. Then, children's cultural capital is converted into educational credentials. Finally, educational credentials serve as a mechanism of cultural reproduction.

This demonstrates that cultural capital confers individuals with educational advantages, and these advantages can be transmitted across generations. In this issue brief, we will explore whether parental education translates into educational advantages for their children, thereby testing the validity of Bourdieu's theory of cultural capital. Previous research and data generally support this idea, showing that parental education is a factor in shaping children's academic outcomes.
Data from the National Center for Education Statistics (NCES) offers great insights into how

parents' education affects their children. The 2018 report points out that among 2002 high school sophomores tracked, 72% of students whose parents had never attended college enrolled in postsecondary education by 2012. In comparison, 84% of their peers whose parents had some college education had done so, as had 93% of those whose parents had earned a bachelor's degree.

Dubow, Boxer, and Huesmann (2009), in their study on the long-term effects of parents' education on children's educational and occupational success, use data from the Columbia County Longitudinal Study, which collected 4 waves of data over 40 years on children who were living in Columbia County, NY in 1960. They indicate that for both genders, parents' educational level during middle childhood was positively correlated with educational aspirations and educational level during late adolescence and educational level during the adult outcomes stage. They also found that although parental educational level during childhood had no direct effects on children's educational level at age 48, significant indirect effects exist through age 19 educational aspirations and educational level. Simply put, children with more highly educated parents developed higher aspirations for their own education and had a higher educational level by age 19, which led to higher levels of adult educational attainment.

Haveman and Wolfe (1998), on the other hand, provide an overview of the factors influencing children's educational, social, and economic outcomes. They point out that "on average, parents with levels of educational attainment far above the mean will produce children who attain high levels of schooling." They state that parental education affects children both directly, through academic expectations and the quality of interactions, and indirectly, through the economic resources it generates. They also found out that being raised by a single parent or stepparents hurts educational attainment, especially for African Americans.

This study seeks to further investigate the impact of parental education on children's educational attainment by utilizing an up-to-date data source to examine recent trends in educational inequality. In addition, unlike most previous research, which typically considers parental education as a unified factor, this analysis adopts a more nuanced approach by exploring whether the educational attainment of fathers and mothers exerts differential effects on the educational outcomes of the next generation, offering a more comprehensive perspective on this issue.

III. Data and Methodology

In this study, we use the High School Longitudinal Study of 2009 (HSLS:09) as the source of data. HSLS:09 is a nationally representative, longitudinal survey of 9th-grade students in the U.S., designed to track their trajectories through high school, colleges, and beyond. The data includes students' academic progress and family background, collected not only from students but also from parents, teachers, and school administrators, offering multiple perspectives on students' experiences. This issue brief primarily utilizes data collected during four key points of the study: the Base Year (BY), First Follow-Up (F1), High School Transcript (HST), and Post-Secondary Transcript (PST) phases.

The analysis will proceed in three stages. In the first stage of the analysis, students are categorized based on their parents' highest level of education and their own college enrollment status. This categorization aims to explore the potential relationship between parental education levels and students' likelihood of college enrollment. Specifically, if parental education significantly influences college enrollment, we would expect to observe distinct distributions of parental education levels between students who enrolled in college and those who did not. This stage serves as a foundational step to examine how family background factors contribute to educational outcomes.

The simple comparison, however, does not account for the influence of family demographics and students' academic abilities on college enrollment. Therefore, in the second stage, regression analysis is employed to address these confounding factors. Specifically, we utilize probit/logit models to estimate the likelihood of students enrolling in college after high school graduation as a function of their parents' highest educational attainment. These models control for key demographic variables, including family income, household size, single-parent status, gender, and race, as well as academic ability, as measured by students' high school GPA. This approach allows us to isolate the effect of parental education while accounting for other relevant factors. The relationship is modeled as follows:

College enrollment = f(parental education, demographics, students' academic ability)

Note that for our key independent variable, we use a binary variable that indicates whether the highest parental education is a bachelor's degree or above. This is because a bachelor's degree is often viewed as a socioeconomic milestone, marking a sharper distinction than other incremental educational levels. Using education level as a continuous variable may fail to capture the potentially nonlinear returns to education.

In the third stage of the analysis, we examine the distinct effects of mothers' and fathers' educational attainment—specifically, whether they have attained a bachelor's degree or higher—on children's educational outcomes. This analysis utilizes the same set of controls as in the second stage but replaces the highest parental education variable with separate binary indicators for each parent's educational level. This approach accounts for potential heterogeneity within households, recognizing that mothers' and fathers' educational backgrounds may influence children differently based on their roles or resources.

IV. Results

Stage 1

Figure 1 illustrates the distribution of children's college enrollment status based on their parents' highest level of education. The size of the circles represents the number of data points, with darker colors indicating a higher count. From the figure, it is evident that among students who enrolled in college, the majority have parents whose highest educational attainment is a bachelor's degree. Conversely, for students who did not enroll in college, the majority of their parents have a high school education as their highest level of attainment.

Additionally, for education levels above a bachelor's degree (e.g., master’s degree, doctorate), the number of parents whose children enrolled in college is consistently higher compared to parents of students who did not enroll in college. This trend suggests a potential positive association between higher parental education levels and children’s likelihood of college enrollment.

Figure 1 *Children’s College Enrollment by Parents’ Education Level*

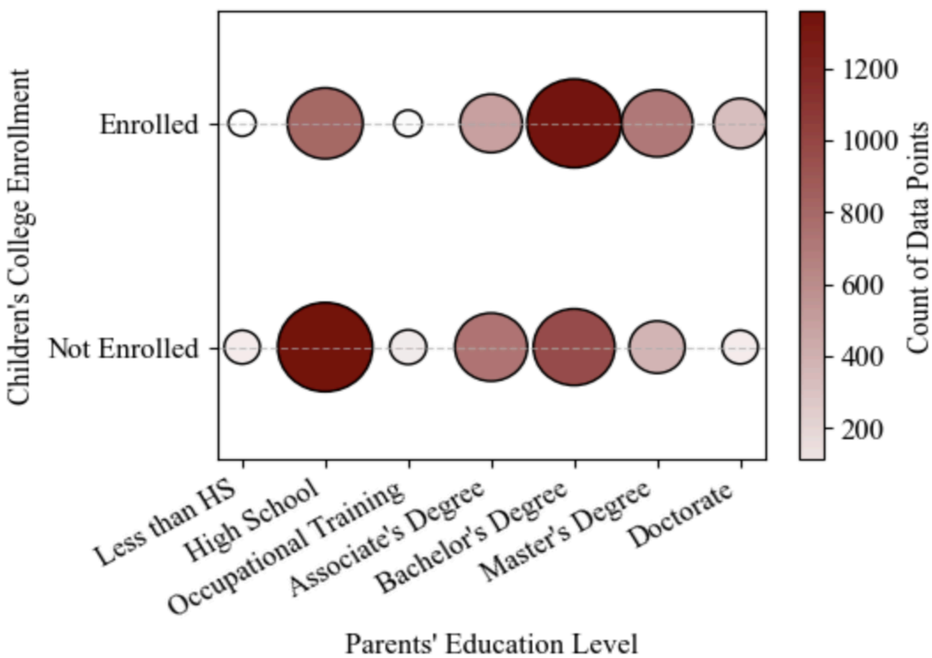

[1]

Source: High School Longitudinal Study of 2009

Stage 2

Based on Table 1, having parents with an educational attainment at or above a bachelor's degree is associated with a 6 percentage point increase in the likelihood of enrolling in college after high

[1] To ensure an unbiased analysis, the data was processed by randomly sampling equal numbers of observations from both enrolled and unenrolled groups. Specifically, the smaller group's size determined the sample size to maintain proportional representation, thereby mitigating potential overrepresentation of the larger group in the analysis.

school graduation, even after controlling for demographics and academic ability. This finding regarding the role of parental education in shaping college enrollment decisions aligns with theories of cultural capital, which posit that knowledge, skills, and education are transmitted across generations. This conclusion is robust across both probit and logit models and yields similar results.

There are two possible explanations for this phenomenon. First, parents with higher educational attainment often have elevated expectations for their children's academic achievements, which may encourage a stronger inclination toward college enrollment. Research suggests that parents with higher levels of education tend to set higher educational aspirations for their children (Rimkute et al., 2011). Additionally, parental expectations have been identified as a key factor influencing students' predisposition to attend college, serving as one of the strongest predictors of enrollment decisions (Stage & Hossler, 1989).

Second, parents with higher education levels often serve as an important source of information during the college application process. Even if their own college experience occurred decades ago, they are likely to have greater familiarity with the application process compared to parents who did not attend college. This familiarity enables them to provide practical knowledge-based support, in addition to emotional and financial support (Ceja, 2006). Furthermore, when students apply to programs of study related to the fields in which their parents hold degrees, parents tend to exert significant control over the decision-making process (Eldegwy et al., 2022).

Table 1 *Regression with highest parental education level*

VARIABLES	(1) Probit - Coef College Enroll	(2) Probit - Avg Margins College Enroll	(3) Logit - Coef College Enroll	(4) Logit - Avg Margins College Enroll
Parent Bachelor	0.2014***	0.0664***	0.3457***	0.0683***
	(0.0330)	(0.0107)	(0.0562)	(0.0108)
Single Parent	-0.0999***	-0.0337***	-0.1685***	-0.0343***
	(0.0310)	(0.0106)	(0.0511)	(0.0105)
Female	0.0143	0.0048	0.0261	0.0052
	(0.0260)	(0.0087)	(0.0433)	(0.0087)
Family Income 1	0.1053***	0.0351***	0.1710***	0.0343***
	(0.0331)	(0.0110)	(0.0542)	(0.0108)
Family Income 2	0.4090***	0.1330***	0.6875***	0.1336***
	(0.0408)	(0.0127)	(0.0682)	(0.0125)
Family Size	-0.0239***	-0.0080***	-0.0409***	-0.0082***
	(0.0091)	(0.0030)	(0.0151)	(0.0030)
Race	-0.0495***	-0.0165***	-0.0840***	-0.0168***
	(0.0060)	(0.0020)	(0.0101)	(0.0020)
GPA	0.5728***	0.1911***	0.9468***	0.1900***
	(0.0206)	(0.0061)	(0.0349)	(0.0061)
Constant	-1.1196***		-1.8398***	
	(0.0809)		(0.1346)	
Observations	11,076	11,076	11,076	11,076
Pseudo R2	0.0946	0.0946	0.0947	0.0947

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1 [2]

Source: Author's calculation using High School Longitudinal Study of 2009

Stage 3

Both probit and logit models, again, yield consistent findings regarding the differing impacts of maternal and paternal education on college enrollment (See Table 2). Specifically, having a mother with a bachelor's degree or higher is associated with a 3.6 to 3.8 percentage point increase in the likelihood of college enrollment, whereas having a father with a bachelor's degree or higher is associated with a 7.2 to 7.7 percentage point increase. This indicates that a father's attainment of a bachelor's degree has approximately twice the impact on a child's likelihood of college enrollment compared to a mother's degree, based on the Average Marginal Effects (AME) results.

[2] Family Income 1 represents families with annual incomes between $35,000 and $115,000 while Family Income 2 represents families with annual incomes above $115,000. The base category consists of families with annual incomes of $35,000 or less.

Table 2 *Regression with mother/female guardian's and father/male guardian's education level*

VARIABLES	(1) Probit - Coef College Enroll	(2) Probit - Avg Margins College Enroll	(3) Logit - Coef College Enroll	(4) Logit - Avg Margins College Enroll
Female Guardian Bachelor	0.1101***	0.0363***	0.1921***	0.0379***
	(0.0405)	(0.0131)	(0.0691)	(0.0134)
Male Guardian Bachelor	0.2215***	0.0721***	0.3987***	0.0774***
	(0.0422)	(0.0133)	(0.0735)	(0.0137)
Single Parent	-0.0883***	-0.0298***	-0.1480***	-0.0301***
	(0.0312)	(0.0106)	(0.0513)	(0.0105)
Female	0.0146	0.0049	0.0267	0.0054
	(0.0260)	(0.0087)	(0.0433)	(0.0087)
Family Income 1	0.1100***	0.0366***	0.1780***	0.0357***
	(0.0331)	(0.0110)	(0.0542)	(0.0108)
Family Income 2	0.4056***	0.1319***	0.6801***	0.1321***
	(0.0408)	(0.0127)	(0.0682)	(0.0125)
Family Size	-0.0235***	-0.0079***	-0.0405***	-0.0081***
	(0.0091)	(0.0030)	(0.0151)	(0.0030)
Race	-0.0490***	-0.0163***	-0.0830***	-0.0167***
	(0.0060)	(0.0020)	(0.0101)	(0.0020)
GPA	0.5711***	0.1905***	0.9436***	0.1893***
	(0.0206)	(0.0061)	(0.0349)	(0.0061)
Constant	-1.1216***		-1.8431***	
	(0.0809)		(0.1345)	
Observations	11,076	11,076	11,076	11,076
Pseudo R2	0.0948	0.0948	0.0951	0.0951

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

[3]

Source: Author's calculation using High School Longitudinal Study of 2009

The first reason may be linked to the previously discussed relationship between parental education and involvement in the college application process. As discussed in Stage 2, parents with higher educational attainment are often more involved in their children's college application process due to their elevated academic expectations and their ability to provide application-related information and resources, which ultimately enhances the likelihood of college enrollment. A possible explanation for the stronger influence of fathers' education could be that fathers, on average, tend to have higher educational attainment than mothers. Consequently, fathers may hold greater decision-making authority during the application process. As shown in Figure 2, the proportion of fathers with higher education levels than mothers (42.8%) significantly exceeds the proportion of mothers with higher education levels than fathers (27.4%).

[3] Family Income 1 represents families with annual incomes between $35,000 and $115,000 while Family Income 2 represents families with annual incomes above $115,000. The base category consists of families with annual incomes of $35,000 or less.

If the observed difference in influence is primarily attributable to unequal educational attainment between parents, there is reason to believe that this disparity may diminish over time. According to the National Center for Education Statistics (NCES), female college enrollment rates among 18- to 24-year-olds have exceeded male enrollment rates since the early 1990s. In 2022, the female enrollment rate was 43.8%, nearly 10 percentage points higher than the male enrollment rate of 34.2%. While this comparison is based on overall gender trends and does not account for the non-random nature of family dynamics, the disproportionate influence of fathers' education on college enrollment—stemming from their educational advantage—holds potential for improvement as women continue to outpace men in college attendance.

Figure 2 *Distribution of Highest Education Attained*

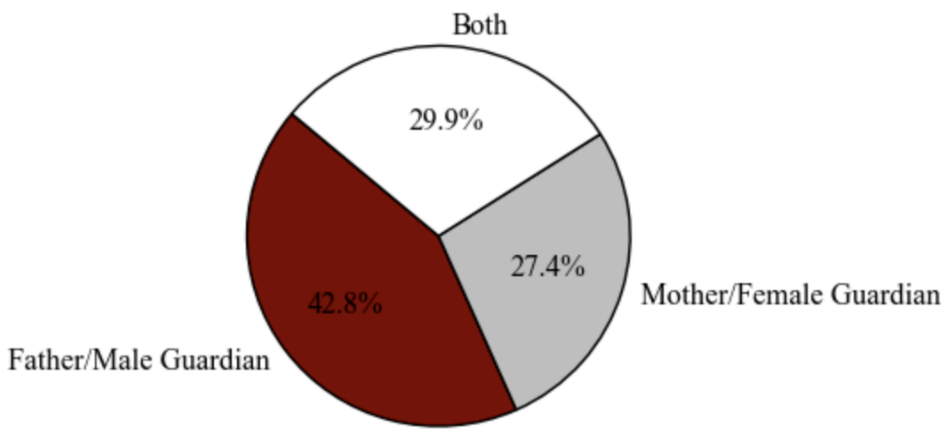

Source: Author's calculation using High School Longitudinal Study of 2009

In contrast, a less optimistic explanation may lie in the influence of traditional gender roles within families. Education represents an investment in human capital (Becker 1967) and involves significant financial considerations such as cost (McPherson and Schapiro 1991), financial aid (Linsenmeier et al. 2006), and student loans (Ionescu 2009). Financial decisions, however, are often viewed as a male-dominated domain. Research indicates that women are less likely to participate in financial decision-making due to gender identity norms, rather than differences in risk preferences, confidence, or financial knowledge (Ke, 2021). This suggests that even when mothers possess financial literacy comparable to fathers—which is often associated with higher education levels (Lusardi & Mitchell, 2014)—they may step back from such decisions due to entrenched gender norms.

Unfortunately, such norms persist. Research by Haines et al. (2016) highlights that perceptions of gender stereotypes have remained largely unchanged over time, with men and women still being assigned distinct roles. For instance, Bartley et al. (2005) found that wives spent 88% of their household labor on low-control tasks, compared to only 51% for husbands.

V. Conclusion

This paper examines the relationship between parental educational attainment and children's educational outcomes and its implications for societal inequality. Using the HSLS:09 dataset, the analysis finds that having parents with at least a bachelor's degree increases the likelihood of college enrollment by 6 percentage points, even after controlling for demographic and academic factors. Notably, the influence of fathers' education is nearly double that of mothers', potentially due to the higher college attendance rates of fathers in this generational cohort, which may enhance their familiarity with the college application process. Alternatively, this disparity could stem from the persistence of traditional gender roles, where fathers are more likely to dominate significant financial decisions, such as supporting children's college education.

Two key forms of inequality, thus, emerge from this study. First, intergenerational educational inequality reflects the transmission of cultural capital across generations. Parents with greater cultural capital provide their children with distinct advantages, increasing their likelihood of obtaining academic credentials and perpetuating this cycle across generations. The other is the unequal parental influences within households. Fathers' education appears to have a disproportionately larger impact on children's college enrollment compared to mothers', underscoring the gendered dynamics of parental roles in shaping educational outcomes.

Although these inequalities are difficult to tackle in a short period of time, educational institutions could help remedy this issue by establishing support mechanisms for first-generation college students. For example, targeted secondary school counseling programs for students without familial higher education backgrounds.

There are, however, two potential limitations in our research. First, the use of a binary variable (bachelor's degree or higher vs. less than a bachelor’s degree) may oversimplify the nuanced effects of varying parental education levels on college enrollment. For instance, the impact of a doctorate might differ significantly from that of a bachelor's degree. Second, despite controlling for socioeconomic status, there remains a risk of omitted variable bias. Factors such as the quality of schools attended, which may strongly influence college enrollment, are not accounted for in the analysis.

Bibliography

Bartley, Sharon J., Priscilla W. Blanton, and Jennifer L. Gilliard. 2005. “Husbands and Wives in Dual-Earner Marriages: Decision-Making, Gender Role Attitudes, Division of Household Labor, and Equity.” *Marriage & Family Review* 37(4): 69–94.

Becker, Gary Stanley. 1967. *Human capital and the personal distribution of income: An analytical approach*. Ann Arbor, MI: Institute of Public Administration.

Bourdieu, Pierre. 1984. *Distinction: A Social Critique of the Judgment of Taste.* Cambridge, MA: Harvard University Press.

Bourdieu, Pierre. 2018. “Cultural Reproduction and Social Reproduction.” In *Knowledge, Education, and Cultural Change*, edited by Brown, Richard, 71-112. London, UK: Routledge.

Ceja, Miguel. 2006. “Understanding the Role of Parents and Siblings as Information Sources in the College Choice Process of Chicana Students.” *Journal of College Student Development* 47(1): 87-104.

Dubow, Eric F, Paul Boxer, and L Rowell Huesmann. 2009. “Long-term Effects of Parents' Education on Children's Educational and Occupational Success: Mediation by Family Interactions, Child Aggression, and Teenage Aspirations.” *Merrill-Palmer quarterly* 55(3): 224-249.

Eldegwy, Ahmed, Tamer H. Elsharnouby, and Wael Kortam. 2022. “Like Father like Son: The Role of Similar-Education Parents in Their Children’s University Choice.” *Journal of Marketing for Higher Education* 34(2): 458–77.

Haines, Elizabeth L., Kay Deaux, and Nicole Lofaro. 2016. “The Times They Are A-CHanging ... or Are They Not? A Comparison of Gender Stereotypes, 1983–2014.” *Psychology of Women Quarterly* 40(3): 353–63.

Haveman, Robert, and Barbara Wolfe. 1995. “The Determinants of Children's Attainments: A Review of Methods and Findings.” *Journal of Economic Literature* 33(4): 1829-1878.

Ionescu, Felicia. 2008. “The Federal Student Loan Program: Quantitative Implications for College Enrollment and Default Rates.” *Review of Economic Dynamics* 12(1): 205–231.

Ke, Da. 2021. “Who Wears the Pants? Gender Identity Norms and Intrahousehold Financial Decision‐Making.” *The Journal of Finance* 76(3): 1389–1425.

Linsenmeier, David M., Harvey S. Rosen, and Cecilia Elena Rouse. 2006. “Financial Aid Packages and College Enrollment Decisions: An Econometric Case Study.” Working Paper No.76. Cambridge, MA: National Bureau of Economic Research.

Lusardi, Annamaria, and Olivia S. Mitchell. 2014. “The Economic Importance of Financial Literacy: Theory and Evidence.” *Journal of Economic Literature* 52 (1): 5–44.

McPherson, Michael S., and Morton Owen Schapiro. 1991. “Does Student Aid Affect College Enrollment? New Evidence on a Persistent Controversy.” *The American Economic Review* 81(1): 309–318

National Center for Education Statistics. 1970-2022. *Percentage of 18- to 24-Year-Olds Enrolled in College, by Level of Institution and Sex and Race/Ethnicity of Student: 1970 through 2022.* Table 302.60. Washington. DC.

National Center for Education Statistics. 2011. *High School Longitudinal Study of 2009 (HSLS: 09): Base-Year Data File Documentation.* Report No. NCES-2011-328. Washington DC.

National Center for Educational Statistics. 2018. *First-Generation Students – College Access Persistence, and Postbachelor's Outcomes*. Report No. NCES-2018-421. Washington, DC.

National Center for Education Statistics, *High School Longitudinal Study of 2009 (HSLS: 09): Base-Year Data File Documentation*. 2011. Washington, DC.

Rimkute, Laura, Riikka Hirvonen, Asko Tolvanen, Kaisa Aunola, and Jari-Erik Nurmi. 2011. "Parents' Role in Adolescents' Educational Expectations." *Scandinavian Journal of Educational Research* 56(6): 571–590.

Stage, Frances K., and Don Hossler. 1989. "Differences in Family Influences on College Attendance Plans for Male and Female Ninth Graders." *Research in Higher Education* 30(3): 301–315.

Sullivan, Alice. 2001. "Cultural Capital and Educational Attainment." *Sociology* 35(4): 893–912.

“Rent Deflation from Vacant Taxation: Vancouver’s Empty Homes Tax”*

Marco Luo†

May 1, 2025

Abstract

Vacancy taxes have been proposed in a number of regions around the world to make housing more affordable. I model the effect using a search-and-match framework, predicting that vacancy taxes decrease equilibrium rental prices. I measure the effect of Vancouver’s Empty Homes Tax (EHT) on rental market prices using a difference-in-differences approach. I use a panel of rental prices for Vancouver — which adopted the EHT in 2017 at a 1% rate and later raised the rate to 1.25% in 2020 and 3% in 2021 — and a panel of rental prices for the neighboring Greater Vancouver Region. I create two model specifications that examine the EHT’s general heterogeneous treatment effects and its effects by number of bedrooms. I find a significant negative effect on rental prices for the EHT’s 1.25% and 3% rates, with decreases of 1.1% and 1.8%, respectively. This corroborates the findings of previous literature. However, I find no statistically significant effect at the 1% level, suggesting that the EHT is effective only at higher rates. Normatively, my findings suggest that the tax can serve as an effective curb on rising rental prices, potentially improving housing affordability.

*I would like to thank my two advisors, Professor Grubb and Professor McCullagh for their invaluable encouragement, knowledge, and guidance. I’ve thoroughly enjoyed our conversations and our time spent connecting on a personal level. I extend my gratitude for CoStar Group Inc. for giving me access to their product and platform. Lastly, I would also like to thank Mom, Dad, and Vicky for helping me be the person I am today.

†Boston College. Email:langxuan.luo@bc.edu

1 Introduction

As a way to tackle the issues of affordability, accessibility, and speculation, Vancouver's city council enacted the Empty Homes Tax (hereafter EHT), in November of 2016. Vancouver's implementation of a vacancy tax is not novel. France adopted vacancy taxes in 1999, taxing housing vacancies in densely populated areas (Segú (2020)). Since then, many other jurisdictions around the world have enacted similar policies, including Washington, DC (2010), and the United Kingdom (2015). Others, such as New York City and Toronto, have proposed vacancy tax policies in their legislation. Clearly, as housing supply and prices become an issue for many governments worldwide, vacancy taxes have emerged as a widely considered solution.

As a city famous for its green glass apartment complexes, Vancouver is also famous as one of the most unaffordable housing markets in the world. With a house-price to income ratio of 12.3 in 2023, it ranks third in the world, ahead of cities like New York City and San Francisco (Cox, 2024). Rental units have also become increasingly expensive, with average prices for apartments in multi-family homes increasing by over 60% in the past decade (CoStar Group Inc., 2013-2024). Comparatively, median household income has only increased by around 40% in a similar time frame, resulting in an increasingly higher proportion of household income committed towards paying for housing (Statcan, 2024). This is reflected in Census data, with around a third of all households spending more than 30% of their income on shelter (Government of Canada, 2022). Simply, Vancouver's housing market has become unaffordable for many.

Housing accessibility is also a pressing issue. With a fast growing population and jobs becoming increasingly white collar (Government of Canada, 2022), competition for housing in central and urban districts of Vancouver is fierce, with a vacancy rate of only 1.8% (CoStar

Group Inc., 2013-2024). For comparison, the city of Boston has a vacancy rate of around 4.5% and New York City has a rate of around 1.4%. Vancouver's tight housing market is worsened by speculative investors, who have contributed to the overvaluation and scarcity of homes in the city's most dense areas (City of Vancouver, 2017). Just as affording a home is getting more impossible, getting access to renting one is difficult too.

Vancouver's EHT taxes properties deemed "vacant" – empty for cumulatively more than 6 months of the fiscal year (City of Vancouver, 2018). The rate began at 1% of a home's taxable assessed value annually in 2017, but increased to 1.25% in 2020 and 3% in 2021, as shown in Figure 1. The EHT's purpose was to expand the city's rental housing supply, limit speculative investment, and contribute to the overall goal of "[ensuring] that renters have access to safe, secure, and affordable rental housing" (City of Vancouver, 2018). In the wider province of British Columbia, the provincial government separately enacted the Foreign Buyer's Tax in 2016 (hereafter FBT) and the Speculation and Vacancy Tax in 2019 (hereafter SVT). These taxes affect all municipalities within the province, including the city of Vancouver, while the EHT was Vancouver City specific. Homeowners in Vancouver City are subject to all taxes and must declare for each tax separately.

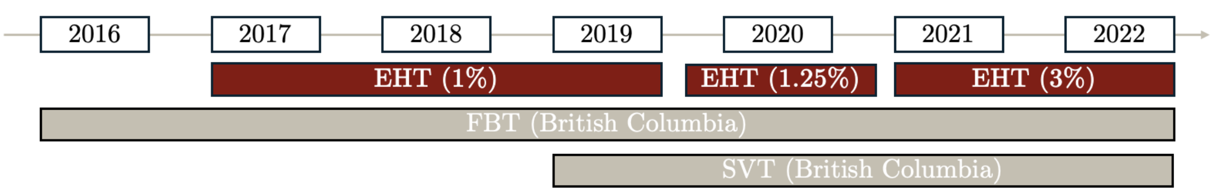

Figure 1: Timeline of the taxes

Identifying the impacts of the EHT on the Vancouver rental housing market is critical for several reasons. First, vacancy taxes serve as a market intervention that allows governments to generate revenue as a source of income to invest in future affordable housing initiatives. This contrasts with traditional forms of affordable housing policy that focus on expenditures such as subsidizing the development of new housing stock or providing direct cash assistance to tenants (Thakral, 2017). Determining whether taxation on vacancy can effectively increase rental market supply and subsequently decrease prices can offer valuable insights into market design and the flexibility of options available to governments in addressing housing crises.

The second reason to investigate Vancouver's EHT is that many cities have looked to Vancouver's implementation as a precedent, such as Oakland, which enacted a vacancy tax in 2019. By measuring the impact of the policy, this study can inform policymakers on whether a similar tax is worth adopting, especially as existing literature and theory warn of potential long-term negative welfare effects in the housing market due to vacancy taxes. Contributing to vacancy tax literature, particularly by examining a well-known case like Vancouver's EHT, will undoubtedly aid other cities in their decision-making processes.

I create a simple theoretical model using a search and match framework, derived from labor market models to complement Han et al. (2023)'s theory. I model the effects of a vacancy tax on the flow equations of tenants and landlords in a basic search and match model with rents from Nash Bargaining. My model predicts a decrease in rental prices as a result of the vacancy tax increasing the number of rental homes in the market.

The empirical method I use is as follows. I employ a difference-in-difference framework, first constructing a panel of quarterly average rental prices for studio, one-bedroom, two-bedroom, three-bedroom, and four-plus-bedroom housing units in multifamily commercial real estate buildings in Vancouver and the Greater Vancouver Region (GVR) (excluding the City of Vancouver) using time series data from CoStar Group (CoStar Group Inc., 2013-2024). This separation of geographical area can be seen in Figure 2. I also control for

Figure 2: Map of the areas of analysis: Vancouver and the GVR

seasonality and the number of bedrooms.

Using the difference-in-difference approach, I analyze the impact of each of the EHT's tax rates on the whole market as well as by unit size, accounting for other changes to the area's housing and rental markets through comparison with the GVR. This analysis enables policymakers to implement vacancy taxes more effectively by considering specific tax rates and their varying effectiveness across different types of homes.

I contribute to the existing literature on Vancouver's EHT in two ways. First, I build on Han et al.'s conclusions by directly testing the effects of the EHT on Vancouver's rental prices using a set of multifamily commercial real estate time series data from CoStar Group spanning 2013 to 2024 (CoStar Group Inc., 2013-2024). Han et al. develop a theoretical model that simulates the tax's effects on the city's rental, owner-occupied housing, and developer markets (Han et al., 2023). They then calibrate the model only for home markets. My results support their theoretical findings that rental prices decrease after the implementation of a vacancy tax. Secondly, I build on Caracciolo and Miglino (2024)'s study by using higher frequency data. Caracciolo and Miglino (2024) study the *cumulative* effect of the EHT using a difference-in-differences framework, analyzing census data from 2011, 2016, and 2021, with Burnaby—a neighboring city—as a control group. I estimate the impact of *each* EHT tax rate using quarterly data while controlling for the number of rooms in each unit. Additionally, I use a broader control group, the Greater Vancouver Region (hereafter GVR), to account for the effects of the COVID-19 pandemic as well as the impacts of the SVT and FBT. In contradiction to their results, I find statistical significant decreases in rental prices after the 2020 increase of the EHT to a 1.25% rate.

The results of my empirical approach are as follows. I estimate a significant negative effect on rental prices of -1.1% at the 1.25% tax rate and an effect of -1.8% at the 3% rate, which remains consistent across all bedroom sizes. My model also assigns statistically insignificant coefficients to the 1% EHT rate, implying that vacancy taxes may only be effective at higher rates. This does not rule out economic significance however. Lastly, the model suggests that the extent of rent decreases varies by home size, with Studios, 3-bedroom, and 4+ bedroom units decreasing at rates higher than 2% at the 1.25% tax rate, while 1- and 2-bedroom bedroom unit rents decreasing at less than 2%. This indicates that vacancy taxes do not affect all subgroups of the rental market equally. These findings reinforces previous research on the rental market and Vancouver's EHT.

My paper proceeds as follows. In Section 2, I provide a concise background on related literature. In Section 3, I proceed with a simple theoretical model that uses search and match theory. Then in Section 4, I discuss the data and methodology I use for my empirical analysis. Lastly, in Sections 5 and 6 respectively, I present my results and conclude with how the implications of these findings relate to evaluating affordable housing policy.

2 Related Literature

A common way to investigate housing markets is to use labor market models, particularly because of similarities between the markets in regard to the short run inelasticities of supply, where prices are the adjustment response to demand (Wheaton, 1990; Halket and Pignatti Morano di Custoza, 2015). The Diamond, Mortensen, and Pissarides (hereafter DMP) search and match model is essential to relevant literature including Mortensen and Pissarides (1994), Han and Strange (2015), and Feijoo-Moreira (2020).

2.1 Borrowing Concepts from Labor Economics

Wheaton (1990) is, to my knowledge, the first to explore ownership housing markets through a DMP model, highlighting similarities with the labor market, such as its static number of jobs and workers and the housing market's "stock-flow" attributes, as well as parallels between wage determination mechanisms and housing price bargaining. However, Wheaton also highlights important differences between labor and housing markets, noting that while unemployment during search is common in the labor market, a similar search process in the housing market often results in households occupying two units. This is because the cost of homelessness is too high. He predicts that vacancy rates result from long-term equilibrium, where the marginal cost of new housing stock adjusts to match the marginal cost derived from a home's expected value. Since this paper, research on housing market search-and-match frameworks has expanded significantly, including but not limited to Díaz and Jerez (2013),Genesove and Han (2012), Gabrovski and Ortego-Marti (2018), and Gabrovski and Ortego-Marti (2022).

Naturally, literature has expanded beyond the housing market to also examining the rental market. Desgranges and Wasmer (2000) apply Nash Bargaining, a core component of the DMP model, to examine rental price determination. They analyze both short-term and medium-term equilibrium, with the latter equilibrium endogenizing the entry of new tenant. Their findings suggest that market tension, θ, defined as the ratio of vacancies to seekers, plays a crucial role in setting rental prices. Dong, Shoji, and Teranishi (2020) adapt housing search-and-match frameworks to account for rental frictions and vacancies. Their model treats entry into the market as endogenous, showing that increased entry into the rental market leads to significant marginal changes in rental prices. My model contributes to this literature by examining the impact of a vacancy tax on θ and the subsequent effect on equilibrium rental prices. The wealth of literature in this area of research underscores the importance of understanding how interventions in the housing market can significantly impact rental market outcomes, and vice versa.

2.2 Vacancies and Vacancy Taxes in the Housing Market

Housing market vacancies are the result of search and matching frictions, where market tightness and the probability of finding a match result in a flow share of stock being empty Han et al. (2023) Moreover, these vacancy rates are the product of incidence and duration, which Gabriel and Nothaft (2001) define as the percentage of stock that is vacant and on average how long that stock is vacant for, respectively. They also suggest that rental market vacancy *incidence* is observed to be related to tenants' search costs as well as the pricing of units, whereas vacancy *duration* is related to frictions in the market itself, implying that duration is related to the efficiency of market clearing. Desgranges and Wasmer (2000) corroborate these findings through a model that simulates rental prices as a product of Nash Bargaining between stakeholders in the rental market. They predict that vacancy rates, as a key part of search and match frictions, impact the value functions of which tenants and landlords use to negotiate. Further, they predict that altering vacancy can be a tool to adjust the rental market. My model builds on these findings by analyzing how vacancy taxes can alter vacancies. Vacancies are both a window in which we can look through to diagnose market dynamics and also a tool in which we can alter those market dynamics.

The relationship between rents and vacancies has also been widely studied as well. Rosen and Smith (1982) describe a stock-flow rental market characterized by high short-term supply inelasticity, a key feature of DMP models. Their model shows that changes in vacancy rates are indicative of changes in supply, demand, and market equilibrium prices. Conversely, fluctuations in vacancy rates also influence those factors, namely rental prices. Abramson et al. (2021); Han et al. (2023); Kumar (2024) explicitly model the link between rental prices and vacancy rates, predicting that if vacancy taxes reduce vacancy rates, rental prices adjust accordingly as the market reaches a new equilibrium. This body of literature highlights the importance of vacancy rates, and in extension, vacancy taxes in shaping rental market dynamics.

Speculation is another factor that impacts vacancy rates, where owners may choose to withhold housing stock from the unit for financial or personal gain. Research by Segú (2018) analyzing France's "Taxe sur les Logements Vacants" differentiates between *frictional* vacancies and *speculative* vacancies, which Segu refers to as extra vacancies from owners strategically withholding their units from the market. An example of this may be vacation-home owners, second homes, or contexts in which owners gain more from the appreciation of an empty house than renting it out, like investment properties. Segú (2020) uses a difference-in-differences approach, utilizing municipalities in France that fall under certain population density thresholds as a counterfactual for those where the tax was imposed. Using household-level taxation data, Segú (2020) finds a 13% decrease in vacancies, which increases the housing stock in the rental market. She predicts that theoretically, this would lead to a decrease in rental prices. In another article, Segú (2018), she warns that in the long term, a theoretical increase in the cost of investment in housing units — due to the lower return on rents — could lead to an offsetting effect of lower supply and higher rental prices. More recent studies complement this finding, such as Kumar (2024), whose model observes that decreasing rental vacancies are associated with opposite effects in the housing market. This suggests a need for further research into the long-term impacts of a vacancy tax and whether its net welfare effect is ultimately positive or negative.

Han et al. (2023) builds on Segu's work on the effects of vacancy taxes, defining *speculative* vacancies, which are owned by investors, as different from the traditional *structural* vacancies, which are part of the market clearing process. These speculative vacancies reduce welfare through an artificially inflation of rental prices, causing the rental market to lack efficiency. Han et al. (2023) create a model to show the impacts of a vacancy tax on both types of vacancies, as well as the rental, developer, and owner-occupied housing markets. They theorize that a tax like the EHT could decrease speculative vacancies by increasing the cost of holding a home vacant. This reduction of speculative vacancies creates downwards pressure on rental prices through introducing off-market vacant second homes into the supply of rental homes. Like Segú (2018), Han et al. (2023) also warn of a potential long term offsetting effect. This is a result of property developers, whom are now more willing to supply to the rental market, supplying less to the owned housing market and causing inflationary pressure there.

2.3 Vancouver's Empty Homes Tax

Han et al. (2023) create a model that evaluates the impact of the EHT specifically in Vancouver, considering local laws like exemptions given during transfer of ownership. They then empirically calibrate the model's predictions on the owned housing market using a dataset of housing transactions from 2014-2018, differencing out the impact of other laws like the FBT and SVT through their counterfactual of Burnaby, BC – a city adjacent to Vancouver not impacted by the EHT. Their model predicts that the EHT, and vacancy taxes in general, results in higher affordability in the rental market but an opposing welfare effect in the home-sales market as a result of lower supply. Empirically, Han et al. (2023) use a set of home transaction data find a decrease home sales activity after the implementation of the EHT in 2017, but do not have rental price data to test the impact of the EHT on the rental market. My research complements their findings by focusing on specifically modelling the relationship between the EHT and the rental market, as well as empirically estimating the EHT's effects on rental market prices.

Recent research on the EHT by Caracciolo and Miglino (2024) uses annual average rents reported in Canadian census data with 67 distinct areas in the city of Vancouver and the neighboring city of Burnaby. Using a difference-in-difference framework, they compare Vancouver to the neighboring city of Burnaby as a control, finding a decrease in the vacancy rate of between 1.5% and 1.4%. They conclude that the EHT's implementation in 2017 significantly reduces vacancies in Vancouver and improves housing availability. However, their research is inconclusive in finding a precise impact of the EHT on rental prices. My results corroborate their finding of a negative effect on rents from the implementation of the EHT. I build upon Caracciolo and Miglino (2024) research by using data with a higher frequency of rental prices and examining price changes per unit-size.

3 Theoretical Model

Consider a housing rental market in which u searching tenants are looking at a supply of v vacant units, with matching function $m(u,v) = A\sqrt{\frac{1}{\theta}}$ as per Duffy and Jenkins (2024), where $\theta = \frac{v}{u}$ is market tension. The probability of a tenant finding a suitable unit is $P_u = \frac{m(u,v)}{u} = \frac{\delta}{\delta + A\sqrt{\theta}}$ and vice versa, the probability of a vacant unit matching with a suitable tenant is $P_v = \frac{m(u,v)}{v} = \frac{\delta\theta}{\delta + A\sqrt{\theta}}$. Tenants naturally leave units with a Poisson process at rate δ. This results in steady state:

$$\delta(1-u) = m(u,v) = uP_u = vP_v \tag{1}$$

I assume that $dP_u/d\theta > 0$ and $dP_v/d\theta < 0$ because as θ increases, P_u rises and P_v falls. To avoid contradiction, u must decrease, increasing $\delta(1-u)$. This raises vP_v, which, since P_v drops, implies an increase in v.

3.1 Stakeholders

Landlords, the owners of the housing units, either have their unit occupied by a satisfied tenant, providing a value of R, an unsatisfied tenant, providing a value of S, or by nobody, providing a vacant value of V. Per period, landlords receive r as rent and pay c as the cost of vacancy, which under a vacancy tax increases by t.

Tenants either satisfactorily occupy a unit full time, Y, or are unsatisfactorily occupying a unit, N. Unsatisfactorily occupying a unit means they are searching for a unit that would be satisfactory *while* living in their unsatisfactory unit. I assume tenants are never unhoused as the cost of homelessness is extremely high, as per Wheaton (1990). Tenants receive a flow benefit of y from being housed in a satisfactory unit, flow benefit n from being housing in an unsatisfactory unit, and pay r as rent per period.

3.2 Inter-temporal Functions

Consider the market level Bellman equations of a landlord when their unit is occupied by a satisfied tenant

$$d \cdot R = r - \delta(R - S) \tag{2}$$

when occupied by an unsatisfied tenant

$$d \cdot S = r - P_u(S - V) \tag{3}$$

and when a unit is vacant

$$d \cdot V = -c + \delta(R - V) \tag{4}$$

where d is the discount rate, R is the expected discounted utility of having a satisfied tenant, S is the expected discounted utility of having an unsatisfied tenant, and V is the expected discounted utility of having a vacant unit.

Also consider the Bellman equations of a tenant when occupying a satisfactory unit

$$d \cdot Y = (y - r) - \delta(Y - N) \tag{5}$$

and occupying an unsatisfactory unit

$$d \cdot N = (n - r) + P_u(Y - N) \tag{6}$$

where Y is the expected discounted utility of living in a satisfactory unit and N is the expected discounted utility of living in an unsatisfactory unit.

3.3 Second Homes

I assume that the tax will not affect the flow value of landlords looking for a tenant because the EHT only comes into effect after 6 months. Therefore, landlords only pay vacancy taxes when keeping a unit vacant as a second home or speculative investment property. The flow value of a second home is

$$d \cdot H = b - c - t \tag{7}$$

where b is the flow benefit of the vacant home and H is the expected discounted utility of owning a second home.

The total number of units on the rental market is the number of tenants, 1, plus the number of vacancies, v, which must equal the total number of units, N_H, multiplied by the fraction of unit owners with second-home benefits less than b^*, which is the value where the flow value of keeping a second home vacant and searching for a tenant are equal. At b^*, the value of a second home is equal to Equation (4):

$$b^* - c - t = -c + \delta(R - V) = d \cdot V \tag{8}$$

$$b^* = \delta(R - V) + t = d \cdot V + t \tag{9}$$

Therefore, the fraction of total units in the rental market is $F(d \cdot V + t)$ where F is the cumulative distribution function of flow benefits b among owners. I assume a uniform distribution. Thus, the second home equation is:

$$1 + v = N_H F(d \cdot V + t) \tag{10}$$

Rewriting Equation (10) as

$$\phi = 1 + v - N_H F(d \cdot V + t) = 0 \tag{11}$$

and by the Implicit Function Theorem:

$$d\theta/dt = -\frac{\partial\phi/\partial t}{\partial\phi/\partial\theta} = \frac{N_H f(d \cdot V + t)}{dv/d\theta - N_H \cdot d \cdot f(d \cdot V + t) dV/d\theta} \tag{12}$$

I find that an EHT-style vacancy tax increases $\theta = v/u$ because $dV/d\theta < 0$, shown by Equations (21) and (22) in the appendix.

3.4 Nash Bargaining

Consider the Nash Bargaining process for determining rent in the rental market where I assume equal bargaining power

$$r^* = \underset{r}{argmax}\ (R_c - V)^{1/2}(Y_c - N)^{1/2} \tag{13}$$

where

$$R_c = \frac{r_c - \delta(R-S)}{d} = R + \frac{(r_c - r_m)(d+\delta+P_u)}{(d+\delta)(d+P_u))} \tag{14}$$

and

$$Y_c = \frac{(y - r_c) - \delta(Y-N)}{d} = Y + \frac{r_m - r_c}{d+P_u} \tag{15}$$

which are inside options for tenants and renters in Nash Bargaining for rental price , where R_c is the Bellman equation for a landlord with a satisfied tenant and Y_c is the Bellman equation for a tenant living in a satisfactory unit. r_c represents the individually negotiated rent, while r_m represents the market level rent that results from Nash Bargaining. This relationship between R_c and R is shown in the appendix via Equations (23) to (26).

Finding the first order condition of (13), setting $r_c = r_m = r$, solving for r, and substituting for matching function $m(u,v) = A\sqrt{\frac{1}{\theta}}$ yields:

$$\frac{A(n-y)\left(2A^2\theta + A\cdot d\cdot(\theta+3)\sqrt{\theta} + A\delta(\theta+3)\sqrt{\theta} - \delta^2(\theta-1) + d^2(\theta+1) + 2d\cdot\delta\right)}{2\theta^{3/2}\left(A\sqrt{\theta}+d+\delta\right)^3} \tag{16}$$

Assuming a splitting rate of once every two years ($\frac{1}{24}$), an interest rate of 5%, A of 0.1, negative $n-y$, and θ of 0.5

$$A = 0.1, \quad \delta = \frac{1}{24}, \quad d = 0.05, \quad \theta = 0.5, \quad (n-y) < 0$$

and substituting into (16) predicts:

$$\boxed{\frac{\partial r}{\partial \theta} < 0} \tag{17}$$

This implies that EHT-style vacancy taxes have negative effects on rents with conservative assumptions.

3.5 Theoretical Model Summary

Holding δ to a certain low level, I find that the rental effect of EHT-style vacancy taxes is negative as vacancy taxes increase θ. This model corroborates the predictions for the rental market discussed in Han et al. (2023), who also offer further discussions on how a vacancy tax might impact home-sales and developer markets. These findings imply an economic effect of vacancy taxes on decreasing rental prices, which is empirically explored in the next section.

4 Data and Empirical Strategy

This section describes the data I utilize and my empirical approach to measuring the EHT's impact on the Vancouver rental housing market. I first summarize my data source and the panel that I use. Then, I discuss my methodology and the variables that I use in my regression.

4.1 Data Summary

My main data source is from CoStar Group, one of North America's top providers of commercial real estate information services. The firm's business model is comprised of three main pillars. First, CoStar offers subscriptions to access its commercial real estate database (Costar Group Inc., 2025). Secondly, the firm provides research services for specific client needs. Lastly, CoStar occupies the online housing marketplace, owning subsidiaries such as LoopNet, Homes.com, and Apartments.com.

Using CoStar's commercial real estate database, I build a balanced panel dataset that includes average rental data of housing units inside over 7,000 multifamily property buildings in Vancouver and the GVR, spanning from Q1 2013 to Q3 2024. These properties report their average rental prices by quarter via surveys from CoStar. My data is comprised of a panel of quarterly rental price averages for s-bedroom units in Vancouver and GVR where $s \in \{0, 1, 2, 3, 4+\}$ (CoStar Group Inc., 2013-2024). Importantly, the number of time series is less than the number of years, which changes empirical strategy discussed later. My panel is summarized in Table 1.

To get a sense of how representative the data is, I compared rental prices in CoStar's data to Canadian Census data in 2021 (Government of Canada, 2022). Of 255,055 multi-family residences in Vancouver in 2021, CoStar's data represented around a third of total units. Furthermore, Census data shows an average monthly spend by tenants for shelter to be 1,660$ CDN, but CoStar shows a higher amount, fluctuating around a monthly spend of 1,800$ CDN throughout the four fiscal quarters (Government of Canada, 2022; CoStar Group Inc., 2013-2024). This can be the result of CoStar's position in the multifamily commercial real estate market, where rents are expected to be higher than average given management and amenity costs. Comparisons between the numbers in the Census for the GVR and CoStar's data are similar. Overall, it appears that the data used in this paper is of a large portion of housing stock in the City of Vancouver and the GVR.

Table 1: Summary Statistics, Mean Rent, CDN$

	Bedrooms				
	Studios	1	2	3	4+
2013-GVR	958	1,215	1,534	1,854	1,511
2013-Vancouver	1,184	1,427	2,035	2,688	3,033
2014-GVR	972	1,234	1,557	1,882	1,543
2014-Vancouver	1,206	1,452	2,070	2,733	3,088
2015-GVR	987	1,252	1,580	1,909	1,574
2015-Vancouver	1,225	1,478	2,106	2,777	3,133
2016-GVR	1,021	1,297	1,632	1,973	1,634
2016-Vancouver	1,262	1,532	2,174	2,834	3,218
2017-GVR	1,073	1,366	1,711	2,070	1,725
2017-Vancouver	1,327	1,612	2,286	2,948	3,359
2018-GVR	1,145	1,449	1,816	2,184	1,835
2018-Vancouver	1,387	1,691	2,389	3,088	3,520
2019-GVR	1,206	1,518	1,904	2,291	1,949
2019-Vancouver	1,442	1,760	2,482	3,211	3,693
2020-GVR	1,265	1,587	1,991	2,392	2,012
2020-Vancouver	1,493	1,815	2,553	3,327	3,832
2021-GVR	1,323	1,656	2,076	2,488	2,065
2021-Vancouver	1,523	1,877	2,626	3,447	3,931
2022-GVR	1,427	1,803	2,246	2,634	2,216
2022-Vancouver	1,625	2,038	2,782	3,549	4,047
2023-GVR	1,562	1,952	2,434	2,807	2,320
2023-Vancouver	1,745	2,220	3,014	3,724	4,184
2024-GVR	1,634	2,029	2,512	2,912	2,405
2024-Vancouver	1,831	2,307	3,152	3,892	4,321

4.2 Variable Description

My primary regression, which I name the "Heterogeneous Treatment Effects" model, examines the effect of the EHT on the log of rental prices for each variation of EHT rate in one regression. This associates a change in rental prices to each EHT rate. The model is as follows:

$$\log(Rent_{r,t}) = \beta_0 + \beta_1 UnitSize + \beta_2 Quarter + \beta_3 \log(Rent_{r,t-1}) + \beta_4 EHT + \sum_{k=1}^{3} \gamma_k * T_K + \sum_{k=1}^{3} \alpha_k EHT * T_K \quad (18)$$

Where:

- $Rent_{r,t}$ is the average rent for units with s bedrooms in region r and time t.
- $UnitSize$ is an indicator variable for the size of a housing unit, measured by the number of bedrooms (0,1,2,3,4+)
- $Quarter$ is an indicator variable for the fiscal quarter that an observation is located in (Q1, Q2, Q3, Q4)
- $Rent_{s,r,t-1}$ is a lag of the previous quarter's rent to control for autocorrelation
- EHT, the treatment variable, is an indicator for the region that was affected by the EHT (Vancouver)
- $\sum_{k=1}^{3} \gamma_k * T_K$ is a set of three an indicator variables for the time period that is associated with a certain EHT rate. For example, "Between 2017-2020" indicates the time period in which the EHT was at 1%.

- $EHT * \sum_{k=1}^{3} \gamma_k * T_K$ is a set of interaction variables that considers the impact of the EHT in the period of a certain EHT rate in Vancouver. For example, EHT*Between 2017-2020 refers to the effect on the log of rental prices in Vancouver when the EHT was at 1%.

I utilize another model that complements my primary regression. This second model, which I call "Heterogeneous Treatment Effects, by Bedroom Number", approaches analysis similarly to the primary regression, but for each bedroom size available in the data. This model follows 5 regressions

$$\log(Rent_{s,r,t,}) = \beta_0 + \beta_1 Quarter + \beta_2 \log(Rent_{s,r,t-1}) + \beta_3 EHT + \sum_{k=1}^{3} \gamma_k * T_K + \sum_{k=1}^{3} \alpha_k EHT * T_K \quad (19)$$

where s is Studio, 1-, 2-, 3-, or 4+ bedroom units.

Descriptions of construction and explanations for variable choices that are used in my methodology are as follows. I assume that the change in average rent from a percentage change in the EHT will be multiplicative. Therefore, I decide to take the log of the average rents as my dependent variable. This allows the approximation of the percentage change of rents due to the EHT and provides better insight into the EHT's effects in contrast with examining a specific dollar amount, as net rental price varies inherently between unit size. This variable is constructed by taking the log of average asking rental prices for each observation.

It is important to consider variables that measure the effect of the EHT itself. In my model, they are the set of $\sum_{k=1}^{3} \gamma_k * T_K$ variables, the singular EHT indicator, and set of $EHT * \sum_{k=1}^{3} \gamma_k * T_K$ variables. The $\sum_{k=1}^{3} \gamma_k * T_K$ variable is an indicator of when the EHT changes. Therefore, the first, $Between2017-2020$ is generated as "1" for all values including and between Q1 2017 to Q4 2020, and "0" elsewhere. This is expanded to the other EHT rates and their corresponding time periods. These indicators allow comparison between the period before and after each change, which isolates the rate hikes instead of the net rate, for example an increase in 1.75% in 2021 instead of the total 3% EHT.

The EHT variable controls for the geography in which the EHT was imposed, Vancouver. This allows for the regression to explore differences in places that were affected by the EHT and geographies that were not. EHT is generated as "1" in Vancouver, and "0" in the GVR. This dichotomy allows for direct comparison between the test group (Vancouver) and the control group (GVR).

Lastly, the $EHT * \sum_{k=1}^{3} \gamma_k * T_K$ represent the specific impact of each EHT rate on Vancouver's rental prices. These variables also difference out confounding market factors, like the FBT, SVT, and the COVID-19 pandemic.

Given the inherent premiums of rental prices between units with varying bedroom numbers and the panel itself spanning from 2014-2024, I also control for the heterogeneity of rental prices for different sized units and seasonality, using the set of indicators $UnitSize$ and $Quarter$. This assumes that new units rented during the quarter of the year follow seasonal pricing trends, which influences the average market rent. These sets of singular indicators account for heterogeneity that could otherwise bias my coefficients of interest.

The $UnitSize$ variables are constructed as indicators of how many bedrooms units have. For instance, the indicator 1-*Bedroom* implies that there is one bedroom in the unit. For this paper, studio units are assumed to have 0-bedrooms.

The $Quarter$ variables are constructed as indicators of the fiscal quarter that observations are in. This is important because housing markets are seasonal in nature, where search and match occur more in certain times of the years, consistently. An example of this is that more people look to sell and purchase a new home during the warmer months of the year, as it makes the moving process easier.

4.3 Robustness

Given that my data have $T > N$, using standard time series cross sectional regression analysis may not yield correct standard errors, as per Beck and Katz (1995). Standard OLS panel regression may therefore lead to overconfidence in the significance of certain coefficients. I control for this potential overconfidence by using panel corrected standard errors (PCSEs), which make estimates more robust under panel heteroskedasticity and contemporaneous correlation. Consider the following adjustment to basic OLS standard error calculations:

$$\mathrm{Cov}(\hat{\beta}) = (\mathbf{X}'\mathbf{X})^{-1}\{\mathbf{X}'\boldsymbol{\Omega}\mathbf{X}\}](\mathbf{X}'\mathbf{X})^{-1} \tag{20}$$

Additionally, I correct for autocorrelation in my independent variable of $log(Rent)$. According to Table 4 in the Appendix, I find a significant possibility of an autocorrelation process of lag-1. Intuitively, rents in the previous period can be explanatory of current rent prices, as it highlights the underlying value of living in a specific area. Therefore, I specify an AR(1) process during my regression analysis.

4.4 Building Confidence in the Parallel Trends Assumption

A core assumption in Difference-in-Differences models is parallel trends. In my data, this would mean that before the implementation of the EHT in 2017, Vancouver's average market rents are comparable to that of the GVR's. I run a formal test to examine parallel trends

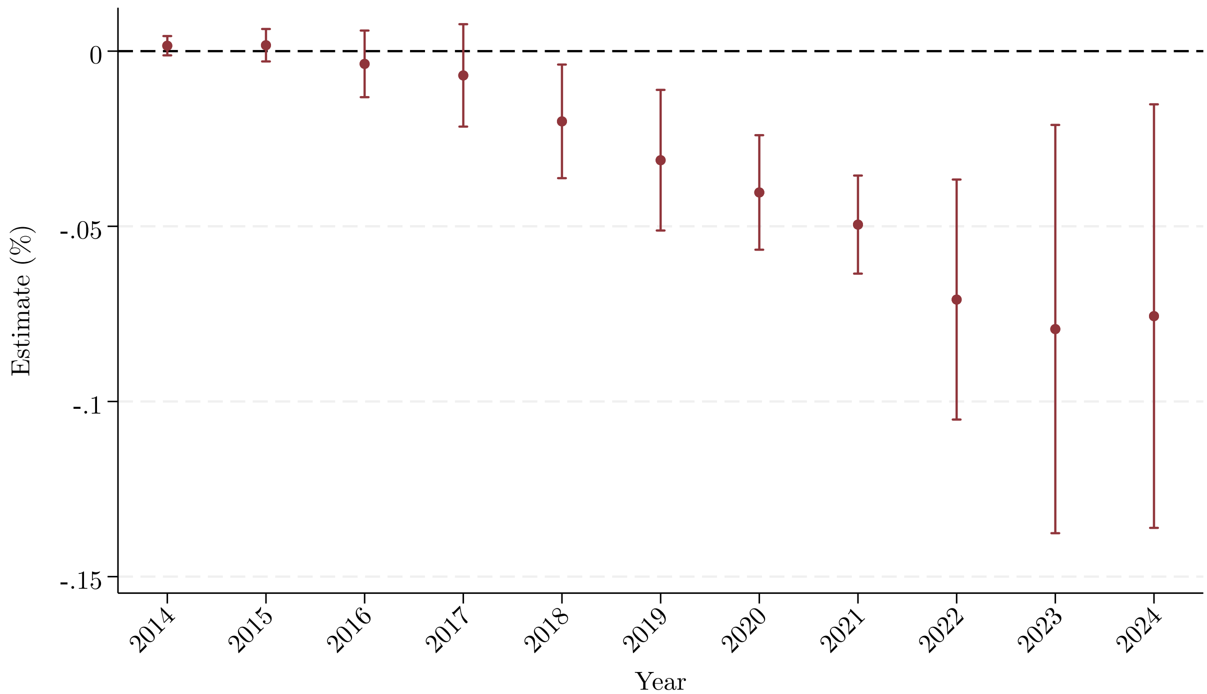

Figure 3: Event Study of the EHT's effects

before the implementation of the EHT using an event study method. Interacting the EHT variable on each year's log of average market rental price shows that Vancouver's rental prices were not significantly different before 2017, as seen in Figure 3. This extends until 2018, but it is important to consider the 6-month period before units are eligible as "vacant", the start of the COVID-19 pandemic, and tax reporting and education delays for the first year of a new tax. A visual inspection of parallel trends can be seen in Figure 4.

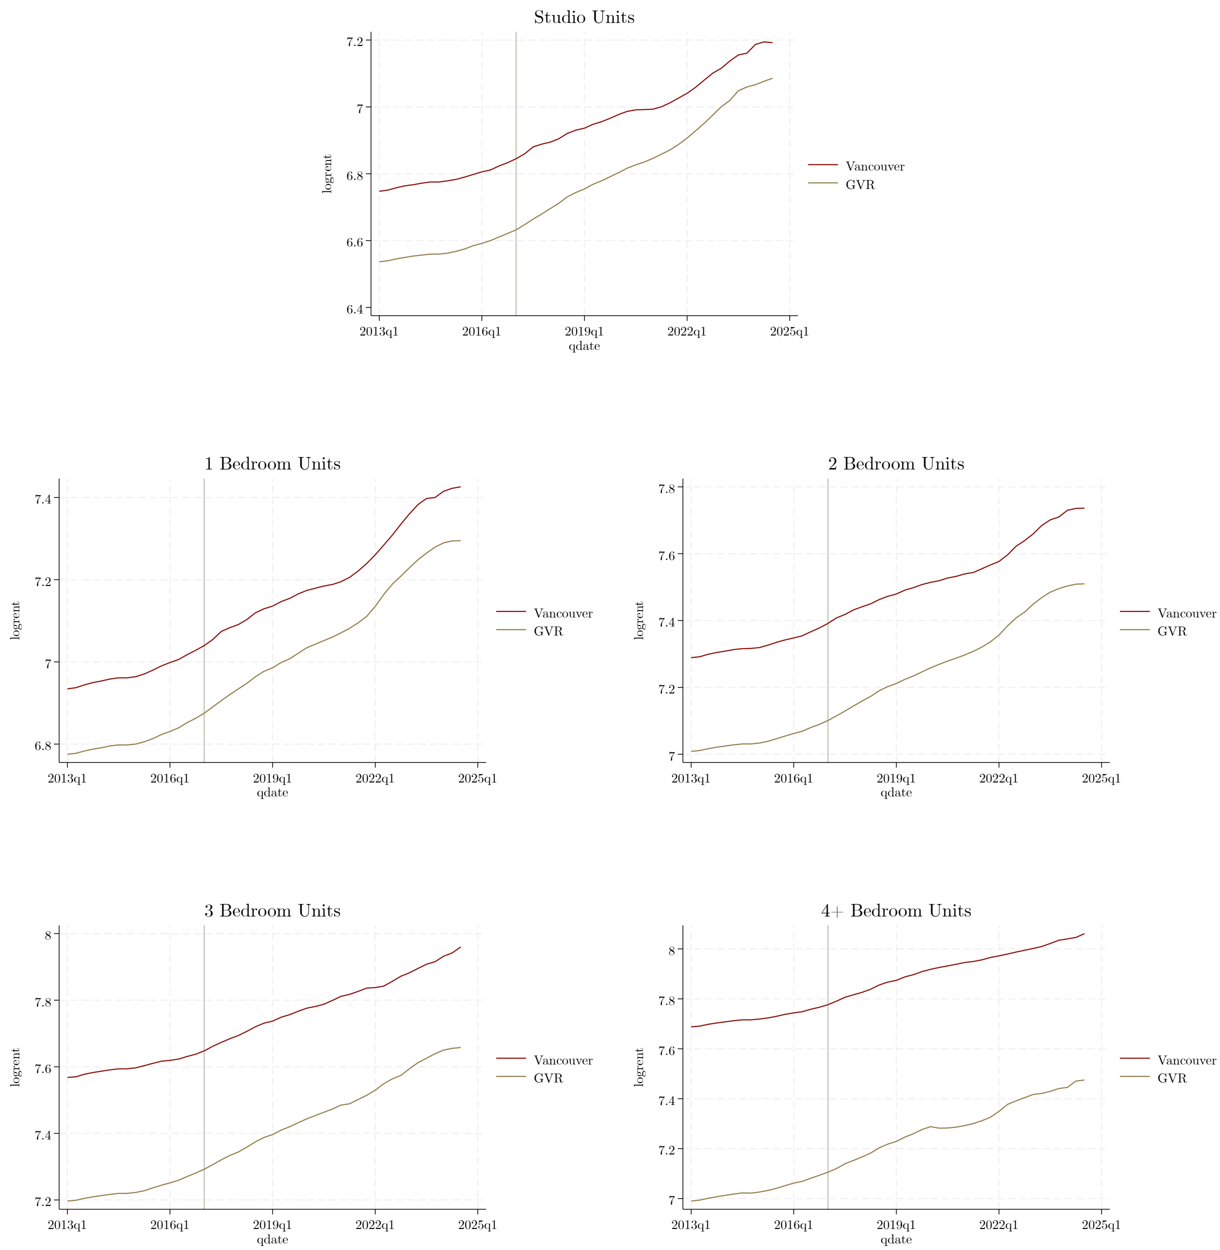

Figure 4: A Visual Inspection of Parallel Trends by Unit Size

5 Results

To understand the effect of the EHT on Vancouver's housing market, I highlight the results of three different model specifications. Examining the model results, I look at the implications for the rental market and how the EHT impacted rental prices.

5.1 Model Outputs

Results for the "Heterogeneous Treatment Effects" are displayed in Table 2. I choose to display the results with and without the AR(1) process. Excluding the AR(1) process associates statistical significant effects of -1.9%, -4%, and -6.8% on rental prices from the 1%, 1.25% and 3% EHT rates respectively. By including the AR(1) process, I find no statistically significant negative effects on change of rent at the 1% EHT rate. At the tax rate of 1% from 2017-2020, the EHT is estimated to decrease average rental price by -0.4% , with a confidence interval that does not fail to reject 0. However, this does not rule out economic significant results. A -0.4% decrease in rents saves the average 1-Bedroom tenant around 110 CDN$ a year, using average rents from CoStar's 2024 Quarter 1 data (CoStar Group Inc., 2013-2024). During the years 2020-2021, when the EHT was at a tax rate of 1.25%, the model estimates -1.1% decrease in Vancouver's average rental price at the 95% confidence level. Lastly, the model associates a significant 1.8% decrease in rental prices after the tax hike to a total of 3% in 2021. This negative coefficient is strongly significant at the 99% confidence interval.

Table 2: Heterogeneous Treatment Effects, With and Without AR(1) Specification

Variables	No AR(1)	AR(1)
log(Rent)	.	.
Quarter 2	0.010	0.003
	(0.018)	(0.004)
Quarter 3	0.021	0.007
	(0.018)	(0.005)
Quarter 4	0.023	0.009*
	(0.019)	(0.005)
1 Bedrooms	0.218***	0.218***
	(0.001)	(0.004)
2 Bedrooms	0.500***	0.498***
	(0.001)	(0.004)
3 Bedrooms	0.718***	0.713***
	(0.003)	(0.008)
4+ Bedrooms	0.689***	0.679***
	(0.004)	(0.010)
Between 2017-2020	0.144***	0.041**
	(0.018)	(0.018)
Between 2020-2021	0.241***	0.084***
	(0.027)	(0.025)
Between 2021-2024	0.375***	0.125***
	(0.017)	(0.031)
EHT	0.345***	0.320***
	(0.002)	(0.007)
Between 2017-2020 * EHT	-0.019***	-0.004
	(0.003)	(0.004)
Between 2020-2021 * EHT	-0.040***	-0.011**
	(0.005)	(0.005)
Between 2021-2024 * EHT	-0.068***	-0.018***
	(0.003)	(0.006)
Constant	6.498***	6.643***
	(0.017)	(0.037)
R-squared	0.921	0.996

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

These coefficients of interest have important implications. The statistically insignificant coefficient during the years 2017-2020 suggest that a 1% tax on vacancy did not have pronounced downwards pressure on market rents. Perhaps the tax rate was not high enough to incentivize second home owners to supply their units to the rental market. Following this, the significant decrease in market rents after the tax increased to 1.25% and 3% may suggest that there exists a range of tax rates that would be effective in incentivizing unit-owners to switch from voluntary vacancy to entering the rental market. Further, by multiplying the -1.8% rental price decrease at the 3% EHT rate by the average Vancouver monthly home rent of 2,301 CDN$ in April of 2025 on CoStar, I predict that the average tenant saves 471 CDN$ per year.

An additional model specification, shown in Table 3 and Figure 5, titled "Heterogeneous Treatment Effects, by Bedroom-Size" examines the EHT's effects on rental prices for units of different bedroom sizes. It is similar to the model used in Table 2, and provides the two following insights on how vacancy taxes can impact different categories in rental markets.

First, this model again provides consistency to previous results, associating negative effects of the EHT at its 1.25% and 3% rates across all bedroom types. This model associates coefficients of varying degrees across bedroom sizes with larger effects on Studios and 4+ Bedroom units. This suggests a higher sensitivity to vacancy taxes at the beginning and end of home size ranges. Secondly, another interesting result of this model is that there is a significant effect at the 1% EHT rate for 3 Bedroom and 4+ Bedroom units. Intuitively, the *absolute* tax value should be larger in homes of larger sizes, and may therefore be more salient for homeowners.

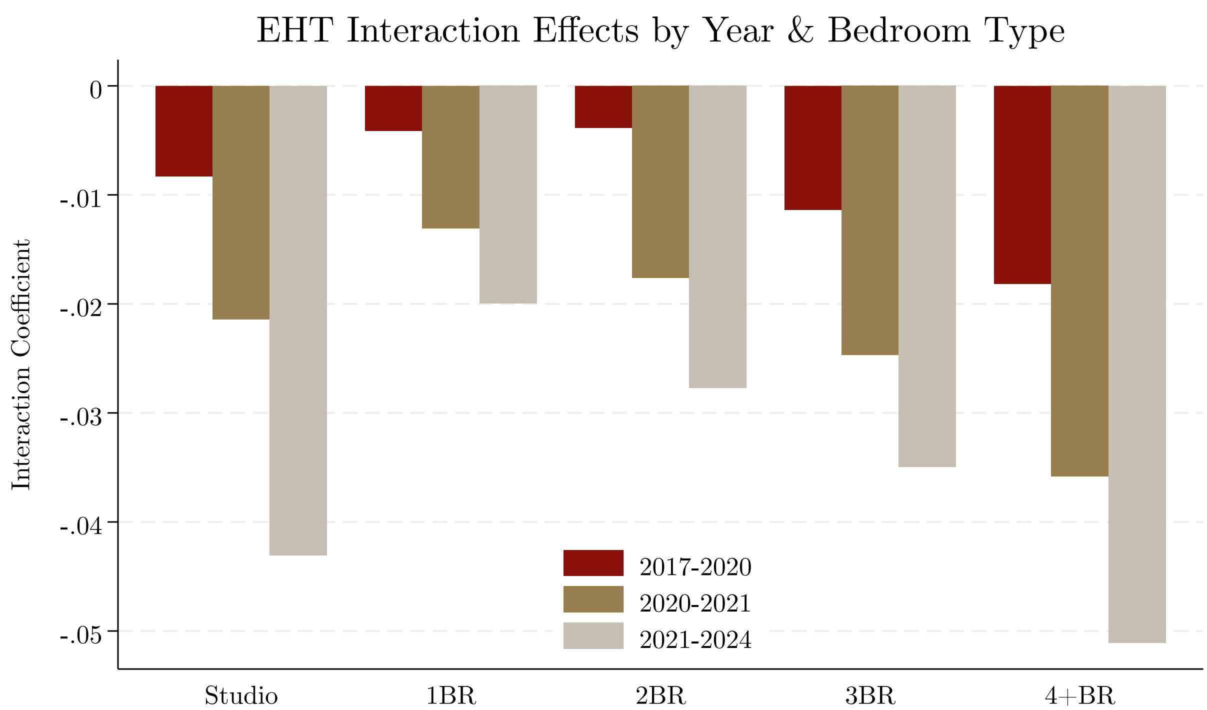

Figure 5: A graphical representation of the EHT's effects by bedroom size

Table 3: Heterogeneous Treatment Effects, by Bedroom Number Results, with AR(1)

Variables	Studio	1 BR	2 BR	3 BR	4+ BR
log(Rent)	.	.	.	.	.
Quarter 2	0.006	0.006	0.006	0.005	0.006
Quarter 3	0.012*	0.012*	0.012*	0.011*	0.013**
Quarter 4	0.015**	0.015**	0.014**	0.015***	0.017***
Between 2017-2020	0.070***	0.072***	0.068***	0.071***	0.090***
Between 2020-2021	0.147***	0.148***	0.140***	0.145***	0.174***
Between 2021-2024	0.222***	0.221***	0.210***	0.217***	0.251***
EHT	0.191***	0.156***	0.270***	0.354***	0.673***
Between 2017-2020 * EHT	-0.008	-0.004	-0.004	-0.011**	-0.018**
Between 2020-2021 * EHT	-0.021**	-0.013***	-0.018***	-0.025***	-0.036***
Between 2021-2024 * EHT	-0.043***	-0.020***	-0.028***	-0.035***	-0.051***
Constant	6.658***	6.889***	7.120***	7.292***	7.085***
R-squared	0.996	0.996	0.997	0.998	0.998

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

5.2 Results summary

These results support the theoretical model laid out in Section 3. Intuitively, because the value of holding homes vacant is decreased via a vacancy tax, there is a higher supply of homes in the active rental market, pushing average rental price downward. Interestingly, the three model specifications add nuance to the theory, suggesting that variations in EHT rates are important in determining the outcome and incentives for owners to list their units on the rental market and how rents in those markets are affected by categorical characteristics.

6 Conclusion

As populations become denser and the demand for housing impedes on the affordability of rental housing around the world, it is pertinent that policymakers make the necessary adjustments to combat these issues. Vacancy taxes have been touted by cities like Vancouver and Oakland to alleviate some of the inflationary pressure in rental markets.

I evaluate Vancouver's implementation case of a vacancy tax, the EHT, over the course of its three iterations in tax rates. I create a model in Section 3 that examines the relationship between a vacancy tax that increases the cost of holding a second-home vacant and the equilibrium rental price in a search and match rental market. To summarize, I find a negative effect on rents with a positive vacancy tax rate that increases θ. This corroborates previous theory and models from Han et al. (2023) that explore the EHT's impact on the rental and home-sales markets. Their theoretical model finds positive effects on housing affordability in the rental market through increased rental stock.

I examine this theoretical relationship between vacancy taxes and rental prices through the use of CoStar Group's commercial real estate data for multifamily units in Sections 4 and 5. I first conduct parallel trend tests to strengthen my assumption that Vancouver and GVR have similar markets before the EHT. I then create a quarterly time series panel of average rental prices for studio, 1-bedroom, 2-bedroom, 3-bedroom, and 4+ bedroom units in Vancouver and the GVR from 2013 to 2024. To measure the EHT's effects on the rental market, I use a difference-in-difference framework to examine the EHT's effects on Vancouver's rental market prices. I create three different model specifications that examine the EHT's three tax rates in combination, separately, and by room number.

This analysis shows a statistically significant decrease in rents of -1.8% for Vancouver at the 3% EHT rate, a statistically significant decrease of -1.1% at the 1.25% EHT rate, and statistically insignificant results at the 1% EHT rate. Thus, it appears that the EHT, and vacancy taxes in general, are most effective at higher levels. I also find that the effects on rental price differ between home sizes. Specifically, the effect is larger in homes with low or high amounts of bedrooms.

To conclude, the results of my analysis imply that vacancy taxes have a deflationary effect on rental prices in the rental housing market at higher rates, around 3%. This suggests that vacancy taxes can be used by policy makers to deflate rental prices. However, though this rent decrease is large by absolute terms, it is small compared to the large problem that is housing affordability. This suggests the need for further intervention at the root causes of housing affordability to complement vacancy taxes.

References

Abramson et al.. 2021. "Search to Rent or Search to Own: Housing Market Churn in the Cross Section of Cities." 10.25740/pd540hs7174.

Beck, Nathaniel, and Jonathan N. Katz. 1995. "What To Do (and Not to Do) with Time-Series Cross-Section Data." *American Political Science Review* 89 (3): 634–647. 10.2307/2082979.

Caracciolo, Gherardo Gennaro, and Enrico Miglino. 2024. "Ripple Effects: The Impact of an Empty-Homes Tax on the Housing Market."

City of Vancouver. 2017. "Housing Vancouver Strategy Presentation."

City of Vancouver. 2018. "Empty Homes Tax Annual Report (2017-2018)."

CoStar Group Inc.. 2013-2024. "Multifamily Commercial Real Estate Data, 2013-2024."

Costar Group Inc.. 2025. "CoStar | # 1 Commercial Real Estate Information Company." `https://www.costar.com`.

Cox, Wendell. 2024. "Demographia International Housing Affordability, 2024 Edition."

Desgranges, and Wasmer. 2000. "Appariements sur le marché du logement." *Annales d'Économie et de Statistique* (58): 253. 10.2307/20076236.

Díaz, Antonia, and Belén Jerez. 2013. "House Prices, Sales, and Time on the Market: A Search-Theoretic Framework." *International Economic Review* 54 (3): 837–872, `https://www.jstor.org/stable/24517067`, accessed on 2025-03-11.

Dong, Mei, Toshiaki Shoji, and Yuki Teranishi. 2020. "Search and Matching in Rental Housing Market." *Working Papers on Central Bank Communication* (015): , `https://ideas.repec.org//p/upd/utmpwp/015.html`, accessed on 2025-03-11.

Duffy, John, and Brian C. Jenkins. 2024. "Search, Unemployment, and the Beveridge Curve: Experimental Evidence." 87 102518. 10.1016/j.labeco.2024.102518.

Feijoo-Moreira, Sergio. 2020. "Search-and-Matching: The Mortensen-Pissarides Model."

Gabriel, Stuart A., and Frank E. Nothaft. 2001. "Rental Housing Markets, the Incidence and Duration of Vacancy, and the Natural Vacancy Rate." *Journal of Urban Economics* 49 (1): 121–149. 10.1006/juec.2000.2187.

Gabrovski, Miroslav, and Victor Ortego-Marti. 2018. "Housing Market Dynamics with Search Frictions."

Gabrovski, Miroslav, and Victor Ortego-Marti. 2022. "Efficiency in the Housing Market with Search Frictions."

Genesove, David, and Lu Han. 2012. "Search and Matching in the Housing Market." *Journal of Urban Economics* 72 (1): 31–45. 10.1016/j.jue.2012.01.002.

Government of Canada, Statistics Canada. 2022. "Profile Table, Census Profile, 2021 Census of Population - Vancouver, City (CY) [Census Subdivision], British Columbia;Metro Vancouver A, Regional District Electoral Area (RDA) [Census Subdivision], British Columbia." `https://www12.statcan.gc.ca/census-recensement/2021/dp-pd/prof/index.cfm?Lang=E`, accessed on 2024-12-10.

Halket, Jonathan, and Matteo Pignatti Morano di Custoza. 2015. "Homeownership and the Scarcity of Rentals." *Journal of Monetary Economics* 76 107–123. 10.1016/j.jmoneco.2015.08.003.

Han et al.. 2023. "Frictional and Speculative Vacancies: The Effects of an Empty Homes Tax."

Han, Lu, and William C. Strange. 2015. "Chapter 13 - The Microstructure of Housing Markets: Search, Bargaining, and Brokerage." In *Handbook of Regional and Urban Economics*, edited by Duranton, Gilles, J. Vernon Henderson, and William C. Strange Volume 5. of Handbook of Regional and Urban Economics 813–886, Elsevier, . 10.1016/B978-0-444-59531-7.00013-2.

Kumar, Nitish. 2024. "A Search and Matching Model of Housing and Rental Market Interactions." *SSRN Electronic Journal*. 10.2139/ssrn.4787212.

Mortensen, D. T., and C. A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *The Review of Economic Studies* 61 (3): 397–415. 10.2307/2297896.

Rosen, Kenneth T., and Lawrence B. Smith. 1982. "The Price Adjustment Process for Rental Housing and the Natural Vacancy Rate." `https://escholarship.org/uc/item/5284v24v`, accessed on 2024-10-19.

Segú, Mariona. 2018. "Taxing Vacant Dwellings: Can Fiscal Policy Reduce Vacancy?." October, `https://mpra.ub.uni-muenchen.de/89686/`, accessed on 2024-11-06.

Segú, Mariona. 2020. "The Impact of Taxing Vacancy on Housing Markets: Evidence from France." *Journal of Public Economics* 185 104079. 10.1016/j.jpubeco.2019.104079.

Statcan. 2024. "Median Total Family Income in British Columbia." https://www.statista.com/statistics/582845/median-total-family-income-british-columbia/, accessed on 2024-12-10.

Thakral, Neil. 2017. "The Design of Public-Housing Policies."

Wheaton, William C. 1990. "Vacancy, Search, and Prices in a Housing Market Matching Model." *Journal of Political Economy* 98 (6): 1270–1292, https://www.jstor.org/stable/2937758, accessed on 2024-11-05.

Appendices

.1 $dV/d\theta < 0$

$$dV/dP_u = -\frac{P_v(y-n)}{d(d+\delta+P_u)^2} < 0 \tag{21}$$

$$dV/dP_v = \frac{y-n}{d(d+\delta+P_u)} > 0 \tag{22}$$

We know that $dP_u/d\theta > 0$ and $dP_v/d\theta < 0$ by assumption so $dV/d\theta < 0$ follows.

.2 R_c and Y_c

$$R_c = R + (R_c - R) = \frac{(r_c - r_m)(d+\delta+P_u)}{(d+\delta)(d+P_u)} - \frac{c\delta P_u - (r_m(d+P_u)(d+P_v)) - \delta r_m(d+P_v)}{(d+\delta)(d+P_u)(d+P_v) - \delta P_u P_v} \tag{23}$$

$$R_c - R = \frac{(r_c - r_m)(d+\delta+P_u)}{(d+\delta)(d+P_u)} \tag{24}$$

$$Y_c = Y + (Y_c - Y) = \frac{d(y-r_m) + \delta(n-r_m) + P_u(y-r_m)}{d(d+\delta+P_u)} + \frac{(r_m - r_c)(d+\delta+P_u)}{(d+\delta)(d+P_u)} \tag{25}$$

$$Y_c - Y_m = \frac{(r_m - r_c)(d+\delta+P_u)}{(d+\delta)(d+P_u)} \tag{26}$$

.3 Woolridge Test

Below is a Wooldridge test for autocorrelation of lag-1 in my panel data. This test associates a significant autocorrelation of lag-1 in my data.

Table 4: Wooldridge Test for Autocorrelation in Panel Data

Null Hypothesis ($\mathbf{H}_0$)	No first-order autocorrelation
F(1, 9)	9069.217
Prob > F	0.0000

What is influencing Intergenerational Income Mobility?
By Xiancheng Huang

Introduction

Intergenerational income mobility—the ability of children to achieve a higher economic status than their parents regardless of their background—reflects the extent to which opportunities are distributed equitably across society. Understanding intergenerational income mobility is not only a measure of economic opportunity, but also a reflection of the broader social structure. Persistent inequalities in mobility hinder the realization of equitable growth and exacerbate disparities in life outcomes.

This paper seeks to address a critical question: What factors most strongly influence intergenerational income mobility across U.S. counties, and how can we identify areas most in need of intervention? Understanding the determinants of mobility is crucial for informing policies aimed at reducing inequality and ensuring that economic opportunity is accessible to all. By leveraging data on income rankings, crime rates, education levels, and unemployment, this research examines the interplay between these predictors and mobility outcomes.

The study proceeds in 3 parts. In the first part I identify a proxy representing the intergenerational income mobility level for counties in the United States and create a heatmap to give a brief visualization about the regional disparities for the inequality. In the second part, I created a model consisting of potential key driving factors through regression. Strongest predictors of intergenerational mobility among hypotheses would be identified. In the final part I discuss the real-world implications for those factors, and certain policies that would greatly boost equality.

Background

How exactly has intergenerational income mobility been measured over the past few decades? Through IPUMS surveys, studies by Aaronson and Mazumder (2008) revealed that mobility steadily increased from 1950 to 1980 during the post World War II boost, where family income moved out quickly from poverty level to mean level. Yet in the 1980s, the mobility fell sharply and failed to revert to its original trend in later years. [1]

On the other hand, Chetty (2014) found that intergenerational mobility remained stable for children in birth cohorts since the 1980s, with an almost linear slope around 0.34 between parental-children income raise. A significant contribution was assumed to be a higher college attendance rate funded by parental income.

Chetty also identified the high geographical and racial disparities features for the intergenerational mobility. Chetty et al. (2019) reveals Black Americans face significantly lower upward mobility compared to their white counterparts, even when starting at similar income levels. Regions like the American Midwest shows higher mobility, whereas the Southeast shows lower.

[1] The exact stat of mobility lacks official data and continue to rely highly on research's measurements.

Erikson and Goldthorpe (2002) emphasize that intergenerational mobility extends beyond economic metrics, reflecting deep-rooted social class structures. By examining mobility through a "class origin to class destination" framework, they argue that societal inequalities are embedded in employment relations and access to stable, high-quality employment. He also stated the pivotal role of education as a mediator in intergenerational mobility, which aligns with the theory of Chetty.

Recent analysis suggests that intergenerational mobility poses impacts beyond income, including various aspects of life. Gene Heyman's 2024 study establishes a strong correlation between low mobility and higher rates of drug overdose deaths in U.S. Midwest counties between the year 2003 and 2022. This finding suggests that limited economic opportunities can lead to adverse social and health consequences, which could continue to involve more aspects of our life if continue developing.

Because of the wide regional variation, it is plausible to assume that certain persistent inequalities in access to resources such as quality education, stable employment, and safe living conditions, have been posing a lasting effect on intergenerational income mobility over time. For instance, counties with higher crime rates often show lower mobility, as unsafe environments can hinder educational attainment and economic progress. Similarly, limited access to higher education and high student-to-teacher ratios exacerbate inequalities, reducing the likelihood of upward mobility for children in disadvantaged areas. This study would like to closely examine these predictors and seek to uncover the most influential ones driving the disparities.

Empirical analysis and Results

Proxy construction and regional overview

The dataset I've been conducting research on is sourced from Chetty's Opportunity Insights. Fixing on county geographical level, we included several socioeconomic factors that's influential according to our hypotheses. It includes information for 3,109 [2]U.S. counties, capturing key variables such as the income ranks of children from families at various parental income levels (P1, P25, P50, P75, P100), crime rates, educational attainment, unemployment, and more. The children, or the sample in the dataset of Chetty all belong to the birth cohorts ranging from 1980 to 1990, and their income rankings were measured around 2014-2015, when they were in age 24-34.

To examine intergenerational mobility, a proxy was calculated as the range of child income ranks between the top (P100) and bottom (P1) [3]parental income levels in each county. This proxy serves as the dependent variable in the analysis, offering a measurable indicator of mobility disparities across regions. A weighted population proxy was also calculated to help us more accurately present the inequality status for each state in the United States. If the value of proxy is high for a certain county, it means that there is low level of intergenerational income mobility.

In general, the value of proxy falls between 0.2 and 0.4 for all 48 states, with Ohio as a high for 0.38 and Utah as a low for 0.21. Figure 1 visualizes the weighted proxy for

[2] Hawaii, Alaska, Washington DC, and several invalid counties' data are not included.

[3] Normally, a more precise calculation would require parental income rank for each county, but our dataset does not include such info. Range methods were used as an alternative.

each state, with each state's exact proxy slightly adjusted. We found that states showing higher levels of inequality were mainly around Southeastern areas, where states in the Midwest were generally on a lower portion of inequality. This result is consistent with Chetty's geographical theory of mobility, which emphasizes the critical role of local factors in shaping economic outcomes. The average proxy for 11 states that has a weighted proxy above 0.35 (marked in maroon) is 0.036 higher than the average for all 48 states, showing significant geographic disparities for mobility.

Southeastern states, often marked by higher poverty rates, lower educational attainment, and weaker social infrastructure, present systemic barriers that hinder intergenerational mobility. These structural challenges create a cycle where children from disadvantaged families have fewer opportunities to improve their economic standing.

In contrast, many Midwestern states benefit from a combination of stronger public services, lower crime rates, and more equitable access to education, which contribute to higher mobility outcomes. This aligns with Chetty's assertion that the 'place effect', or the influence of where a child grows up, is a key determinant to opportunity. The Midwest is also the region showing less racial disparity, where white population is more dominant than other races. Less racial factors could also potentially be the reason for higher regional equality.

Figure 1 *Map of Intergenerational Mobility Proxy for American States*

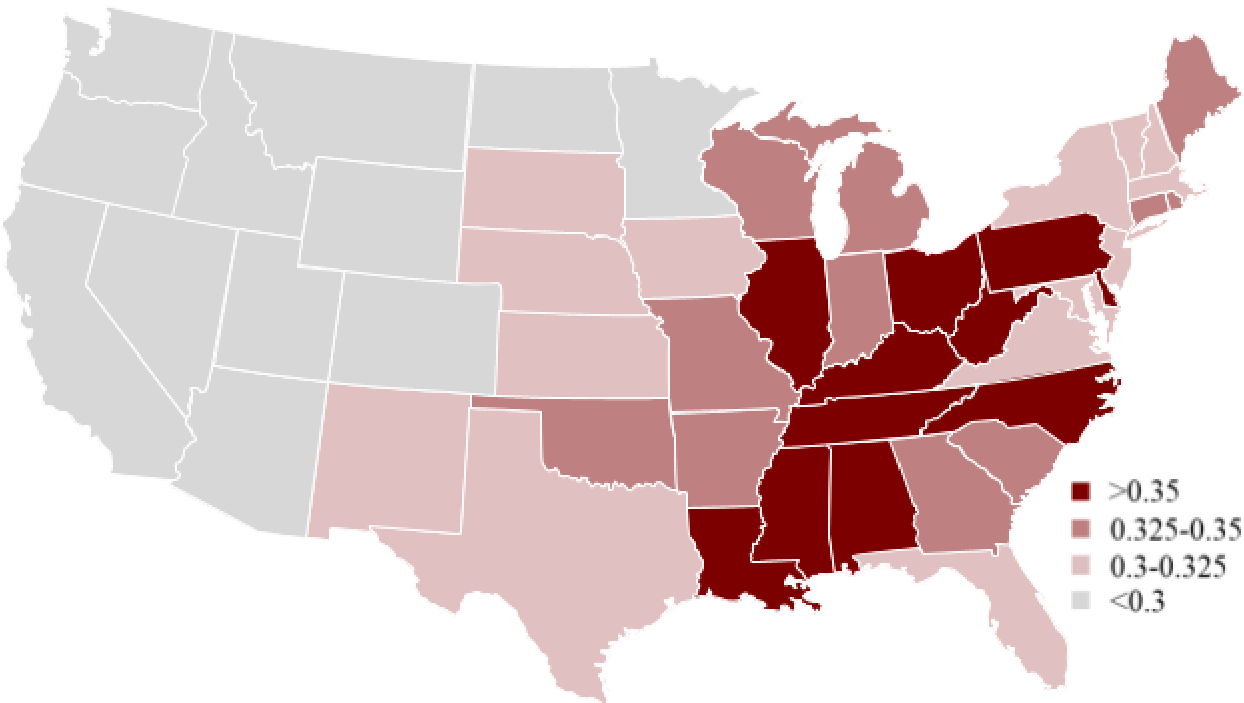

Note: Hawaii, Alaska, and DC are not included
Source: Authors' calculations from OpportunityInsights' data

Correlation Check

Since large varieties of socioeconomic factors were involved in mobility and there's limitation for our study to access everything related at once, a correlation check was made before the main regression model to ensure we retain the key factors with high significance and low bias. Several variables from the initial hypothesis were dropped during the check, and the final selected factors' correlation with the proxy are provided in Table 1 in the Appendix.

It is noticeable that two selected factors presenting different levels of education were slightly more correlated with each other than acceptable level.[4] But given their high correlation with the proxy and the importance of education's contribution to upward mobility, we include both in the regression model.

Regression results

The hypothesized model for regression, based on the background discussion, was constructed as:

Mobility Proxy (County)=$\beta 0+\beta 1$(Employment)+$\beta 2$(Education)+$\beta 3$(Crime Rate)+e

Employment consists of work participation for individuals in the dataset as well as the unemployment level for each county. Education includes several factors together contributing to upward mobility in the study.

The main takeaways for the regression are highlighted in Figure 2. The overall model is a good fit with high F-stats and $R^2 = 0.1716$. Given that the dependent variable range of 0.2 to 0.4 for all States under study, we could retrieve a detailed measurement for each factor. All independent variables in the regression are statistically significant at 95% level, with a proxy mean of 0.324. Full regression can be found in Table 2 in the Appendix.

The county unemployment rate, which is considered a significant regional factor for upward mobility, reflects the young adults' job-seeking environment. Its Beta 0.486 suggests a 10% increase in unemployment would result in an upward proxy shift of 0.05, suggesting a strong positive relationship. Higher unemployment rates correlate with greater inequality.

The labor force participation rate shows a reverse effect on mobility than unemployment. For children from lowest quartile (P25), participating in work greatly increased their upward mobility, posing a strong negative relationship to the proxy. It's fitful to believe stable income earnings provide opportunity for people to rise from their original ranking.

The Education factor is determined by several predictors together in our study. The measure of Pupil-Teacher ratio is a numeric value presenting the average number of students that one instructor teaches. There's also the college degree completion rate as well as some college education rate for children at lowest quartile. Higher number of students per teacher indicates a lower regional quality of education, while higher

[4] Some college education's correlation with Full College education is 0.7, greater than 0.5

college education percentage indicates a greater high-level education participation for children from lower ranking families.

All education predictors have negative coefficients in the regression, reflecting their positive contribution to equality. For full college and some college education, their coefficient were -0.17, -0.11, respectively. Despite the higher correlation between different education levels in the previous check, we could still tell that higher levels of college education contribute more to equality and upward mobility by their difference in values. The Pupil-Teacher ratio suggests that every 10 more students per teacher result in a 0.025[5] increase in proxy, which suggests that aside from degrees, quality education is also crucial in the path of intergenerational equality.

The local crime rate for each county is relatively small in nature, around 0.4% to 1% for most counties. Yet their impact determining the safety of the neighborhood poses great importance to children's quality of living. Beta of 1.36 suggests 3% increase in crime rate, causing the proxy to increase by 0.04, highlighting its significant positive contribution to inequality.

Figure 2 *Betas of socioeconomic factors on mobility proxy regression*

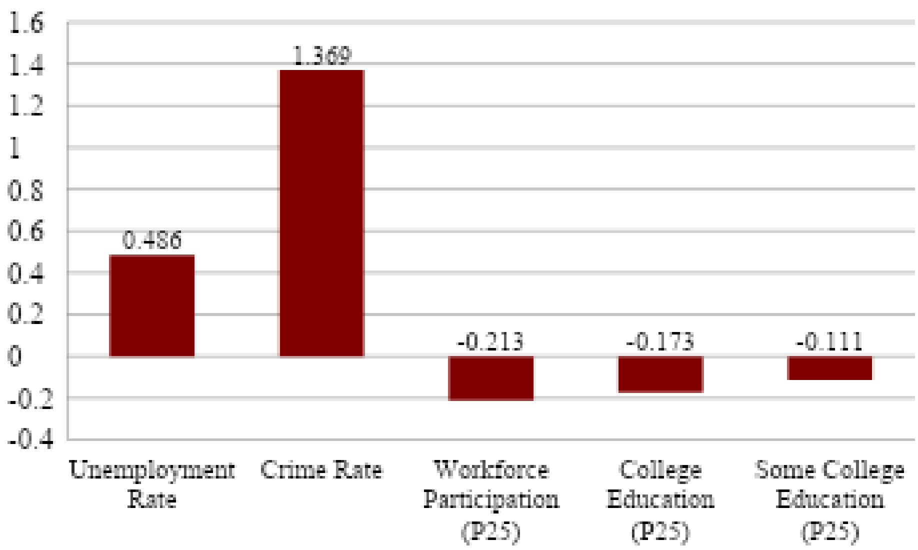

Note: Pupil-Teach Ratio were not included because it changes numerically
Source: Authors' calculations from OpportunityInsights' data

The regression findings highlight actionable opportunities for addressing intergenerational inequality. To mitigate the impact of high unemployment rates on mobility, targeted job creation programs and workforce training initiatives should be prioritized in disadvantaged counties. Additionally, improving the quality of education by reducing pupil-teacher ratios and expanding access to higher education for low-income families is essential. The significant role of crime rates in perpetuating inequality underscores the need for community safety programs, such as crime prevention lectures and neighborhood revitalization. Together, these interventions can address systemic barriers and foster pathways to achieve upward mobility.

[5] Same as unit level per change of color on Figure 1.

Limitations/Potential Improvements

Since the topic of intergenerational income mobility was not structurally explained with official posts, my attempts on at giving a brief discussion of the contributing factors seems immature in many ways. The limited dataset makes it hard to track changes in mobility over time, and to assess the effects of recent policy interventions. We use a proxy to measure inequality, which is only a simplified measure that might not fully capture the complex nature of mobility. Some dataset factors that are likely to influence mobility, such as access to healthcare, early childhood interventions, or family wealth, are not included in the regression. Their exclusion could result in omitted variable bias. Many more border factors including urban-rural differences or racial differences would require even greater effort to apply in the study.

Some further approaches to the topic include but are not limited to:

Gathering more information on the parental side of the sample to bring a more scientific proxy; utilize multi-factor models as well as other statistical tools to create higher adaptive models and graphics; including other demographic factors such as race or gender to reflect inequality in wider ranges. All would be proceed based on the current foundations and results.

Conclusion

Though intergenerational income mobility shows great regional disparities, we managed to retrieve valuable info through regression. The certain socioeconomic factors highlighted in this study addresses the targeted interventions we could make to achieve higher levels of equality for disadvantaged regions.

Education programs seem particularly effective, with better public school teaching resources, a certain region could achieve significant higher upward mobility. Career resources such as workforce training or business incentive programs would also provide more opportunities for people from lower ranking families. Expanding violence prevention initiatives and fostering community-led policing can help create safer environments for economic advancement. Tackling these interconnected issues is essential to ensuring equitable access to opportunity and breaking the cycle of generational inequality.

References

Chetty, Raj, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter. 2019. "Race and Economic Opportunity in the United States: An Intergenerational Perspective." Working Paper 2019-1. Cambridge, MA: Opportunity Insights.

Heyman, Gene M., Ehri Ryu, and Hiram Brownell. 2024. "Evidence That Intergenerational Income Mobility Is the Strongest Predictor of Drug Overdose Deaths in U.S. Midwest Counties." International Journal of Drug Policy 132: 104558.

Heyman, Gene M., Nico McVicar, and Hiram Brownell. 2019. "Evidence That Social-Economic Factors Play an Important Role in Drug Overdose Deaths." International Journal of Drug Policy 74: 274–84.

Raj Chetty, Nathaniel Hendren, Patrick Kline, Emmanuel Saez, "Where is the land of Opportunity? The Geography of Intergenerational Mobility in the United States", The Quarterly Journal of Economics, Volume 129, Issue 4, November 2014, Pages 1553–1623

Aaronson, Daniel, and Mazumder, Bhashkar (2008). "Intergenerational Economic Mobility in the United States, 1940 to 2000." Journal of Human Resources, 43(1), 139–172

Erikson, Robert, and John H. Goldthorpe. 2002. "Intergenerational Inequality: A Sociological Perspective." Journal of Economic Perspectives, 16 (3): 31–44.

I'd like to thank Professor Gene Heyman and BC Psycho department in the help of sorting the dataset, as well as the inspiration / explanation to the topic.

Appendix

Table 1 *Correlation check for the Socioeconomic factors*

Socioeconomic factors	Correlation with Proxy
Unemployment	0.2313
Work Participation (P25)	-0.2708
Full College Education (P25)	-0.3514
Some College Education (P25)	-0.3497
Pupil-Teach Ratio	0.1781
Local Crime Rate	0.1683
Single Mother Households	0.3403*

Note: *Single Mother Households isn't studied in this research, but its correlation is noticeable
Source: Authors' calculations from OpportunityInsights' data

Table 2 *Regression Results for the Model of Mobility Proxy*

Socioeconomic factors	Coefficient (Standard error)
Unemployment	0.486***
	(0.0799)
Work Participation (P25)	0.213***
	(0.0378)
Full College Education (P25)	-0.173***
	(0.0241)
Some College Education (P25)	-0.111***
	(0.0188)
Pupil-Teach Ratio	-0.00244***
	(0.00057)
Local Crime Rate	1.369***
	(0.383)
Constant	0.580***
Observations	3109
R-squared	0.1856
Proxy mean	0.324

Note: ***Indicates significance at the 1-percent level and ** at the 5-percent level.
Source: Authors' calculations from OpportunityInsights' data

Building Novel Regressive Models to Debias Recidivism Predictions

Cameron Craig
Boston College
ECON3339
December 6th, 2024

I. Introduction

Pretrial risk assessment instruments (PRAIs) are predictive algorithms used to determine criminal law decisions such as bail and sentencing. Initially, these algorithms were regarded as a potential eliminator of bias in criminal decisions due to the removal of a human component. However, these tools have been criticized for implicit racial biases, maintaining the same issue as human decisions throughout history. COMPAS, a widely used PRAI, functions by assigning individual criminal defendants a risk level based on a 137-question survey. This survey data is combined with a predictive model trained on historical criminal data, demographic information, educational and occupational history, and many other data points. One of the major criticisms of this algorithm is the lack of transparency in the model; COMPAS does not share the proprietary modeling process that the algorithm uses to determine risk assessments. COMPAS's proprietary model is considered a trade secret by Northpointe, the current owner of the software. The result is an algorithm that achieves around 65% accuracy in predictions with no clear method to understand said predictions. These results of validation tests are publicly available from Northpointe.

The literature has taken issue with PRAIs such as COMPAS over the implicit racial biases seemingly present in COMPAS's predictions. ProPublica, an investigative journalism organization, has conducted analyses on COMPAS's performance that suggest that COMPAS's algorithm has a Type 1 error rate of 45% for black defendants. In this paper, I aim to create a novel predictive algorithm with the goal of minimizing this Type 1 error rate when predicting recidivism while maintaining the same accuracy as COMPAS. Further, I intend to create said algorithm with fewer inputs. One of the major criticisms in the literature is the excessive number of input variables that construct COMPAS and similar PRAIs. Various sources have proposed that similar accuracy can be achieved with minimal inputs. In the construction of my algorithm, I have used a smaller number of inputs directly related to predicted risk level. In my model, I intend to apply two debiasing techniques. These techniques include residualizing inputs such as predicted risk after regressing said variables on an indicator for race and using several interacted inputs to capture an amplified effect from key explanatory variables. In assessing these models, the predictive accuracy is determined based on how frequently the model correctly predicts an individual's risk of recidivism at greater than 50%. The level of bias is determined and measured by the frequency of Type 1 and Type 2 errors for Black and White defendants. Following my analysis, residualizing variables as a debiasing technique had no effect on accuracy or the level of bias. However, using models with fewer inputs as well as including interacted variables both improved model accuracy and the level of bias present in the data. Based on these conclusions, I have determined that it is both possible and important to construct predictive models that rely on fewer inputs. Additionally, I propose the simple technique of including interaction terms as a method through which to reduce bias and improve accuracy.

II. Scope of Research

My research includes a literature review synthesizing the existing criticisms and analyses of predictive algorithms used by the justice system, as well as the analysis and results of novel predictive models built with publicly available COMPAS data. Through the literature review, I have identified two areas for potential improvement and expansion. These include the continued development of clear predictive models that stray from black-box algorithms and further exploration of debiasing techniques in said predictive models. In my creation of a novel predictive model, I have utilized maximum likelihood estimation to predict risk level with ordered multinomial logistic regression models, as well as logistic regression models to predict the likelihood of recidivism for each individual. My research is limited in that I am only exploring the estimation of the likelihood of recidivism. While I have conducted my own risk assessment for the individuals in the data, this paper explores minimizing Type 1 errors in predicting recidivism rather than debiasing comprehensive risk assessments. Based on the accuracy levels determined by ProPublica's critique of COMPAS, I intend to predict the likelihood of recidivism with an accuracy higher than 61%. Further, I intend to improve upon the racial bias identified in the ProPublica article by decreasing the percentage of Type 1 errors for black defendants. According to ProPublica, COMPAS's Type 1 error rate for Black defendants is 45%. I also intend to apply two different debiasing techniques. These include residualizing key input variables as well as the novel approach of interacting related input variables. It is important to note that a large portion of criminal statistics are unavailable to the general public, and thus my research is limited in scope due to the lack of access to statistical data from individual jurisdictions. Because of this, I will be conducting my analysis using the available data from Broward County, Florida, utilized in ProPublica's critique.

III. Review of Relevant Law and Economics Literature

1. Introduction and Overview

In conducting a preliminary review of existing literature, there are several key aspects that stand out regarding COMPAS and other similar PRAIs. These include criticisms over racial biases and transparency, as well as potential areas for algorithmic improvement. The literature also addresses the broader applications in criminal law, which highlights the existing biases present in human judicial decision-making and indicates the potential benefits of improving upon existing PRAIs. The general opinion on algorithmic risk predictions is fairly consistent; the existing literature identifies and emphasizes the presence of bias and the need for further improvement. However, authors vary regarding their optimism about the potential of removing biases and the broader applications of the PRAIs in the criminal justice system.

2. Critiques of Existing Algorithms (COMPAS as a Case Study)

ProPublica's analysis of COMPAS revealed foundational issues in algorithmic risk assessments, including racial bias and moderate predictive accuracy (61% for general recidivism). Black defendants were significantly more likely to be misclassified as high risk, raising concerns about fairness in judicial outcomes (Larson et al. 2016). Vaccaro (2019) corroborated these findings, noting that while COMPAS did outperform human decision-making in accuracy, it failed to improve fairness. Her dissertation revealed that COMPAS's predictions anchored human judgments, exacerbating biases. This stems from the historical data used in COMPAS's analysis, which includes years of racially biased decisions (Vaccaro 7, 22). Similarly, Hill (2021) critiqued PRAIs, such as COMPAS, for upholding dangerousness predictions based on biased historical data, which disproportionately harm Black and Latino individuals (Hill 54, 63).

3. Transparency and Accountability

A consistent theme across the existing literature is the opacity of proprietary algorithms. Tashea (2017) and Washington (2019) both emphasized the challenges this poses for due process and accountability. In the example of COMPAS, the back-end algorithm is not available to the public; the widely available component of the algorithm is the questionnaire provided to criminal defendants (Tashea 56-57) (Washington 7, 36). A recent court case, *State v. Loomis*, highlighted the inability of defendants to contest COMPAS predictions due to its proprietary nature.

Defendants lack the means to object to the algorithm's methodology due to the lack of clarity (Washington 4-5). This raises significant legal and ethical concerns and decreases public trust in the judicial system. Washington argued for "data science reasoning" to interrogate the inner workings of such algorithms, advocating for procedural fairness through transparency. Lyn (2020) extended this critique, highlighting how opaque algorithms obscure the normative judgments crucial to their design, which can serve to perpetuate existing inequalities and biases (Lyn 18-19).

4. Technical Solutions to Algorithmic Bias

Arnold et al. (2024) provided a more technical perspective. In their two subsequent articles, they develop tools to measure and mitigate disparate impacts in algorithms. They proposed several input adjustments, such as residualizing individual inputs. This process involves regressing each input variable on race while controlling for the true outcome of interest. The race component is then subtracted from the input to create the residualized version; this process would prevent inputs from exhibiting conditional disparities related to race (Arnold et al. 6). Arnold et al. describe this process as regressing "each algorithmic input on race while controlling for pretrial misconduct potential in this sample. By subtracting the race component of these regressions from each algorithmic input, we create pre-processed inputs that have no conditional input disparities by construction. We can then use these pre-processed inputs to build a non-discriminatory algorithm using, for example, a linear regression of true misconduct potential on the pre-processed inputs" (Arnold et al. 1-2). These pre-processing adjustments then "residualize each algorithmic input on race without controlling for true misconduct potential" (Arnold et al. 2). These proposed methods aim to reduce discrimination without sacrificing accuracy and attempt to illustrate how algorithmic fairness can be achieved through careful redesign. Engel et al. (2023) similarly explored algorithmic corrections for COMPAS, finding that anti-Black and anti-youth biases could be mitigated by utilizing machine learning methods with existing data (Engel et al. 9). However, these corrections occasionally introduced trade-offs, such as increased false negatives. Engel et al. thus argue that such normative decisions should not be embedded within the algorithm and should instead be transparent and subject to judicial insight (Engel et al. 9-10).

5. Policy Implications and Broader Social Concerns

Bagaric & Wolf (2018) explored the potential for computerized risk assessments to enhance consistency and reduce judicial biases. They noted, however, that algorithms must be carefully designed to maintain consistent sentencing decisions (Bagaric & Wolf 33-35). Their work tied into broader discussions on the balance between efficiency and equity in criminal justice tools. Lyn (2020) echoed this concern, emphasizing the importance of aligning algorithmic tools with broader social goals, such as reducing incarceration rates and addressing systemic racism (Lyn 17). Based upon these articles, the social costs of utilizing algorithmic predictions bear a similar weight to the existing criticisms of the criminal justice system. With the goals of reducing systemic racism in mind, incorporating judicial opinions into existing algorithmic predictions seems to be a necessary step for improving long-standing biases.

6. Alternative Models Emphasizing Transparency

Recent studies have introduced alternative approaches to recidivism prediction that emphasize transparency, simplicity, and fairness. Rudin et al. (2020) critique the proprietary nature of COMPAS, emphasizing the risks associated with both opacity and lack of accountability. Their findings reveal inconsistencies in COMPAS's documented methodology and actual implementation, particularly its nonlinear dependency on age (Rudin et al. 4-6). These inconsistencies, coupled with the algorithm's reliance on historical data embedded with systemic biases, reinforce the transparency concerns highlighted by Tashea (2017) and Washington (2019). Rudin et al. argue that interpretable models not only enhance procedural fairness but also enable the public to scrutinize and challenge predictions (Rudin et al. 29). This aligns with the broader call for transparent PRAIs by Lyn (2020) and the critiques of biased outcomes by ProPublica and Hill (2021). Similarly, Dressel and Farid (2018) demonstrate that COMPAS does not achieve better predictive accuracy than non-expert humans or simpler statistical models using only two variables: age and prior convictions (Dressel & Farid 3). Their findings support the argument made by Rudin et al. (2020) and Vaccaro (2019) that complex black-box algorithms are not inherently more accurate than interpretable models. Furthermore, Dressel and Farid underscore the racial disparities identified by ProPublica, showing that Black defendants are disproportionately labeled as high risk (Dressel & Farid 3). These results align with those of Engel et al. (2024), who also highlight trade-offs in addressing algorithmic biases, such as increased false negatives when reducing anti-Black or anti-youth biases.

Investigating further into the modifications that can be made to the algorithmic models themselves, Zeng et al. (2015) propose Supersparse Linear Integer Models (SLIM). These models generate transparent and interpretable scoring systems for recidivism prediction. SLIM models perform as accurately as black-box algorithms like COMPAS but are fully interpretable, allowing for manual computation of risk scores (Zeng et al. 691). This approach resonates with the work of Arnold et al. (2024), who advocate for input adjustments to reduce disparities in PRAIs. Zeng et al.'s findings also align with the policy implications raised by Bagaric et al. (2018) and Lyn (2020), emphasizing the importance of designing models that are fair, understandable, and legally defensible.

7. Synthesis and Consensus in the Literature

Taken together, these studies underscore a growing consensus within the literature: transparency and accountability are crucial for addressing biases in PRAIs while maintaining public trust. By demonstrating that simpler, open-access models can achieve comparable accuracy to proprietary algorithms, Rudin et al., Dressel & Farid, and Zeng et al. challenge the prevailing reliance on opaque tools like COMPAS. These findings further support the need for judicial systems to adopt PRAIs that prioritize fairness, accountability, and societal equity.

8. Gaps and Contribution

There are several gaps in the literature when considering debiasing techniques. Arnold et al. (2024) have proposed and tested debiasing techniques through residualizing inputs, however, additional techniques would further improve the public's trust in PRAIs. In my own research, I aim to both create clear predictive models and explore additional debiasing techniques to reduce the rate of Type 1 errors for Black defendants.

9. Conclusion

Following this review of the existing literature, utilizing predictive models that are both clear and understandable is a major improvement to the use of PRAIs. It is arguable that black-box algorithms protected by trade secrets do encourage innovation and design, however, I believe that the loss of trust in judicial systems and the lack of accountability resulting from black-box algorithms outweigh the detriments to innovation and design. It is important to maintain clarity in a field as significant as criminal justice; incarceration and the freedom of low-risk defendants are too important to sacrifice.

IV. Methodology:

To conduct my analysis, I first collected relevant data on criminal defendants, prior criminal history, and recidivism. I utilized the publicly available ProPublica COMPAS data. This data includes the COMPAS screening profiles of 7,214 defendants from Broward County, Florida ranging from January 1, 2013 to December 31, 2014. The dataset includes information on race, age, sex, prior offenses, juvenile offenses, charge degrees, the COMPAS assigned risk-level, whether or not the defendant was arrested within two years of the screening, and various other data points. Offenses include various charges, such as battery and possession, each classified as a misdemeanor or a felony.

Second, I identified key variables in the dataset for predicting recidivism. The key dependent variable for this prediction is an indicator for whether or not a defendant was arrested within two years (*two_year_recid*). This variable is inherently limited; it is a proxy variable for convictions. It identifies recidivism through arrests, which does not guarantee whether the individual is truly guilty of committing a crime. Ideally, the dependent variable would identify recidivism through convictions. The key explanatory variables present in the dataset that I have chosen to focus on include the number of prior offenses (*priors_count*), the number of juvenile felonies a defendant has (*juv_fel_count*), the number of juvenile misdemeanors a defendant has (*juv_misd_count*), and the number of miscellaneous juvenile offenses a defendant has (*juv_other_count*). These variables were chosen due to their indications of criminal history and attempt to capture individual propensity towards criminal activity. For the sake of my analysis, I created several variables using information provided in the dataset. To control for Black-defendant Bias, an indicator for whether a defendant is Black (*black*) was used in each regression. To control for gender effects, an indicator for whether a defendant is female was used in each regression. To control for propensity towards severe criminal activity, a categorical variable for the charge degree of past offenses (*charge*) was included in each regression, equal to 1 for a misdemeanor and 2 for a felony.

To extend my analysis, I mirrored COMPAS's predictions of recidivism risk. The rationale behind including an additional predictor of risk was to gain explanatory power on top of existing predictors. To estimate this risk level, I created a variable titled *risk*, which is equal to one for low risk, two for moderate risk, and three for high risk. I then ran two separate models to obtain predicted risk levels. First, I ran an ordered logistic regression on *risk* using the variables *black, age, charge, juv_misd_count, juv_other_count, female, priors_resid,* and *juv_fel_resid*. These variables were selected to maintain predictive elements similar to those of the larger models. Second, I ran a multinomial logistic regression on *risk* using the same variables. While the ordered logistic model fits the ordinal nature of the *risk* variable, I included the multinomial logit due to its focus on the probability of each category rather than the cumulative probabilities.

Third, I applied the debiasing technique proposed by Arnold et al. (2024) to the data. Following the creation of these variables, I ran several simple regressions to capture the residuals of several variables. These models attempt to capture bias by regressing different variables on the race indicator. The output is then captured as residuals to eliminate the portion of the variable explained by whether an individual is Black in an attempt to eliminate the component of racial bias. The residuals of an OLS regression of *priors_count* (*priors_resid*) on *black* were used to eliminate bias towards Black defendants in prior convictions. The residuals of an OLS regression of *juv_fel_count* (*juv_fel_resid*) on *black* were used to eliminate bias towards Black defendants regarding the consideration of severe convictions in youth. I then applied the same technique to the predicted risk variables mentioned in the previous paragraph. In order to remove the racial bias of these predicted risks, I ran two additional OLS regressions of each predicted risk variable individually regressed on *black.* These residuals were captured as *risk_resid2* and *risk_resid3* for ordered logistic and multinomial logistic respectively.

Finally, I introduced various interaction terms as a method to further debias predictions. These interaction terms aim to weight effects differently by combining the effects of several important variables. To add more weight to the impact of an individual being higher risk and having more prior convictions, I interacted *risk_resid* with *priors_count* (*risk_priors*). I also interacted *priors_count* with *predicted_risk* (*risk_priors2*) for a comprehensive analysis that includes both residualized and non-residualized interactions. To add more weight to the impact of an individual being higher risk and having a propensity towards severe crime (more felonies), I interacted *predicted_risk** with *charge* (*risk_charge*). Using the same reasoning as the previous interaction term, I also interacted *risk_resid* with *charge* (*risk_charge2*). To add more weight to the impact of an individual having a high number of prior felony convictions relative to total convictions, I interacted *priors_count* with *charge* (*priors_charge*). To add more weight to individuals with conviction records that start during youth, I interacted several variables with *priors_count*. These include *priors_juv_fel*, an interaction of *priors_count* and *juv_fel_count*; *priors_juv_misd*, an interaction of *priors_count* and *juv_misd_count*; and *priors_juv_other*, an interaction of *priors_count* and *juv_other_count*. The summary statistics for each variable are listed below in Table 1:

Table 1: Summary Statistics

Summary Statistics					
Variable	Observations	Mean	Std. Dev	Min	Max
two_year_recid	7214	0.4506515	0.4975933	0	1
risk	7214	1.654283	0.7843774	1	3
black	7214	0.5123371	0.4998824	0	1
female	7214	0.193374	0.394971	0	1
age	7214	34.81799	11.88892	18	96
priors_count	7214	3.472415	4.882538	0	38
charge	7214	1.646798	0.477998	1	2
juv_fel_count	7214	0.0672304	0.4739715	0	20
juv_misd_count	7214	0.0909343	0.4852385	0	13
juv_other_count	7214	0.1093707	0.5015858	0	17
risk_resid	7214	5.32E-09	0.7538953	-0.8655303	1.567652
priors_resid	7214	-1.55E-08	4.780982	-4.438853	33.56115
juv_fel_resid	7214	-1.02E-09	0.4727718	-0.1001082	19.96731
predicted_risk1	7214	1.433047	0.7052084	1	3
predicted_risk2	7214	1.420571	0.7167934	1	3
risk_resid2	7214	7.1E-09	0.6511367	-0.6972402	1.844514
risk_resid3	7214	0.00000001	0.6685335	-0.6599026	1.839397
priors_charge	7214	6.053507	9.097401	0	76
risk_priors	7214	1.370238	5.417954	-25.96591	51.73252
priors_juv_fel	7214	0.6465207	7.371719	0	480
priors_juv_misd	7214	0.8929859	7.893294	0	336
priors_juv_other	7214	0.6334904	4.58989	0	204
risk_charge	7214	2.40768	1.553341	1	6
risk_priors2	7214	7.31924	13.72158	0	114
risk_charge2	7214	0.0488424	1.324675	-1.731061	3.135304

Based on the summary statistics, we can draw several conclusions about the data. On average, 45.1% of the individuals in our dataset were arrested within two years of their COMPAS screening. On average, COMPAS assigned individuals with a recidivism risk level of 1.654, placing the average individual slightly above "low" risk. On average, 51.23% of the individuals in Broward County screened by COMPAS were Black. On average, only 19.34% of these individuals were women. The average individual screened by COMPAS was 35 years old. On average, individuals screened by COMPAS had 3.472 prior convictions. On average, the charge degree for these convictions was 1.647, indicating that felony charges were more common than misdemeanor charges. Based on the three variables for juvenile offenses, juvenile offenses were relatively uncommon on average. However, based on the maximum numbers of each of these variables, we can see that individuals in Broward County were convicted of as many as 20 juvenile felonies, 13 juvenile misdemeanors, and 17 juvenile miscellaneous charges.

Following the creation of these variables, I ran eight different models. Models one and two consisted of logistic regressions without predicted risk levels, models three and four consisted of logistic regressions with the ordered logistic predicted risk, and models five through eight consisted of logistic regressions with the multinomial logistic predicted risk. The models using multinomial logistic predicted risk had the most success with debiasing techniques; I suggest that this is due to the emphasis on the probability of each category rather than cumulative probabilities. The relationship between each risk category is arbitrary, and thus unique probabilities for each category improve predictions.

Models one and two focus on predicting recidivism without a variable dedicated to risk level. Model one consisted of a logistic regression with *two_year_recid* as the dependent variable and the following independent variables: *black, age, priors_count, charge, juv_fel_count, juv_misd_count, juv_other_count,* and *female*. I selected these variables based on both their indications of criminal history and their demographic indicators. This model attempted to capture a similar prediction of recidivism to COMPAS without using a risk assessment. Model two introduced the residualized variables into the regression on *two_year_recid*; the independent variables are as follows: *black, age, charge, priors_resid, juv_fel_resid, juv_other_count, juv_misd_count,* and *female.* These variables were chosen to again identify indicators of past criminal history and demographics. The goal of using residualized variables was to eliminate the components of prior convictions and juvenile felonies that were correlated with race. This model, however, performed identically to the first model.

Models three and four introduce a variable dedicated to risk level. In each of these models, I included the ordered logistic predicted risk (*predicted_risk1)* in the regression. The goal of this was to increase explanatory power while avoiding COMPAS's risk predictions. Model three consisted of a logistic regression on *two_year_recid* with the following independent variables: *black, age, predicted_risk, charge, priors_count, juv_fel_count, juv_other_count, juv_misd_count,* and *female*. While this model improved upon the Type 1 error frequency for black individuals, this model saw marginally lower predictive accuracy. Model 4 reintroduced the residualized variables into the regression; the independent variables are as follows: *black, age, risk_resid2, charge, priors_resid, juv_fel_resid, juv_other_count, juv_misd_count,* and *female*. These residualized variables were included to again eliminate components of racial bias; however, this model performed identically to model three.

Models five through eight shift to the multinomial logistic predicted risk. In each of these models, I included the multinomial logistic predicted risk (*predicted_risk1*) in the regression. Model five consisted of a logistic regression on *two_year_resid* without residualized variables, with the following independent variables: *black, age, predicted_risk2, charge, priors_count, juv_fel_count, juv_other_count, juv_misd_count,* and *female*. These variables were chosen to identify factors related to criminal history, as well as key demographic information. This model further reduced the Type 1 Error frequency for Black individuals and improved upon predictive accuracy when compared to models three and four. This model, however, had lower predictive accuracy than models one and two. Model six reintroduced the residualized variables to the regression, including the residuals of *predicted_risk2,* titled *risk_resid3*. This model consisted of a logistic regression on *two_year_resid* with the following independent variables: *black, age, risk_resid3, charge, priors_resid, juv_fel_resid, juv_other_count, juv_misd_count,* and *female.*

The goal was to reduce the Type 1 error frequency for Black defendants by including residualized variables. Model six, however, performed identically to model five.

Models seven and eight move away from residualized variables and introduce interacted terms for a new approach to debiasing the data. Model seven consisted of a logistic regression on the following variables: *black, age, predicted_risk2, charge, priors_count, risk_priors, priors_charge, juv_fel_count, juv_misd_count, juv_other_count, priors_juv_fel, priors_juv_misd, priors_juv_other,* and *female.* These variables were selected to account for the criminal history of individuals, demographic-related effects, as well as the weighted effects from combining related variables. This model had higher accuracy than the previous models, albeit still marginally lower than the two initial models. The introduction of interaction terms proved to be relatively successful in removing bias, as evidenced by the lowest Type 1 error frequency for Black defendants. Finally, model eight consisted of the same logistic regression with the addition of two additional variables, *risk_charge2* and *risk_priors2*. These variables were introduced to further account for varied weighting effects. This model was the most successful predictor of recidivism and had the highest accuracy rate. This model, however, had a marginally higher frequency of Type 1 errors for Black defendants when compared with the first models. This suggests a small tradeoff between decreasing racial biases and maintaining better predictive accuracy.

Finally, I created metrics to assess the validity and accuracy of each model. I took the predicted probabilities from each logistic regression model and created a variable titled *recidivism_predicted* using these probabilities. I then created a variable titled *predicted_class*, which was equal to 1 if the predicted probability of recidivism was greater than 0.5. To test the accuracy, I created a variable titled *correct_prediction*, which was equal to 1 if *predicted_class* was equal to the actual *two_year_recid* variable. Thus, the mean of *correct_prediction* provides the predictive accuracy of each model. To test for false positives, I created a variable titled *false_positive*, which was equal to 1 if *predicted_class* = 1 and *two_year_recid* = 0. I then created two variables for Black and White defendants, titled *black_false_positive* and *white_false_positive*. Both of these variables were equal to 1 if *false_positive* = 1 and *black* = 1 or 0, respectively. I then took the mean of both of these variables if *black* = 1 or 0, respectively, for each race, providing the percentage of false positives for Black and White defendants. To test for false negatives, I created a variable titled *false_negative*, which was equal to 1 if *predicted_class* = 0 and *two_year_recid* = 1. I then created variables for each race, titled *black_false_negative* and *white_false_negative*. These variables were set equal to 1 if *false_negative* = 1 and *black* = 1 or 0, respectively, for each race. I then took the mean of both of these variables if *black* = 1 or 0, respectively, for each race, providing the percentage of false negatives for Black and White defendants.

V. Results

The results of models one through eight show relatively consistent accuracy percentages when predicting recidivism with slight variations in Type 1 error frequency for Black individuals. Of these models, model seven achieved the highest accuracy while simultaneously minimizing the Type 1 error frequency for Black individuals. This suggests that including interaction terms is a valid debiasing technique and should be investigated further. Based on the results of these models, residualizing variables proved to be ineffective in changing predictive accuracy or bias. This technique also requires further investigation to identify potential benefits of utilization. In this paper, I will not be addressing the average marginal effects that provide a more clear idea of each variable's specific impact. This paper will instead focus on accuracy and frequency of Type 1 and Type 2 errors. The coefficients and accuracy metrics of models one and two are depicted below:

Table 2: Models without Predicted Risk Level

	(1) Logit	(2) Logit Residualized
Variable	two_year_recid	two_year_recid
black	0.1026112*	0.4027002***
	(0.0527462)	(0.0519776)
age	-0.0437812***	-0.437812***
	(0.0025344)	(0.0025344)
priors_count	0.14591***	
	(0.0071112)	
charge	0.1734687***	0.1734687***
	(0.0543482)	(0.0543482)
juv_fel_count	0.1620938**	
	(0.0824003)	
juv_misd_count	0.0201056	0.0201056
	(0.0722242)	(0.0722242)
juv_other_count	0.1903947***	0.1903947***
	(0.0627147)	(0.0627147)
female	-0.2860603***	-0.2860603***
	(0.0657569)	(0.0665931)
priors_resid		0.14591***
		(0.0071112)
juv_fel_resid		0.1620938**
		(0.0824003)
Constant	0.4957455***	0.8595565***
	(0.1329246)	(0.1367929)
Observations	7,214	7,214
Pseudo R2	0.1097	0.1097
Prediction Accuracy	0.67771	0.67771
Type 1 Error Frequency, Black	0.1515152	0.1515152
Type 1 Error Frequency, White	0.0727686	0.0727686
Type 2 Error Frequency, Black	0.1704545	0.1704545
Type 2 Error Frequency, White	0.2498579	0.2498579

Standard Errors in Parentheses
***p<0.01, **p<0.05, *p<0.1

Models one and two performed with decent accuracy, providing validity to predictive models of recidivism that rely on fewer input variables. Model one consisted of a logistic regression of *two_year_recid* on the variables *black, age, priors_count, charge, juv_fel_count, juv_misd_count, juv_other_count,* and *female*. Regarding model one, we can make several conclusions based on the coefficients of each variable. The coefficient on *black* is relatively high in magnitude, suggesting that an individual being Black is a significant predictor of recidivism. This coefficient is significant at the 10% level. The coefficients on *priors_count*, *charge, juv_fel_count,* and *juv_other_count*, are all higher in magnitude than the coefficient on *black*, suggesting that previous criminal history is a stronger predictor of recidivism when compared with an individual being black. The coefficients on *priors_count, charge,* and *juv_other_count* are all significant at the 1% level, while *juv_fel_count* is significant at the 5% level. The coefficients on *age* and *female* are both negative, suggesting that both women and older individuals have a lower likelihood of recidivism on average. These coefficients are both significant at the 1% level. Model one predicted recidivism with an accuracy of 67.77%, which is slightly higher than COMPAS's accuracy. This model saw a Type 1 error frequency of 15.15% for Black defendants and 7.28% for White defendants. Further, this model saw a Type 2 error frequency of 17.05% for Black defendants and 24.99% for White defendants.

Model two consisted of a logistic regression of *two_year_recid* on the variables *black, age, charge, juv_misd_count, juv_other_count, female, priors_resid,* and *juv_fel_resid*. Regarding model two, the coefficient on *black* has increased in magnitude to 0.4027. This coefficient is significant at the 1% level. This suggests that by residualizing variables correlated with *black*, the coefficient on *black* has captured additional race-specific effects. This is the only coefficient that changed between the two models; everything else has remained the same. Further, each of the accuracy and bias metrics has remained identical. This suggests that residualizing input variables does not have an effect on accuracy or bias present in predictions. Residualizing input variables, however, has improved the statistical significance of the coefficient on *black*. Following models one and two, I introduced the variable for predicted risk into the regressions. Models three and four are depicted below:

Table 3: Ordered Logistic Predicted Risk

	(1)	(2)	(3)
	Ordered Logit	Logit	Logit Residualized
Variable	risk	two_year_recid	two_year_recid
black	1.05274***	0.0105562	0.2009608***
	(0.0521399)	(0.0558035)	(0.0562367)
age	-0.0804279***	-0.0372979***	-0.0372979***
	(0.0029718)	(0.0027723)	(0.0027723)
priors_count		0.1137004***	0.1137004***
		(0.0093288)	(0.0093288)
charge	0.3191774***	0.1241721**	0.1241721**
	(0.0546202)	(0.0551546)	(0.0551546)
juv_fel_count		0.0566145	0.0566145
		(0.0791185)	(0.0791185)
juv_misd_count	0.3370161***	-0.0651342	-0.0651342
	(0.0747188)	(0.068804)	(0.068804)
juv_other_count	0.2833191***	0.1008085	0.1008085
	(0.0556466)	(0.0632914)	(0.068804)
female	0.1090229*	-0.2966002***	-0.2966002***
	(0.0640179)	(0.0659715)	(0.0659715)
predicted_risk		0.3514594***	
		(0.0671477)	
risk_resid2			0.3514594***
			(0.0671477)
priors_resid	0.21156***		
	(0.0067527)		
juv_fel_resid	0.4829419***		
	(0.0859461)		
Constant		0.0408512	0.4469577**
		(0.1584727)	(0.1329743)
Cut1	-1.452694		
	(0.1427992)		
Cut2	0.3042793		
	(0.1425722)		
Observations	7214	7214	7214
Pseudo R2	0.2021	0.1124	0.1124
Prediction Accuracy		0.6774328	0.6774328
Type 1 Error Frequency, Black		0.1382576	0.1382576
Type 1 Error Frequency, White		0.0508812	0.0508812
Type 2 Error Frequency, Black		0.1845238	0.1845238
Type 2 Error Frequency, White		0.2714611	0.2714611

Standard Errors in Parentheses
***p<0.01, **p<0.05, *p<0.1

Prior to estimating models three and four, I first ran an ordered logistic regression on *risk* in order to create a predicted risk variable. This process consisted of an ordered logistic regression of *risk* on *black, age, charge, juv_misd_count, juv_other_count, female, priors_resid,* and *juv_fel_resid*. This regression used the residualized versions of *priors_count* and *juv_fel_count* to improve statistical significance as seen in the previous models. In this regression, the coefficient on *black* had the highest magnitude, suggesting that black had the strongest effect on risk level. This coefficient was significant at the 1% level. The coefficients on *charge, juv_misd_count, juv_other_count, priors_resid,* and *juv_fel_resid* were all relatively close in magnitude; these coefficients were each significant at the 1% level. The coefficient on age is negative, suggesting that being older reduces the likelihood of being deemed higher risk on average. This coefficient is significant at the 1% level. The coefficient on *female* is positive in this regression, suggesting that being a woman increases the likelihood of being deemed high risk. This coefficient is significant at the 10% level.

Model three features a logistic regression of *two_year_recid* on *black, age, priors_count*, *charge, juv_fel_count, juv_misd_count, juv_other_count, female,* and *predicted_risk*. These variables were selected to maintain consistency between risk prediction and recidivism prediction. These variables also intend to capture the effects of criminal history, severity, and key demographics. Based on the coefficients in this model, *predicted_risk* had the largest impact on the individual's likelihood of recidivism. This coefficient is significant at the 1% level. The coefficients on *priors_count, charge,* and *juv_other_count* are all close in magnitude, demonstrating similar effects on the individual's likelihood of recidivism. The coefficient on *black* is relatively small, suggesting low race-specific effects in this regression model. However, this coefficient is not statistically significant. The coefficients on *age, female,* and *juv_misd_count* are all negative, suggesting that these inputs actually lower the individual's likelihood of recidivism. The coefficients on *age* and *female* are both significant at the 1% level, while the coefficient on *juv_misd_count* is not significant. The negative effect from *age* and *female* makes intuitive sense, however, one would expect *juv_misd_count* to have a positive effect. This negative effect could be explained by a tendency towards minor crimes, suggesting that individuals with higher juvenile misdemeanors engage in petty activity. However, the presence of statistical selection bias is also possible, suggesting that individuals with higher juvenile misdemeanors are different from others in ways not measured in the data. Regarding accuracy, model three performed marginally worse than models one and two. The prediction accuracy of this model was 67.74%, a slight decrease from the previous 67.77% accuracy. Model three was less biased towards Black defendants, with a Type 1 error frequency of 13.83% for Black individuals and 5.1% for White individuals. This model did see an increase in Type 2 error frequencies for both White and Black individuals. The frequency for Black individuals increased to 18.45% while the frequency for White individuals increased to 27.15%. This suggests that there is a tradeoff between false positives and negatives where minimizing one may increase the other.

Model four is again very similar to model three, suggesting a low impact of using residualized variables. The main differences between these two models are the coefficient on *black*, the constant, and statistical significance. First, the coefficient on *black* has increased from 0.011 to 0.2, suggesting a much larger increase in the likelihood of recidivism for Black individuals. This coefficient is statistically significant at the 1% level in this model. The constant has similarly increased from 0.041 to 0.45, suggesting a major increase in baseline odds of recidivism. The constant in this model is statistically significant at the 5% level. These effects further suggest that residualizing variables does not change predictive accuracy, but it improves statistical significance and allows the coefficient on *black* to better capture the effect of racial biases. The predictive accuracy and Type 1 and Type 2 error frequencies have not changed between models. Next, I move to the models using *predicted_risk* from the multinomial logistic regression models, as well as the incorporation of interaction terms. Models five through eight are shown below:

Table 4: Multinomial Logistic Predicted Risk

			(5)	(6)	(7)	(8)
	Multinomial Logit Outcome 2	Multinomial Logit Outcome 3	Logit	Logit Residualized	Logit Interactions 1	Logit Interactions 2
Variable	risk	risk	two_year_recid	two_year_recid	two_year_recid	two_year_recid
black	0.9776356***	1.561278***	0.0356991	0.1979754***	0.0514114	0.1222566**
	(0.0642756)	(0.0818339)	(0.0544804)	(0.0561684)	(0.0552961)	(0.0567919)
age	-0.0706337***	-0.122806***	-0.038343***	-0.038343***	-0.0365927***	-0.0344824***
	(0.0033887)	(0.0052403)	(0.0027089)	(0.0027089)	(0.0027206)	(0.0028437)
priors_count			0.1136815***	0.1136815***	0.1553414***	0.267219***
			(0.0093811)	(0.0093811)	(0.0280895)	(0.0291786)
charge	0.2514948***	0.4106036***	0.1383071**	0.1383071**	0.217442***	0.1348851**
	(0.0656893)	(0.0841797)	(0.0547788)	(0.0547788)	(0.0675689)	(0.068279)
juv_fel_count			0.0541712	0.0541712	0.1760533	-0.0249826
			(0.0786527)	(0.0786527)	(0.1301876)	(0.1508993)
juv_misd_count	0.3935733***	0.6944069***	-0.0937187	-0.0937187	0.0031578	-0.1900222
	(0.1388686)	(0.1377414)	(0.0686673)	(0.0686673)	(0.1150448)	(0.1239303)
juv_other_count			0.1754258***	0.1754258***	0.1362036	0.0125931
			(0.0624547)	(0.0624547)	(0.0860895)	(0.0876633)
female	0.2909511***	0.0209856	-0.2977091***	-0.2977091***	-0.2996253***	-0.3196723***
	(0.0762195)	(0.1015454)	(0.0658518)	(0.0658518)	(0.0661948)	(0.0675267)
predicted_risk			0.3250076***		0.3139466***	0.5747147***
			(0.0635356)		(0.0691242)	(0.080162)
risk_resid3				0.3250076***		
				(0.0635356)		
priors_resid	0.2483561***	0.3456665***				
	(0.0104485)	(0.0118123)				
juv_fel_resid	0.2136514	0.613725***				
	(0.1562167)	(0.1508379)				
risk_charge						0.2396732***
						(0.0291786)
risk_priors					0.0251173***	-0.0711191***
					(0.0073225)	(0.0088159)
priors_charge					-0.0294111*	-0.0137633
					(0.0152107)	(0.013697)
priors_juv_fel					-0.0128829	0.0053577
					(0.0079646)	(0.0128786)
priors_juv_misd					-0.0076333	0.0093199
					(0.007487)	(0.0093982)
priors_juv_other					-0.0097846	0.0051734
					(0.0109243)	(0.0114514)
Constant	0.8981777***	0.1339122***	0.0770968	0.4543015***	-0.0954912	-0.3921525**
	(0.1667648)	(0.2247921)	(0.155797)	(0.1329674)	(0.1728598)	(0.0675267)
Observations	7214	7214	7214	7214	7214	7214
Pseudo R2	0.2093	0.2093	0.1123	0.1123	0.1148	0.1326
Prediction Accuracy			0.6756307	0.6756307	0.6775714	0.6869975
Type 1 Error Frequency, Black			0.1214827	0.1214827	0.1195887	0.1520563
Type 1 Error Frequency, White			0.054008	0.054008	0.0525867	0.0818647
Type 2 Error Frequency, Black			0.2080628	0.2080628	0.2080628	0.1650433
Type 2 Error Frequency, White			0.2649233	0.2649233	0.2643547	0.2268334

Standard Errors in Parentheses
***p<0.01, **p<0.05, *p<0.1

Prior to estimating models five through eight, I first ran a multinomial logistic regression on *risk* to create a predicted risk variable, similar to the process for models three and four. This consisted of regressing *risk* on *black, age, charge, juv_misd_count, female, priors_resid,* and *juv_fel_resid*. These variables were chosen to maintain consistency in inputs and to improve statistical significance through the inclusion of residualized variables. In this regression, I set the base outcome as *risk* = 1. This means that the baseline outcome for individuals is low risk, and the coefficients in outcomes two and three suggest independent impacts of different variables. Based on the coefficients on the variables for each outcome, the coefficient on *black* again has the largest impact on the individual's risk level. This coefficient is significant at the 1% level for each outcome. The coefficients on *charge, juv_misd_count, priors_resid,* and *juv_fel_resid* are all relatively close in magnitude and increase by around 200% from outcome 2 to outcome 3. The coefficients on *charge, juv_misd_count,* and *priors_resid* are all significant at the 1% level. However, the coefficient on *juv_fel_resid* is not significant. Similar to the previous regressions we have seen, the coefficient on *age* is again negative, suggesting a decrease in the likelihood of recidivism as an individual gets older. This coefficient is significant at the 1% level for both outcomes. Similar to the ordered logistic regression on *risk*, the coefficient on *female* is once again positive. This coefficient is significant at the 1% level in outcome 2 and is not significant in outcome 3. This coefficient suggests that if an individual is a woman, she has increased odds of being assigned a higher risk level.

Model five consists of a logistic regression of *two_year_recid* on *black, age, priors_count, charge, juv_fel_count, juv_misd_count, juv_other_count, female,* and *predicted_risk*. The coefficients suggest that *predicted_risk* has the highest impact on the likelihood of recidivism, followed by *juv_other_count, charge,* and *priors_count*. The coefficients on *predicted_risk, juv_other_count,* and *priors_count* are all significant at the 1% level. The coefficient on *charge*, however, is significant at the 5% level. The coefficients on *age* and *female* are both negative, suggesting diminishing likelihood of recidivism as the individual gets older as well as lower likelihood of recidivism for women. These coefficients are both significant at the 1% level. Surprisingly, the coefficient on *juv_misd_count* is negative. I suggest that higher juvenile misdemeanors suggest a propensity towards minor crimes, which could explain the decrease in likelihood of recidivism. However, this coefficient is not statistically significant. This model saw an accuracy of 67.56%, marginally lower than the previous two models. This model did, however, see a decrease in Type 1 error frequencies. The Type 1 error frequency is 12.15% for Black defendants and 5.4% for White defendants. Model five did see an increase in Type 2 Error frequency for Black defendants, shifting from 18.45% to 20.8%. The same frequency decreased for White defendants, however, shifting from 27.15% to 26.49%. Following the trend in each of the models, model six was identical to model five, aside from the coefficient on *black* and the constant in the model. The coefficient on *black* increased from 0.036 to 0.198; the coefficient also gained statistical significance at the 1% level in model six. The constant increased from 0.077 to 0.454, suggesting an increase in the baseline odds of recidivism. The constant also gained statistical significance at the 1% level.

Finally, models seven and eight introduce the interaction terms into the regression model. Model seven consisted of a logistic regression of *two_year_recid* on *black, age, priors_count, charge, juv_fel_count, juv_misd_count, juv_other_count, female, predicted_risk, risk_priors, priors_charge, priors_juv_fel, priors_juv_misd,* and *priors_juv_other*. The intuition behind including these interaction terms was to capture increased weights of past criminal history by interacting related variables. The results of this regression show *predicted_risk* as having the largest impact on the likelihood of recidivism. This coefficient is significant at the 1% level. The effects and magnitudes of the non-transformed variables are relatively similar to past models.

The coefficients on the various interaction terms reveal interesting information about the joint weights of the variables. For example, the coefficients on *priors_charge, priors_juv_fel, priors_juv_misd,* and *priors_juv_other* are all negative. This suggests that for individuals with a high number of prior offenses and juvenile criminal history, the likelihood of recidivism is lower. This could stem from possible legal and behavioral interventions successfully lowering the risk level of previous offenders. It is important to note, however, that this could indicate the possibility of overfitting in the model. Regarding the performance of the model, model five performed the best when considering both accuracy and Type 1 error frequency. Model five predicted recidivism correctly 67.76% of the time, an increase from all models except one and two. Type 1 error frequency in model five was 11.96% for Black defendants and 5.26% for White defendants. Type 2 error frequency was 20.81% for Black defendants and 26.44% for White defendants. These performance metrics suggest that model five was the most successful in minimizing the bias against Black defendants while maintaining similar accuracy to prior models.

Model eight introduced an additional interaction term, *risk_charge*. This variable aims to capture an increased effect from risk by interacting it with the charge degree. In the case of this variable, the risk level would double for each defendant who had been convicted of a felony in the past. The coefficient on this variable is 0.2397; this coefficient is significant at the 1% level. Including this variable has also increased the coefficient on *predicted_risk* from 0.314 to 0.575; this coefficient is also significant at the 1% level. In general, this model saw an increase in magnitude for several variables when compared to model seven. For example, the coefficient on *black* has increased from 0.051 to 0.122. This suggests that the inclusion of *risk_charge* has isolated some of the racial bias onto the indicator for Black defendants. Of each of the models featured in this paper, model eight had the highest pseudo R-squared, suggesting that this model best explained the variation in the data. Model eight further had the highest accuracy of all of the models at 68.7%. This model, however, saw an increase in the frequency of Type 1 errors. The frequency was 15.21% for Black defendants and 8.19% for White defendants. On the other hand, this model minimized the frequency of Type 2 errors with 16.5% for Black defendants and 22.68% for White defendants. Based on these metrics, there appears to be a tradeoff between accuracy and fairness. When minimizing the frequency of Type 1 errors for Black defendants, there was a consistent decrease in predictive accuracy. However, when maximizing predictive accuracy, there was an increase in the frequency of Type 1 errors for Black defendants.

VI. Discussion

Based on my results, it seems to be the case that black-box algorithms not only pose issues with accountability and due process, but they may also suffer when compared to more basic, clear models. In each of the three models reported in my results, I was able to improve upon COMPAS's accuracy percentage; ProPublica determines that COMPAS predicts recidivism with an accuracy of 61%, while each of my models maintained at least 67% accuracy. In the preliminary models observed in this paper, however, the number of inputs compared with COMPAS's model is minimal; these models also boast a much lower frequency of Type 1 and Type 2 errors (i.e. COMPAS performs at ~45% type 1 error for black defendants, highest Type 1 error in the preliminary models is 15.15% for Black defendants). This suggests that decreased inputs has a significant impact in reducing biases–this implies that the "trade secret" component of PRAIs may be limiting the applicability of computerized assessment tools.

Unfortunately, the use of residualized variables as a debiasing technique had no effect on the accuracy or bias of any of my models. This technique did, however, improve the statistical significance of coefficients and appeared to localize the effects of racial bias on the indicator for Black defendants. Based on the work of Arnold et al. (2024), this technique does have real viability and should be further explored and applied. However, in these models, residualizing variables did not prove to be successful as a debiasing tool. The debiasing technique of using interaction terms to increase the magnitude of important predictor variables proposed in this paper has proved to be successful. In model seven, introducing interaction terms was successful in decreasing Type 1 error frequency to 11.96% for Black defendants. I believe that this technique has viable applications; with more time for research, I would like to explore this technique further and determine whether it can prove to be effective in increasing accuracy as well as reducing bias.

This paper is limited in the use of data solely from Broward County, Florida. Due to the proprietary nature of COMPAS and the difficulty of obtaining data on criminal history, the information available to assess and explore this topic is relatively small. Perhaps the greatest limitation in the data is the use of a proxy variable for recidivism. In the dataset, the variable for recidivism treats arrests as recidivism. This variable should instead use convictions to count recidivism, as arrests do not necessarily mean the defendant will be convicted. This paper is also limited in that it looks only at general recidivism, rather than violent recidivism. With additional time, predicting violent recidivism would further expand upon the analyses present. Future steps for this research include obtaining additional state or county-level data on criminal defendants and applying the techniques utilized in this paper. This would also include collecting data on future convictions and moving away from the proxy variable. The final idea that I would like to explore is proposed by Rudin et al. (2020), in which the authors suggest that risk and recidivism predictions depend nonlinearly on age (Rudin et al. 4-6). The research that I have conducted suggests both that these models can function with minimal inputs and that the inclusion of interaction terms can help to improve predictive accuracy and reduce bias. However, further research is required to further confirm these topics, as well as better identify the specific impacts of different variables.

VII. Conclusion

To conclude, my attempts to create a novel model for predicting the risk of recidivism have been successful. Compared with ProPublica's reports on COMPAS's accuracy, my proposed models have improved on accuracy as well as improved upon the rate of Type 1 and Type 2 errors when predicting recidivism. While my results suggest that the debiasing technique of including residualized variables as model inputs is not successful, further research should be conducted regarding this technique. My proposed technique of including interacted variables to add weight to major predictors of recidivism, however, has proved to be a successful technique for both improving accuracy and reducing biases present in predictive models. This technique can account for strong joint effects between multiple variables in the predictive model and has proven to improve both accuracy and bias. Regardless of debiasing techniques, the approach of minimizing inputs into the model has proven to be successful. This suggests that COMPAS's 137+ inputs are unnecessary. Clear, understandable models can be more effective at both conducting risk assessments and predicting the likelihood of recidivism.

Regarding policy implications, the success of my models with minimal inputs suggests that courts should prioritize adopting transparent models. The arguments for black-box models are understandable; Northpointe, the creator of COMPAS, would naturally want to maintain the trade secrets behind their model to secure their intellectual property. The use of trade secrets and preservation of intellectual property are important for innovation and discovery and are beneficial to society. However, considering a topic as detrimental to the individual as past criminal history, maintaining clarity is more important socially than promoting innovation and discovery. By relying on black-box models in which the reasoning behind criminal decisions is unclear, the public's trust in the judicial system will undoubtedly decrease. As we progress as a society, the use of computerized predictive models should require easily interpretable decision-making in order to maintain accountability and due-process rights.

Works Cited

Arnold, David, Will S. Dobbie, and Peter Hull. *Building Non-Discriminatory Algorithms in Selected Data*., 2024. *ProQuest.* Web. 6 Oct. 2024.

Bagaric, Mirko and Wolf, Gabrielle. *Sentencing by Computer: Enhancing Sentencing Transparency and Predictability, and (Possibly) Bridging the Gap between Sentencing Knowledge and Practice.* July 31, 2017. 25(4) *George Mason Law Review*, *Forthcoming.*

Engel, Christoph, et al. *Code is law: how COMPAS affects the way the judiciary handles the risk of recidivism.* 2024. *Artificial Intelligence and Law*: 1-23.

Hill II, Sean Allan. *Bail Reform and the (False) Racial Promise of Algorithmic Risk Assessment.* 2021. *UCLA Law Review* 68 (4): 910–87.

Dressel, Julia and Farid, Hany. *The accuracy, fairness, and limits of predicting recidivism.* 2018. *Sci. Adv. 4*, eaao5580.

Larson, Jeff, et al. *How We Analyzed the Compas Recidivism Algorithm.* 23 May 2016, *ProPublica*, www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm.

Lyn, Alexandra. *Risky Business: Artificial Intelligence and Risk Assessments in Sentencing and Bail Procedures in the United States.* December 16, 2020. http://dx.doi.org/10.2139/ssrn.3831441

ProPublica. *COMPAS Raw Scores.* 23 May 2016. *GitHub*, raw.githubusercontent.com/propublica/compas-analysis/master/compas-scores-raw.csv.

ProPublica. *Two-Year Recidivism Data.* 23 May 2016. *GitHub*, raw.githubusercontent.com/propublica/compas-analysis/master/compas-scores-two-years.csv. Accessed 19 Nov. 2024.

ProPublica. *COMPAS Questionnaire.* 23 May 2016. *DocumentCloud*, www.documentcloud.org/documents/2702103-Sample-Risk-Assessment-COMPAS-CORE.html. Accessed 19 Nov. 2024.

Rudin, Cynthia, et al. *The age of secrecy and unfairness in recidivism prediction.* Jan 2020. *Harvard Data Science Review*, vol. 2, no. 1, 31. https://doi.org/10.1162/99608f92.6ed64b30.

Tashea, Jason. *CALCULATING CRIME: Attorneys are challenging the use of algorithms to help determine bail, sentencing and parole decisions.,* 2017. *ABA Journal* 103, no. 3: 54-59.

Vaccaro, Michelle Anna. *Algorithms in human decision-making: A case study with the COMPAS risk assessment software.* 2019. PhD diss.

Washington, Anne. *How to Argue with an Algorithm: Lessons from the COMPAS ProPublica Debate*. February 4, 2019. *The Colorado Technology Law Journal.* Volume 17 Issue 1 http://ctlj.colorado.edu

Zeng, Jiaming, Berk Ustun, and Cynthia Rudin. *Interpretable Classification Models for Recidivism Prediction.* June 2017. *Journal of the Royal Statistical Society Series A: Statistics in Society*, Volume 180, Issue 3, Pages 689–722, https://doi.org/10.1111/rssa.12227

The Costs and Deterrence of the Death Penalty: Is There Evidence to Support its Practice?

By: Kristina McKay

Abstract

This paper looks at the death penalty through an economic lens. To do so, there must be a comparison drawn from life-without-parole sentences, as those are the most common alternative to a death sentence. As of 2024, there are 27 states with the death penalty and 49 states with life-without-parole. At the outset of this paper, the goal was to see the efficiency of practicing the death penalty. This means asking if the benefits exceed the costs, making it a rational practice. The costs and benefits can be a sum of many aspects, but the primary factors on each side are the fiscal costs of practicing the death penalty and the deterrence effect. Therefore, the deterrence effect should outweigh the fiscal costs for the death penalty to be efficient. Under the definition outlined in this paper, if the quantifiable deterrent effect exceeds the fiscal costs, then the death penalty is efficient. This paper concludes that there are cost increases that come from the death penalty and insufficient evidence of deterrence, making the death penalty is largely inefficient; understanding the economics behind this can help policymakers decide how to proceed with the practice.

The “data” for this paper’s unique analysis is the results of reports from opting-in states. This data analysis looks at the difference between fiscal costs of life-without-parole and the death penalty in each state. The analysis evaluates the significance of cost differences, finding that states spend, on average, double the amount on the death penalty as they do on life-without-parole cases. There is no raw data modeled, as it had not been made available, so results of fiscal cost reports by states or third parties are assumed valid. In addition, the deterrence studies collected data that appears to be unpublished, so this paper only evaluates their results and the validity of their findings.

Section I introduces the background of the death penalty and where research on the topic is limited, impacting the scope of the conclusions of this paper.

Section II focuses on the fiscal costs of the death penalty by collecting reports of state-by-state conclusions and analyzing their trends. First, the methodology of Maryland's fiscal cost report is summarized to understand how states find their total costs. This one analysis is representative because it gives details about all the different possible fiscal costs. Table 1 shows fiscal costs, their differences, and adjusts them to 2024 dollars. Finally, a brief discussion of why fiscal costs differ looks at what areas generally cause increased fiscal costs.

Section III focuses on deterrence. First, the methodology of three papers are discussed. The first paper concludes that there is a deterrence effect as a result of executions, the second says that the effects depend on the state, and the third finds no deterrence. As their methodology is discussed, the limitations of their methods are also evaluated to understand the scope of each paper. Following this, there is a discussion of what these papers can conclude. Lastly, there is a brief discussion of the qualitative aspect of deterrence, and the questions that prime the future of death penalty literature.

Finally, Section IV concludes that the practice is not economically efficient. Section IV then discusses the trends of reporting versus non-reporting states in respect to fiscal cost analyses. Section IV shows that the states that do not report their fiscal costs are the states with the most executions. The next section discusses the different areas of research that should be pursued in the future. Finally, I discuss how a new perspective might be needed when looking at the death penalty as it is practiced in the United States.

I. Introduction

A. Background

The practice of the death penalty has been a topic heavily debated since its reinstatement in the United States in 1977. Since then, it has been at the forefront of discussions about the American justice system. However, its federal constitutionality does not translate to it being legal everywhere. Further, just because it is legal, does not mean it is practiced. As of 2024, just under half the states have abolished the death penalty. An additional 12 states have not executed anybody in ten years. This indicates that, while legal, it is losing popularity. However, just the existence of the death penalty can be costly to a state that is not executing anybody. Each state has its own rule for capital punishment eligibility, meaning that the prosecution has the option to seek the death penalty. To do so, the prosecution would file a death notice, which establishes that they will pursue the death penalty. However, not all cases with a death notice end in a death sentence. Many cases instead end in a sentence of life-without-parole.

A frequent substitute and alternative to the death penalty is life-without-parole, otherwise known as LWOP. It is a fairly self-explanatory practice, as the defendant is sentenced to life with no opportunity to rejoin society. LWOP acts as a substitute in states where the death penalty is abolished. In states that still have the death penalty, it acts as an alternative sentence for capital-eligible cases. However, LWOP raises many of the same moral objections as the death penalty as it is a sentence to die in prison. They are similar where prisoners know they have no opportunity to rejoin the free world unless they can successfully appeal their case, which may also be seen as cruel and unusual punishment.

While the focus of cost and deterrent studies tends to be the death penalty, LWOP is often used as a comparison or the baseline for capital-eligible cases. To better understand these two forms of punishment, economists and legal professionals alike have studied the variables that go into the costs and benefits of each practice. Because of the complexity of the legal system and the abundance of quantifiable and unquantifiable factors, it is important to determine key factors to be studied. With the death penalty and LWOP, the most important factors are fiscal costs and the deterrence effect.

To determine fiscal costs, many states have compiled data reflecting the amount of public funding the death penalty takes, often that is done in comparison to LWOP cases. These fiscal costs are important to understanding whether or not practicing the death penalty is economically efficient. This efficiency is determined by whether the benefits resulting from the death penalty outweigh its costs. These benefits lie in the second modeled factor, the deterrence effect. The deterrent effect is a measurement of how punishments for crimes deter future offenders from committing crimes. Measuring deterrence has proven difficult because it tends to pair with the assumption of rationality. For crime to be deterred, the rational criminal will not find the crime efficient. The criminals' efficiency is when the expected benefits outweigh the potential risks. When looking at the deterrence effect of the death penalty, only deaths matter. The death penalty is generally reserved for murders, so for the death penalty to be deterring crime it should be deterring murders, therefore, saving lives. This means that the effect is measured in lives "saved" per execution based on murder rates.

Literature on the death penalty tends to focus on either the fiscal costs of practicing the death penalty or the deterrence effect. There is a limited selection of research that combines both factors to better understand the rationale behind the death penalty. This paper combines both areas of research by using all the data reported about fiscal costs and the available studies of deterrence to understand where arguments lie for and against the death penalty.

B. Limitations and Scope of the Paper

The study of the death penalty and LWOP are both hindered by limitations. Primarily, research can only be done based on the information provided by government agencies at the state and federal levels. Due to time and resource limits, only officially reported data will be evaluated. An additional difficulty from this starting point is that data and statistics are often reported differently agency by agency. This means that, while one state provides comparative statistics, others provide pieces of the puzzle. The issue appears when agencies only offer parts of the puzzle that make it difficult, if not impossible, to compare their findings to others. This difficulty with comparison will hurt the validity of any study that seeks to match findings, such as this one.

Further, studying the criminal justice system often requires long-term studies at county, state, and national levels. This poses a challenge because there are always changes that impact criminal activity and policing. For instance, the "war on drugs" increased incarceration rates, especially among minorities, because of changing laws and policing efforts. While it began in the early 1970s, the movement evolved with changing leadership and policies throughout the late 20th and early 21st century. This poses an issue to time-series studies, as they may only collect data from an unrepresentative period. Unfortunately due to findings being reported without raw data, it is outside the scope of this paper to adjust for these potential issues.

The different parts of this paper will have unique assumptions and limitations that will be laid out at the beginning of each section. These assumptions will change because each section uses a different methodology. Therefore, each topic will have individual methodology and limitations. Sections II and III will be quantitative, while Section IV will be partially quantitative and partially qualitative to understand other areas in need of studying.

II. Fiscal Costs

While it is uncomfortable to assign human lives numbers, understanding the death penalty without confounding its fiscal costs with its morality is essential to studying its impacts. Further, it is important to look at its fiscal costs in comparison to LWOP as it is, except for Alaska, the alternative to the death penalty (Death Penalty Information Center). Each state holds this information privately, so research can only be conducted if individual states permit it. Between 2000 and 2024 many states that continue to practice and no longer practice capital punishment have published findings of their fiscal costs through governmental agencies and independent parties. Where the details of states' findings are available, they are consistent.. Complex econometrics is unnecessary for these studies. Because of the relative simplicity of methodology, the most effective use of Section II.A. is to focus on the details of one paper. While the details for each paper lie outside the scope of this paper, following Section II.A. will be Table 1 which shows a comparison of relevant factors that differ among the papers. The paper of focus is "The Cost of the Death Penalty in Maryland" by Roman et al. from 2008. While it is older than other fiscal cost analyses listed, its rigor is unmatched as it details cost areas at every step (EJ USA). Roman et al. collect detailed fiscal costs and uniquely create three categories of cases instead of two, which lends to a more detailed comparison of where fiscal costs differ.

A limitation of the data available for this Section is combining the papers collected by the Death Penalty Information Center and a list created in "The Death Penalty vs. Life Incarceration: A Financial Analysis" by Torin McFarland. The states with available reports are Maryland, California, Connecticut, Indiana, Nebraska, North Carolina, Kansas, Arizona, Montana, Utah, South Carolina, Oklahoma, Ohio, Washington, Nevada, and Oregon, in addition to the federal government. With the DPIC, there is easy access to all of the papers; however, McFarland seems to cite numbers from elusive source materials. The cost difference reported for South Carolina is not found anywhere else, with "South Carolinians for Alternatives to the Death Penalty" concluding that there has never been a comprehensive study done in the state (SCADP 2023). The citations provided by McFarland seem to be missing the South Carolina source unless it is hidden behind a paywall. This raises concerns about the reliability of the South Carolina cost difference reported. A difficulty with DPIC's list is the low bar to be added to the list. For instance, the listed Florida report is a newspaper article from 2000 without source citations. As there is no confirmation of where the data was found, or how they analyzed it, the findings are unusable. Another state that poses a problem is Nebraska where the data comes from a journal behind a paywall. Due to a lack of access, only the reported numbers from McFarland can be used. Next, North Carolina's study focuses on savings, so there is ample data for the difference, but no clear finding for the total fiscal costs of LWOP versus the death penalty. Lastly, the Tennessee report listed by DPIC states insufficient data to conclude any fiscal costs, but the attempt to calculate is admirable. I also ran into many states having a fiscal cost analysis for policy purposes, but using the information of other states instead of their own state. Because of these limitations, the states in Table 1 are the only states with data to analyze. While I share much of the information on Table 1 with the table in "The Death Penalty vs. Life Incarceration: A Financial Analysis" by Torin McFarland, I researched each state McFarland included to ensure the correct information was present. I was only unsuccessful in finding the South Carolina and Nebraska sources, with the other state sources included in this paper's bibliography. All adjustments for inflation, and many of the reports in Table 1 were my independent research.

A. Methodology in Maryland

The study of Maryland took place in 2008 before the state abolished the practice in 2013. A limitation of this paper is its data comes entirely from 1978 until 1999, a common range of dates in death penalty literature. This poses an issue as fiscal costs are among the most variable factors in data over time. While this poses a concern for the validity of using a study with decades old data, it is important to understand how these changes would occur. While fiscal costs of prison have increased (Leigey and Schartmueller 2019, 247), this data remains usable. Because of the cost increases, the data only runs the risk of understating fiscal costs. As for cost comparisons, the same logic applies as a majority of factors are present in both death penalty and capital-eligible cases, simply to different degrees (i.e. inflated attorney fees will increase for both as the cost difference comes from time spent). This rationale allows for the limited supply of research on death penalty fiscal costs to be used for this paper. Further, it is outside the scope of this paper to make conclusions to the degree of specificity that would call the validity of the data in question.

Maryland's report offers a detailed analysis of the fiscal costs involved in capital cases. Roman et al. add a second layer to evaluating fiscal cost differences by splitting cases into three categories. The first category of cases is capital-eligible cases that *do not* file a death notice. The second category is cases that *do* file a death notice, but *do not* result in a death sentence. The final category is cases where the defendant is sentenced to death. The first and third categories are reported in every state's fiscal cost analysis, whereas the second is not as frequently reported. The first category, should the defendant be found guilty, ends in LWOP. This is true for every state besides Alaska.

Maryland studies 1,136 cases of murder that were capital-eligible from 1978 until 1999. The study compiles data from prosecutors, defense counsels, and judges through structured interviews to gather data on trial fees. The remainder of the data comes from the "Maryland Judiciary Case Search database and the federal PACER database" (Roman 2008, 17). 509 of the 1,136 cases had available complete administrative data, which was applied to the remaining cases after sample representation was evaluated. The authors do not explain why there is missing data, but they also do not indicate any self-selection bias of cases with full costs. Fiscal costs are gathered for the trial phase and incarceration phase of sentences. Due to a lack of information, some pretrial costs are unable to be added to the sums in addition to costs to appeal to the United States Supreme Court.

The original sample included 1,227 cases, but Roman et al. eliminated the not-guilty verdicts for the analysis, leaving the 1,136 cases. As previously mentioned, the calculations behind this study are minimally complex. The basic cost equation is:

$$Cost_i = Price\ of\ unit\ of\ input_i \times Quantity\ of\ inputs_i$$

The trial phase can be broken up into three parts: the guilt trial, the sentencing trial, and the state-level appellate costs. The remaining fiscal costs fall into post-conviction costs including health and housing costs. The paper offers a breakdown of where fiscal costs differ; however, for the scope of this paper, the important findings pertain to total fiscal costs. The regression analyses use p-values of 0.10, 0.05, and 0.01. The most commonly used was a $p < 0.01$, reflecting the high level of rigor. The notable finding from this paper is an overall fiscal cost difference of $1.9 million between cases with a death sentence and cases without a notice. In addition to per-case findings, Roman reports the fiscal costs of these 56 cases with death sentences cost the state $186 million. $7 million of this comes from the operation of the

Maryland Capital Defender's Division, which only serves death penalty cases. As it only exists to serve the death penalty, the entire $7 million per year spent on the division is a result of the death penalty's existence.

B. Findings

As discussed in Maryland's methodology, the death penalty cost Maryland an additional $1.9 million dollars per case until its abolition. In this section, the most recent findings from the states that have reported them will be compared to understand why these fiscal costs are important. The finding that remains consistent in every state is an increase in fiscal costs per death penalty case over LWOP cases. The Death Penalty Information Center publishes an updated list of papers that analyze the costs of the death penalty. This list includes the sixteen states in Table 1 plus the federal government. These sixteen states include three states that have since abolished the death penalty (Maryland, Connecticut, and Washington) and ten states that have not executed anybody in the last ten years including the three abolished states (Maryland, Connecticut, Washington, California, Indiana, North Carolina, Kansas, Montana, Nevada, and Oregon). As noted above, there is variation in the data provided by each state, with Maryland being the gold standard. Table 1 includes all the states with reliable reports cited by McFarland and compiled through independent research.

Table 1. State by State Fiscal Cost Information

State	Cost of a case eligible for the death penalty but not pursuing it resulting in LWOP	Cost of a case where a death sentence was handed down	Difference between the costs of LWOP cases and death sentence cases	How many times greater the cost of a death sentence is than the cost of LWOP	Difference between costs of LWOP cases and death sentence cases adjusted to 2024 dollars
Maryland*[1]	$1,100,000	$3,000,000	$1,900,000	2.72	$2,785,663
California`'	$627,322	$2,087,926	$1,460,604	3.33	$2,360,779
Connecticut*	$202,365	$380,000	$177,635	1.88	$87,112
Indiana'^	$42,658	$449,887	$407,229	10.55	$589,516
Nebraska^	No data	No data	$1,500,000	No data	$1,972,842
North Carolina'^	No data	No data	$2,000,000	No data	$2,942,746
Kansas'	$740,000	$1,200,000	$460,000	1.62	$789,160
Arizona`^	$70,231	$143,604	$73,372	2.04	$134,500
Montana'	$800,000	$1,200,000	$441,000	1.50	$606,322
Utah	No data	No data	$1,660,000	No data	$2,403,063
South Carolina^	No data	No data	$1,100,000	No data	$1,592,391
Oklahoma	$1,000,000	$3,500,000	$2,500,000	3.50	$3,387,578
Ohio`	$1,000,000	$3,000,000	$2,000,000	3.00	$2,710,062
Washington*^	$2,014,727	$3,073,612	$1,000,000	1.53	$1,331,820
Nevada'	$775,000	$1,307,000	$532,000	1.69	$709,369
Federal`	$77,618	$620,942	$543,324	79.9	$786,531
Oregon`'	$1,682,282	$2,569,667	$887,385	1.53	$1,183,240
Average			$1,096,621	3.29	$1,882,872

[1] * Indicates the state has abolished the death penalty since their study.
' Indicates the state has not executed a person in 10 years but still has the death penalty (2014-2024).
^ Indicates all life sentences are life without parole sentences.
` Indicates there is a hold/moratorium put on the death penalty by the state.

Table 1 can provide valuable information on how fiscal costs differ among states and the difference between cases. The majority of reports give information on both death penalty cases' fiscal costs and costs of eligible cases without the death penalty sought. These cost differences fell between $73,000 and $2.5 million, leading to an average difference of nearly $1.9 million when adjusted to 2024 dollars. While there is a large range of cost differences, it is important to note that fiscal costs all increase when comparing LWOP cases to death sentences in every study. The minimal increase among available data is $73,000, which is a significant deviation from the average. However, a more valuable finding is the percent increase in costs, as Arizona's monetarily small increase represents a 2 fold increase. When analyzing how many times greater the latter price is, it averages to a 3.29 times increase in costs with a standard deviation of 2.79, which is due to outliers, the federal courts and Indiana. These outliers are harmful to the data set because they are so drastically different from every other state with Indiana's costs increasing 10 times over and federal costs increasing almost 80-fold. Having these in the data set will skew the results to be higher than representative of the set because the mean is pulled up. With both removed from the dataset, the average difference is a 2.21 increase with a standard deviation of .77. This shows that, with outliers removed, a death penalty case costs on average two times more than a similar case that does not pursue the death penalty. Figure 1 shows the two different costs in a visual comparison.

Figure 1: State by State Summary of Costs with and Without a Death Sentence

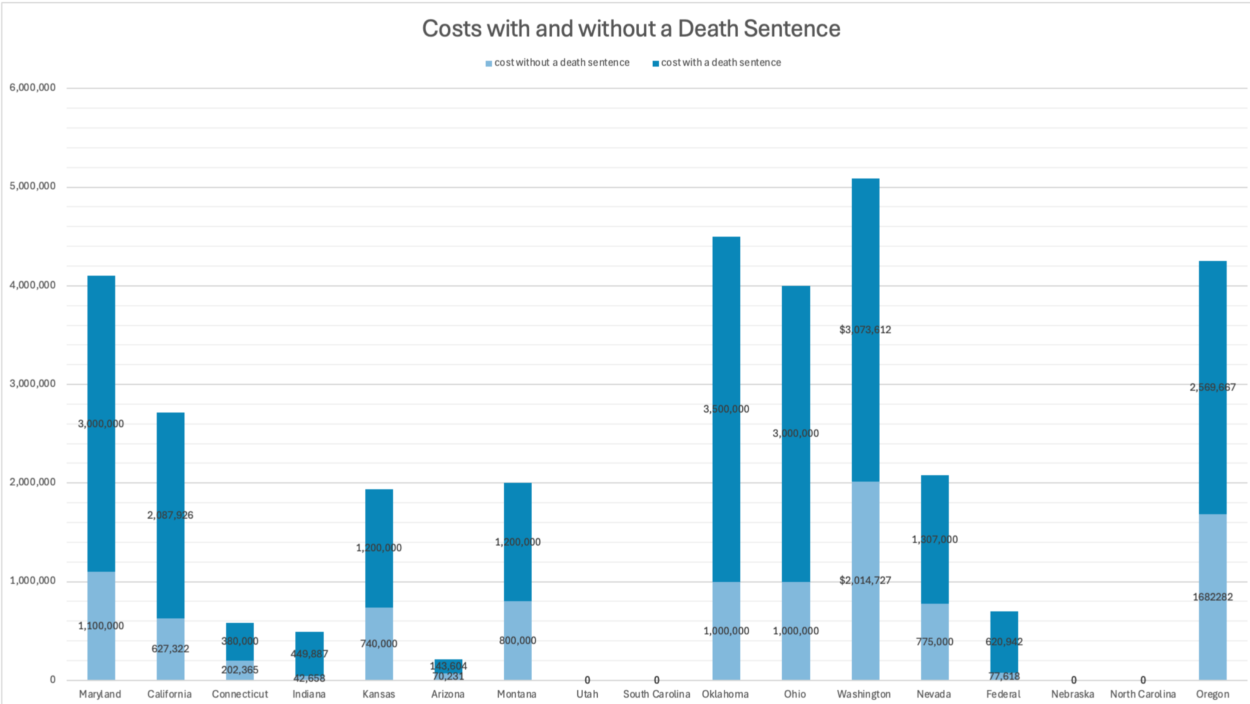

There should be significantly more rigor in future cost determinations to find a true average. However, within the scope of this paper, it can be concluded that pursuing the death penalty increases the fiscal cost of a case by at least 1.5 times in every state that has reported costs. More importantly, this is significant evidence that the death penalty poses greater fiscal costs than the alternative in every state. While this data is not robust enough to conclude specific amounts saved by eliminating the death penalty, it is enough to conclude that, in every state with a report, eliminating the death penalty would save money. It also debunks the misconception that LWOP must be more expensive because there are longer stays in prison with LWOP.

C. What might explain the differences?

As previously stated, it is outside the scope of this paper to conclude which costs have the most impact. However, there are a few areas where generalizations can be made to understand why the death penalty is more expensive. Many believe that LWOP sentences would have higher fiscal costs because people would spend more time incarcerated than when on death row. However, the time differences are smaller than expected.

In 2020, the Bureau of Justice Statistics reported capital punishment statistics on a national scale, including the average time between sentencing and execution. In Table 12 of Capital Punishment Statistical Tables from 2020, every year's average time is listed from 1977 to 2020. The notable trend is a steady increase in time served starting in 1984. While there are years where the average time decreases, these dips do not negate the trend. In 1984, the average time spent was 74 months, or just over 6 years. By 2020, this had surged to 227 months, or 18.92 years, with the highest time being 2019 at 264 months. Figure 2 shows how long prisoners spend on death row on average, refuting the idea that death row saves money by reducing the costs of incarcerating prisoners for life. Figure 2 below shows Table 12 provided by the Bureau of Justice Statistics. Figure 2:

Average elapsed time between sentencing and execution, 1977–2020

Year[a]	Executions	Average elapsed time from sentence to execution[b]
Total	1,529	147 mos.
1977	1	:
1979	2	:
1981	1	:
1982	2	:
1983	5	:
1984	21	74
1985	18	71
1986	18	87
1987	25	86
1988	11	80
1989	16	95
1990	23	95
1991	14	116
1992	31	114
1993	38	113
1994	31	122
1995	56	134
1996	45	125
1997	74	133
1998	68	130
1999	98	143
2000	85	137
2001	66	142
2002	71	127
2003	65	131
2004	59	132
2005	60	147
2006	53	145
2007	42	153
2008	37	139
2009	52	169
2010	46	178
2011	43	198
2012	43	190
2013	39	186
2014	35	218
2015	28	195
2016	20	204
2017	23	243
2018	25	238
2019	22	264
2020	17	227

Note: In 1972, the U.S. Supreme Court invalidated capital punishment statutes in several states (*Furman v. Georgia*, 408 U.S. 238 (1972)), effecting a moratorium on executions. Executions resumed in 1977 when the court found that revisions to statutes in several states had effectively addressed the issues previously held unconstitutional (*Gregg v. Georgia*, 428 U.S. 153 (1976) and its companion cases).

:Not calculated. A reliable average cannot be calculated from fewer than 10 cases.

[a]No executions were carried out in 1978 or 1980.

[b]Average time was calculated from the most recent sentencing date.

Source: Bureau of Justice Statistics, National Prisoner Statistics program (NPS-8), 2020.

In addition, my own data set compiled from data in Texas shows a similar result. This dataset is only composed of prisoners who were executed between 2014 and 2024 in Texas, as the data from the remaining inmates is not readily available. This data set of 83 executions found that the average time for this set of inmates between the arrest and execution was 27 years (Death Penalty Information Center).

This time can be compared to studies done by Louisiana, Pennsylvania, and Massachusetts cited by Leigey and Schartmuelle in "The Fiscal and Human Costs of Life Without Parole." These studies focus solely on incarceration time and fiscal costs for LWOP. Louisiana reports an average LWOP sentence being 16.3 years in 2017 and Pennsylvania reports 20.7 years spent on average in 2016 (Leigey and Schartmuelle 2019, 245), while Massachusetts does not report this finding. These times are similar to the findings of the DOJ and myself. While this data is not robust enough to go so far as to say death row may even tend to have people on it longer, it is evidence that executions do not significantly decrease incarceration time. This results in a lack of money saved from incarceration time.

In addition, Leigey and Schartmuelle report that LWOP prisoners are either in medium or maximum security prisons. However, death row is always maximum security. The cost differences between medium and maximum security can be substantial over time. The average daily cost of an inmate from 2017 federal data was $92 for medium security and $123 for maximum security (Federal Prison System Per Capita Costs 2017), while those costs were $65 and $83 respectively in 2012 (Leigey and Schartmuelle 2019, 247). Further, Massachusetts found their annual incarceration costs increased from $45,917 to $53,041 from 2010 to 2014. The daily costs of housing inmates add up over the decades inmates spend incarcerated, and seem to be increasing.

An additional area of fiscal costs are the trial phases. Looking at the data provided by Roman et al. in Maryland, phases of the trial differ greatly in length in the three different categories. The guilt phase for no death notice filed averaged to 237.2 days, which increases to 262 days with a notice filed and 312.7 days when a sentence is returned. The latter two categories both include a penalty phase that is not present without a notice filed. This phase is 100.2 days when a notice is filed and 152.4 days when a sentence is returned. Under the assumption that more time spent in court will always cost more money, it is clear that death sentence cases are substantially more expensive than those where the option is not pursued at all. This does not account for differing costs that come from attorney fees and expert costs along with other added discovery costs. While there is no obvious data about the impact of this, the Criminal Justice Act outlines that two attorneys are to be appointed in every capital case (Chapter 6, § 620.10.10 Federal Death Penalty Cases (a)), which indicates that attorney fees would automatically increase. One last area that increases costs for death sentences is the automatic appeal requirement by the federal government which is also practiced by some states. This creates fiscal costs for appeals that might otherwise not exist.

III. Deterrence

The deterrence effect has been heavily studied in criminal justice to understand the efficiency of laws. This is no different with the death penalty, as a primary argument in its favor is the increased deterrence provided by the death penalty. This argument lacks conclusive support from economists, as the collection of studies on the effect yields a variety of results. Of these studies, the yes versus no question of whether the effect exists is split, and further divided by the question of if the *executions* create the deterrence or if the existence of the penalty is enough. Ideally, I would perform an analysis myself to remedy the issues that I have with existing literature, but that task is much too demanding to be feasible. Therefore, like the fiscal costs, I turn to existing analyses of deterrence. My independent analysis will focus on what the papers are missing and opportunities for future research in the area.

Unfortunately, from what I can find from my research, there seems to be a lack of recent studies looking at deterrence. Therefore, the papers discussed here are from 2003, 2005, and 2009. More recent literature uses these papers, among others, to discuss the death penalty, but there does not appear to be significantly more recent studies with unique data.

These three papers are not the only ones that embrace the deterrent effect of the death penalty but are three that conclude different findings from each other. The paper "Does Capital Punishment Have a Deterrent Effect? New Evidence from Postmoratorium Panel Data" by Hashem Dezhbakhsh, Paul H. Rubin, and Joanna M. Shepherd finds a significant deterrent effect in the United States in a county-by-county analysis. Joanna Shepherd is also the author of the 2005 paper "Deterrence Versus Brutalization: Capital Punishments Differing Impacts Among States" which concludes that the deterrent effect varies among states. The final paper is by Tomislav Kovandic, Lynne M. Vieraitis, and Denis Paquette Boots called "Does the Death Penalty Save Lives?" in 2009.

A. Methodology

Among the three papers, there are only two methodologies, because Shepherd's second paper uses the same data as the 2003 paper she contributed to with the same analysis, merely applied state-by-state. Consequently, the methodology of Dezhbakhsh et al. can be equally applied to Shepherd's 2005 paper. Each paper has an in-depth methodology explained with modeling and equations where appropriate. For this paper, the importance of methodology comes largely from the variables, so those will be the focus of the following overviews. Subsequently, the discussion of these papers will look at the impact of their findings and whether or not their conclusions can be applied today.

Beginning with the methodology used in 2003 by Dezhbakhsh et al., panel data became the gold standard for researching deterrence, as the variables assessed below have been regarded as the most wholistic representation of the deterrent effect from the death penalty. Many papers about death penalty deterrence follow their formula, but I find it insufficient in evaluating deterrence. However, the field does not appear to have a subsequent methodoology that has fixed the mistakes of Dezhbakhsh et al. That paper was also the first to use county-by-county data, rather than state-by-state data. Data was taken from 3,064 counties in the United States between 1977 and 1996. It is argued that this increases the validity of their findings by removing overgeneralizations. However, this is a bit ironic when the conclusion in their paper is proved to be much more nuanced (to a degree that the prior may even be misleading) in Shepherd's second paper. In the prior, they were searching for a single deterrence value for the whole country, while Shepherd was looking at deterrence state by state, which is more appropriately aided by the county-by-county data. I also worry county-by-county is overly specified, as there is no given reason to believe counties will not impact each other. For instance, in Massachusetts, Middlesex and Norfolk counties surround Suffolk County which contains Boston. These counties are likely highly connected, as American culture often sees people living in the suburbs, which would fall outside Suffolk County, but working in cities, as Boston is in Suffolk. These workers would likely not only be impacted by actions in the county they live in, but surrounding areas. Further, when counties share this proximity, they will have a spillover of recreation and visitation. There is no accounting for the spillover effects from other counties, making the causal relationship questionable. The data may unintentionally isolate impacts that are better explained by surrounding actions.

This panel data uses a linear model, which the authors argue is even less likely to find deterrence based on previous deterrence studies, making it more conservative. The model uses murder arrests as its dependent variable, meaning they also include non-deterrable crime. In addition, they look at aspects within the criminal justice system and law enforcement to understand their effectiveness. They also evaluate political influence by measuring the Republican presidential candidate's percentage of votes. While there is no analysis of the impact of this variable, it seems like a weaker representation of the politics of an area than looking at their local or state representation. This variable is included to understand the presence of "tough on crime" policies, which were "popular with Republican candidates." (Dezhbakhsh et al. 2003, 357)

The model also includes aggravated assault and robbery as they are connected with the death penalty, by providing the background for non-deterrable murders. The idea of non-deterrable crime is that these are crimes that would be committed no matter what preventative measures were in place. With murder, the non-deterrable crime is considered murder that was not part of the intentional crime. This is why aggravated assault and robbery are included, as the planned assault or robbery may not have included a resulting murder. These are used as control variables along with economic and demographic factors. The last control variable is NRA membership to understand gun culture in the counties.

Within the models, the authors use two different lags to measure execution and sentencing probabilities. The prior is found from (t-6)/t while the latter is (t-2)/t. They include t-6 for the lag from arrest to execution assuming the average time between arrest and execution is six years. T-2 is meant to represent the time between arrest and sentencing. These adjustments are used to look at the impacts in real-time, as results that come directly after an arrest will not have any deterrence from that arrest (concerning the death penalty). While t-2 may still be accurate when looking at the total trial time coming in just under 500 days for death sentence cases (Roman et al. 2008, 24), the t-6 assumption is questionable. The authors argue for its use because of a six-year assumption of average time spent on death row. However, the Department of Justice reported the average time spent on death row between 1984 and 1996, which does not reflect his number. Figure 2 above shows this data as it was previously discussed.

While there is no data available from 1977 through 1983, the average time spent from 1984 through 1996 is 100 months or 8.4 years. This is greater than the assumption made by Dezhbakhsh et al., calling into question if this assumption yields correct results. Notably, there is a 50-month difference between 1977 and 1996, which is just over four years. By 1996, the national average was over ten years, substantially different from the six-year assumption made. The shortest year average is just over six years, but the trend indicates the missing years were likely less. The biggest trend shift comes in 1991 with a 20-month jump, meaning the years where the deterrence findings are likely least representative of any year after 1990. When the lag is not accurate, it calls into question if the deterrence can be properly evaluated as deterrence is trying to prove a causal relationship. When the lag is six years, it is saying the arrests this year will be executed in six years, so today's deterrence will reflect the arrests of six years ago that are now being executed. In theory, this should impact the perceived probability of execution. However, when this lag is wrong it may reflect arrests from too large a range of years.

The above lags are used to estimate the perceived rates of offenders. However, they fail to account for a plethora of confounding variables that will impact the offender's perception. Because it does not appear any county is reporting its execution rates, it can be assumed that perpetrators are creating their own expected rates. As read in McAdams, expectations are heavily influenced by heuristics that do not appear to be controlled (McAdams 2008, 5). For instance, how widely publicized the execution is may impact a perpetrator's availability heuristic. The way this heuristic is understood to operate, readily available information is overvalued, might even indicate one largely publicized execution should have a greater impact on perceptions than multiple executions entirely out of the public eye. The deterrence effect assumes that the salience of executions comes from the number of executions rather than the publicity of them.

There may also be the presence of optimism bias that impacts the perpetrator's perception of arrest likeliness and execution probabilities (Jolls 2004, 7). It might make the model additionally robust to include organized crime rates, as that will incite the anecdotal heuristic and may reflect important trends of sentencing (Knoll 2010, 2). The anecdotal heuristic is a tendency for people to overvalue information from their peers as representative of the general population. Here, the anecdotal heuristic within organized crime may be that members of a gang that just had one of its members executed would see their probability of execution heightened because of their proximity to the information. Future research would benefit from looking at the presence of organized crime in this perception as it might impact jury bias or bolster optimism bias for those outside of organized crime. The prior could include surveying for sentencing differences between cases where the only difference is their participation or lack thereof in organized crime. This would give insight into whether juries are biased towards or against sentencing gang members to death. The latter would be valuable in seeing if perpetrators are aware of risks related to being in organized crime. For instance, if the executions are all members of organized crime, then somebody without affiliation may see their likelihood of a death sentence lower than expected.

Once all the variables were established, the authors ran six models to test for sentencing deterrence and execution deterrence. These findings are all statistically significant at $p < 0.05$ level of significance. The paper concludes an expected deterrence of each execution deterring eighteen deaths. However, this comes with a ten-person standard deviation, which is a concern as Shepherd finds that reverse effects and lack of effect are real and the values become negative within two standard deviations of their conclusion. However, the authors perform robustness checks that support their conclusion, and this study is peer-reviewed and accepted, which is why Sheperd opts to use it for her subsequent study.

In this 2005 study, Shepherd takes the same county-by-county data and concludes differently in each state. While Shepherd maintains the original conclusion of an overall deterrent effect, her results are much less favorable to the practice of the death penalty. Her significant findings are in her paper as Figure 2 shown below as Figure 3

Figure 3: Individual State Deterrent Effects: Number of Murders Deterred or Incited

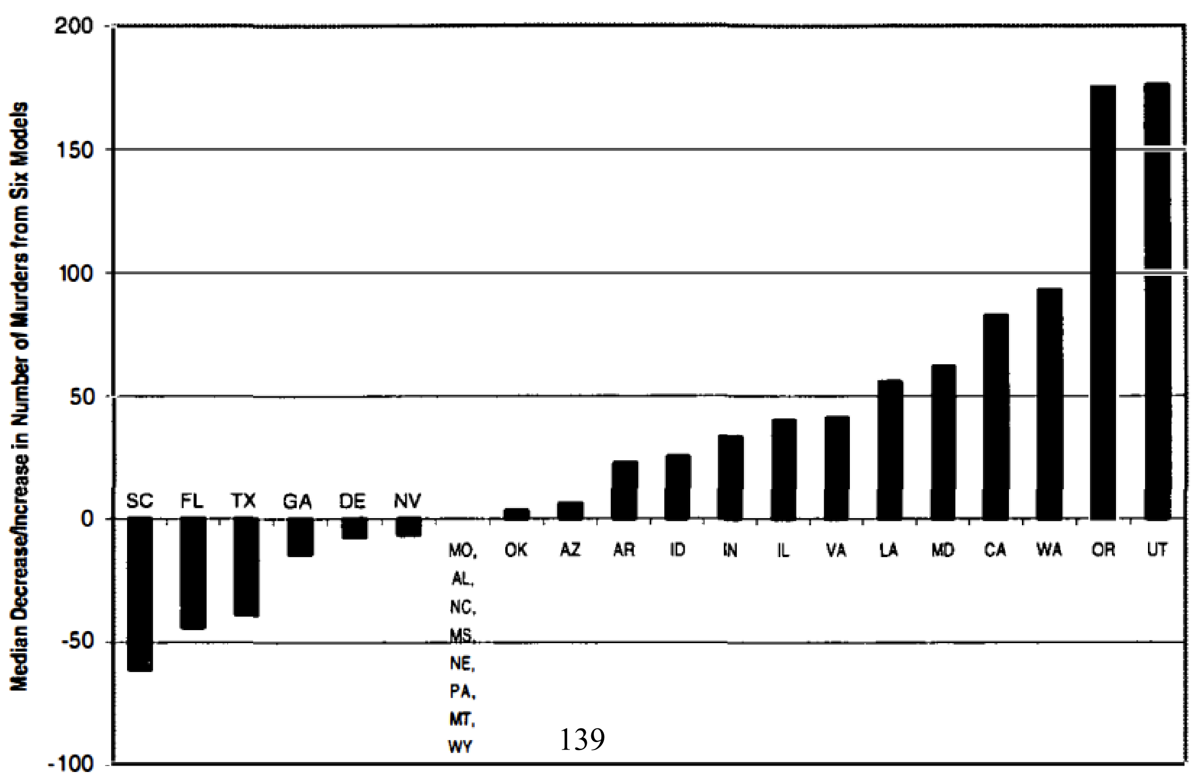

Figure 3 reflects the deterrence effects state by state measured by murder rates among the six performed models. The y-axis is the median change in number of murders from the six performed models, meaning each state's data is a result of Shepherd's analysis results. The x-axis is the states that were involved in the study arranged from lowest to highest number of murders (in value not magnitude). These results paint a very different picture than the previous conclusion, as the number of states with no deterrence exceeds those with it. The only states that are found to have a deterrent effect are South Carolina, Florida, Texas, Georgia, Delaware, and Nevada. The remaining states either show the "brutalization effect"[2] or a lack of impact. Shepherd also embraces the concept of the threshold effect, which is the theory that there is a certain threshold for the number of executions to provide a deterrent effect. Under this, Shepherd believes that executions have to hit a certain amount to have an effect. This hypothesis is supported by the states that find a deterrent effect having high execution rates. Shepherd maintains the conclusion from 2003 because there are more executions that take place in deterrent states than in brutalization states. There were 192 executions that saved 6,918 lives in deterrent states and 112 executions that caused 5,358 murders. This results in the overall deterrent effect found by Dezhbakhsh et al. in 2003, which appears to misrepresent the findings to those who are not interested in the nuances of their findings.

The final paper discussed is from 2009 which is also a panel analysis, but uses somewhat more recent data from 1977 until 2006. This study is done state by state rather than county by county. Again, the variables will be the main focus of this section as the models themselves do not raise any questions, nor does it appear the chosen model impacts the results. The 2006 study uses different variables in their analysis, but the same murder rates variable to evaluate

[2] This paper will not look at the brutalization effect, however, it is the effect of an execution creating a culture of violence that, in theory, encourages homicides. The discussion of the effect in the remainder of the paper to categorize states with increased murders.

deterrence. The authors included a new control variable for the crack epidemic along with policy evolutions over the time data time period. These variables are three-strikes laws and right-to-carry laws. The prior is similar to the tough-on-crime control used by Dezhbakhsh and the latter is likely from the same category as the NRA variable (Kovandzic 2009, 814). The new variable comes from the crack index to control for its impacts on homicides. Finally, there are also demographic socioeconomic variables that serve the same purpose as similar variables in Dezhbakhsh's study. These include income, unemployment rates, and population density, among others that have been previously found to correlate with murder rates.

Kovandzic also uses the gap time of six years from sentence to execution as that appears to be the widely accepted time period. As previously explained, this causes concern as to the validity of probability findings that this gap is used to calculate. By 2006, the average time on death row had grown to 145 months, double the initial 74-month data. The use of t-6 is questionable as it no longer represents the gap between arrest and execution. They do not explain why they choose gaps instead of taking data from executions each year.

The results show no deterrence effect which is far from the eighteen murders deterred found by the other two papers. This paper also checked for the robustness of their study and found no issues. Following their conclusion, they state the different conclusion from Shepherd and Dezhbakhsh is the result of omitted variable bias. They discuss how the crack epidemic variable is unique to their paper, so that may be the omitted variable they mean, but they do not elaborate on which variables are considered omitted from the initial analysis.

B. Discussion

Deterrence is a unique area of criminal law literature because it always seems like there is something not accounted for by the models, as there are variables within variables that impact the results of studies, as detailed below. This is a result of the human component of the crime making the theory of rationality difficult to rely on as seen by the different results from the papers available having a variety of conclusions. The beginning of the difficulty behind deterrence is how the study is rooted in rational choice theory assumptions, including that actors are rational and will, therefore, choose the efficient option. This assumption makes criminal behavior modelable but is likely misrepresentative of actors. The assumption of efficiency would conclude that actors will choose to murder should they perceive that the benefits exceed the expected risks or costs. This is where the probability of arrest comes into play, as the possibility of arrest, and consequently, death row is the primary risk. Because of this, it is assumed that rational actors will choose to not commit a crime should their expected risk of execution exceed the expected benefits. However, this conclusion may only work in theory. All non-detterable crime is outside of rationality, as this crime results from something else. For example, a criminal might rationally decide to commit aggravated robbery but then kills somebody they never planned to kill. This kill was not accounted for in the perceived risks, meaning the murder might fall outside of rationality, but was a choice made in the heat of the moment. The existence of this crime decreases the likelihood that the deterrence models explain the real impact.

Rationality is also heavily influenced by biases and heuristics, as they often push decision-makers away from the rational choice. However, there are no variables that attempt to control for these biases or evaluate their impact. The models seem to assume the existence of recency bias, as deterrence is assumed to be a response to a recent execution. Recency bias claims that people favor recent events when considering probability, which is the root of saying executions will impact murder rates of the year they occur. The models assume that all executions are the same in impact and the most recent ones will have the most impact. A variable that could account for differing levels of publicity of the execution and the case can help explain if all executions really have the same deterrent effect, or if there are certain aspects of executions that are more effective in deterring than others. Additional research is required to learn what might measure this impact and measure the public knowledge of an execution. A possibility may be public opinion on the death penalty being impacted by a high-profile execution or sentencing. Researching this might indicate that availability bias has a greater impact, as this does not rely on when the event occurred but on how quickly a person thinks about it. This specificity is something I take more issue with in the Dezhbakhsh paper, as they emphasize the unprecedented specificity of their analysis. When they make claims about their study being better than the others, it is disappointing to see variables left out that would likely be useful with county-by-county data.

The most interesting results come from the 2005 Shepherd paper as there is more than one finding reported. The state-by-state analysis shows six states have a deterrent effect from executions, eight see no effect, and thirteen show a brutalization effect. These twenty-seven states were not all of the states with the death penalty at the time, as it appears New Jersey, New Mexico, New Mexico, Colorado, and Connecticut are not included among the states, all of which have since abolished the death penalty. Among the states included, Maryland, Washington, Illinois, Virginia, and Delaware, have abolished the death penalty. It begs the question of if this case was part of that choice as all of these states besides Delaware show a brutalization effect. The subset of states that show deterrence is said to be a result of a threshold of executions met.

These states are South Carolina, Texas, Florida, Georgia, Delaware, and Nevada. Since 2014, Nevada has not executed anybody and South Carolina only has two in 2023. As seen in Table 2 (located in Section IV with analysis) Texas, Florida, and Georgia remain the states with the highest number of executions. They are joined by Missouri, Alabama, and Oklahoma with over ten executions since 2014. Should one assume deterrence results from this study are sufficient evidence for the efficiency of the death penalty, Texas, Florida, and Georgia are the only states that can argue their practices are effective, as their executions save lives. However, the more important takeaway is that, assuming these findings are valid, Missouri, Alabama, and Oklahoma are perpetrators of unnecessary executions, costing the states millions of dollars and, for Oklahoma, costing lives of victims from an increase in murders.

C. Qualitative Deterrence

Thus far, studies have assumed that potential perpetrators are deterred more by executions than by spending life in prison. The assumption that a person choosing to commit a capital-eligible crime knows they could qualify for the death penalty or LWOP is fairly safe with the exception of non-deterrable crime. However, the assumption that a person is more deterred by the death penalty than LWOP does not seem to be based on empirical evidence. It is helpful to turn to a qualitative test that challenges this assumption by speaking with inmates. While this survey falls victim to biases with their pool of respondents, it is a challenge to common beliefs to hear that any inmates would have the following views. Because deterrence studies are rooted in this assumption it is important to understand how valid this assumption really is. When interviewed about their opinions on serving LWOP or a death sentence, prisoners tended to respond with a preference for death row, as they saw LWOP as a “harsher punishment than death." (Appleton and Grover 2007, 607) This indicates that their risk assessments may not have increased with the probability of execution.

There is also significant opposition to the death penalty among inmates, but not based on self-preservation. The study of 309 prisoners in Ohio modeled opinions on the death penalty of prisoners. These prisoners were interviewed with their responses being indicative of sanctions having less influence on actions than their beliefs (Steele and Wilcox 2003, 305). These inmates shared a lack of belief in the deterrent effect of executions, rooted in their view of a prevalence of non-deterrable crime. Their opinions indicate that non-deterrable crime makes up a significant amount of capital crimes. However, it is important to note that this sample comes from people who were not deterred by sanctions themselves. Between a potential trend of preference for death over LWOP and criminals indicating a lack of deterrence, there is reason to question the assumption that executions are a worse punishment than LWOP.

It is easy for me to sympathize with the inmates, to the point where I forget that these punishments are in place because people have committed the unthinkable. The families of victims deserve to feel like justice has been served even if nothing can be done to make up for the pain. This retribution feeds into how just people think the system is. I have the privilege of discussing these punishments detached from the reality of it all, so for a moment it is important to discuss the value of retribution for victims. Looking at these studies through the view that retribution is the primary goal, I feel concerns for execution's effectiveness. Families and friends of victims knowing the perpetrator may prefer their death sentence may take away their sense of justice about the sentence. In theory, those seeking retribution want to inflict as much pain and suffering on the perpetrator as they have felt. When the assumption is that execution is worse, those who support the death penalty are indicating a desire for the maximum suffering inflicted on the perpetrator. When there is uncertainty about what truly is the greatest suffering, it may leave families uneasy about whether their desire for retribution has been met. The results of these surveys are in no way claiming that LWOP is a more extreme punishment than execution. However, when something is a fundamental assumption for modeling, it is important to understand to what extent it can be assumed.

IV. Conclusion

The field of death penalty research is far from comprehensive and desperately needs to be updated as trends change. There are some results that can be concluded with the costs. There is also a large opening for future research and the field may benefit from a new way of thinking about the status quo.

A. Findings

The two empirical aspects of the death penalty have very different levels of conclusivity. The first conclusion of this paper is that the death penalty's existence increases the fiscal costs of the state two-fold on average. The increases in fiscal costs vary from $73,000 to $2 million per case. This is concerning as these costs exist at every stage of a case even if no sentence is returned, meaning that the mere existence of the death penalty will increase the state's costs.

The second conclusion of this paper is that deterrence evidence is not sufficient to conclude that a deterrent effect exists as a result of executions. There is reason to believe that some states may see some of this effect. However, the studies are outdated to the point that there may be no viable conclusion to draw at all.

Therefore, the practice of the death penalty is economically inefficient and should not be practiced if the states act rationally. Should there be conclusive evidence that there is a deterrent effect from executions, then there would be efficiency behind the practice where that deterrence is found. However, the states where they do not execute anybody cannot be efficient, as they incur the fiscal costs of the death penalty without the potential deterrence.

B. Reporting Comparison

Criminal justice poses the always difficult issue of significant barriers to research. This requires unusual actions to be taken by researchers to access raw data which makes analysis difficult, and at times impossible. This is an issue posed while studying the fiscal costs of the death penalty when looking at the trends of states that do and do not report. Of the states that still have the death penalty, thirteen states have created reports on fiscal costs in addition to federal data. These states are California, Indiana, Nebraska, North Carolina, Kansas, Arizona, Montana, Utah, South Carolina, Oklahoma, Ohio, Nevada, and Oregon. Florida has a journal article from 2000 that appears not to be from a reputable journal. In addition, the inclusion of South Carolina is based on McFarland's table, rather than an independently read paper, which is contradicted by SCADP saying there has never been a fiscal cost report of South Carolina. In addition to the thirteen states with reports that still practice, three states have since abolished the death penalty, and of those thirteen states, four plus the federal government have placed holds on the practice since their reports. This means of the sixteen reporting states, seven of the sixteen states plus the federal government have altered their practice since their report, with the federal government executing three people in January of 2021 then placing their moratorium in July of 2021. While it cannot be determined as causal, this pattern is notable and might speak to the impact of having a fiscal cost report. However, only three of the twenty-three states that have abolished the death penalty have fiscal cost reports, so the link should not be overstated. More interesting is the difference between the states that still have the death penalty that did and did not report.

The states that do not report, but still have the death penalty are Alabama, Arkansas, Georgia, Idaho, Kentucky, Mississippi, Missouri, South Dakota, Texas, Wyoming, Florida (to my understanding), and Tennessee. Removing Maryland, Connecticut, and Washington from the set, Table 2 shows the execution history of the reporting states, followed by states that did not report but have it.

Table 2: Executions by State 2014-2024

State	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	10 years
California	0	0	0	0	0	0	0	0	0	0	0	0
Indiana	0	0	0	0	0	0	0	0	0	0	0	0
Kansas	0	0	0	0	0	0	0	0	0	0	0	0
Arizona	1	0	0	0	0	0	0	0	3	0	0	4
Montana	0	0	0	0	0	0	0	0	0	0	0	0
Utah	0	0	0	0	0	0	0	0	0	0	1	1
South Carolina	0	0	0	0	0	0	0	0	0	0	2	2
Oklahoma	3	1	0	0	0	0	0	2	5	4	3	18
Ohio	1	0	0	2	1	0	0	0	0	0	0	4
Nevada	0	0	0	0	0	0	0	0	0	0	0	0
Federal	0	0	0	0	0	0	10	3	0	0	0	13
Nebraska	0	0	0	0	1	0	0	0	0	0	0	1
North Carolina	0	0	0	0	0	0	0	0	0	0	0	0
Oregon	0	0	0	0	0	0	0	0	0	0	0	0

The remainder of Table 2 is states that did not report their costs

Alabama	0	0	2	3	2	3	1	1	2	2	5	21
Arkansas	0	0	0	4	0	0	0	0	0	0	0	4
Georgia	2	5	9	1	2	3	1	0	0	0	1	24
Idaho	0	0	0	0	0	0	0	0	0	0	0	0
Kentucky	0	0	0	0	0	0	0	0	0	0	0	0
Louisiana	0	0	0	0	0	0	0	0	0	0	0	0
Mississippi	0	0	0	0	0	0	0	1	1	0	0	2
Missouri	10	6	1	1	0	1	1	1	2	4	3	30
South Dakota	0	0	0	0	1	1	0	0	0	0	0	2
Texas	10	13	7	7	13	9	3	3	5	8	5	83
Wyoming	0	0	0	0	0	0	0	0	0	0	0	0
Florida	8	2	1	3	2	2	0	0	0	6	1	25
Tennessee	0	0	0	0	3	3	1	0	0	0	0	7
Pennsylvania	0	0	0	0	0	0	0	0	0	0	0	0

The red highlighted cells indicate where there were executions in that year. The difference between the states that reported and those that did not is astounding. Future research would benefit from understanding state sentiment about the death penalty, however, my research does not reveal a study that includes results from independent states. (There are independent state reports, however, each was done by a different organization, and it is outside of the scope of this research to review the methodology of each study to determine credibility.) Excluding the federal government from analysis, as the federal government's use of executions has been unique, the thirteen states that reported had a combined 30 executions in the past ten years, while the remaining fourteen states have had a combined 198, half of which are from Texas. Figure 4 shows this below, taken from the data in Table 2. Figure 4:[3]

[3] The (blank) seems to be a processing error and I cannot remove it. It appears to be a duplicate of Florida.

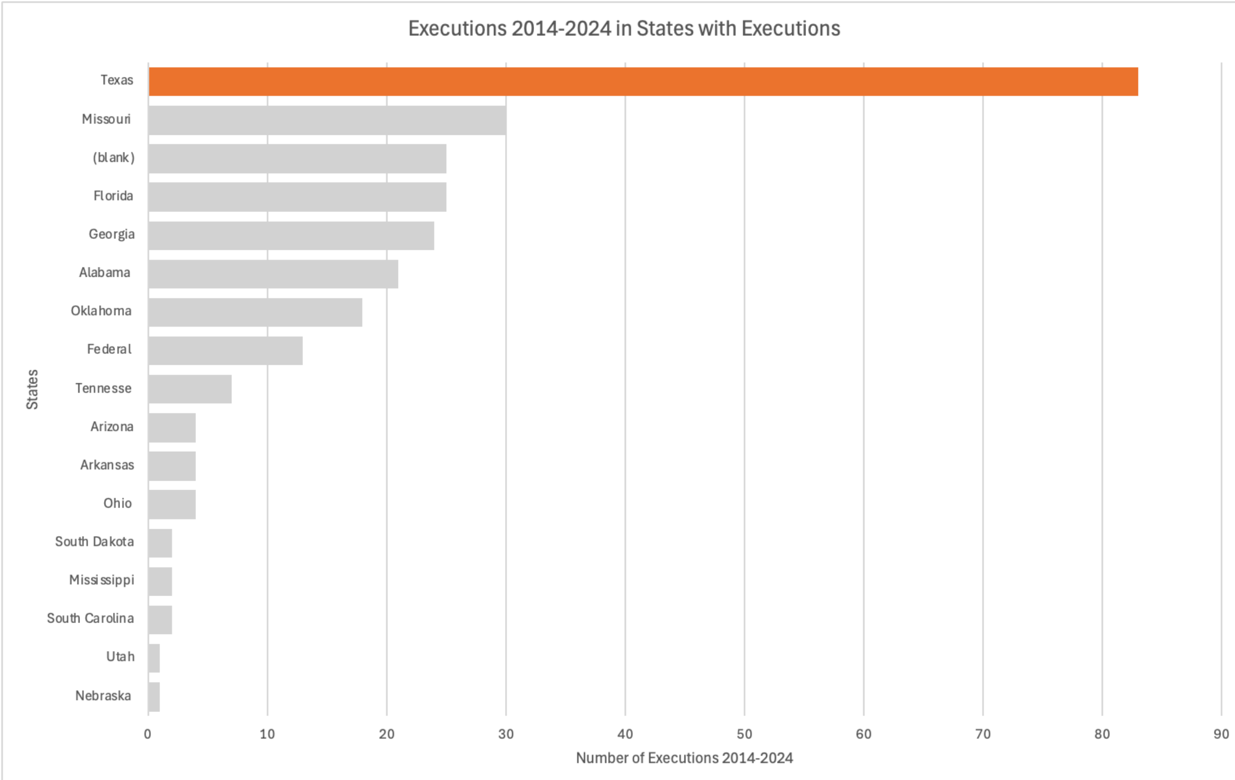

The same pattern follows for the distribution of states with executions in the last ten years as six of thirteen states reporting have carried out executions, while nine of the fourteen not reporting have executions. This indicates that there are fifteen states that carry out executions still, while the other twelve have the death penalty without enforcing it. However, all twenty-seven states had people on death row as of December 31, 2020.

The states that did not report findings have combined executions six times greater than the reporting states, however, that is partially due to an outlier. Texas executed 83 people in ten years, distantly followed by 30 executions in Missouri. If Texas is removed, the 13 states not reporting fiscal costs still have 115 executions which is 3.5 times greater than the others. In addition, only six states have executed more than ten people in ten years, and all but Oklahoma, the lowest of the bunch, did not report costs. To avoid an average skewed by Texas, the medians for both categories were taken and are notably different. The median of reporting states is four executions (excluding states with zero executions), while the median for non-reporting states is twenty-one. It is clear that the states that practice the death penalty the most tend to fall into the non-reporting category, with the exception of Oklahoma. It seems counterintuitive that the data about the death penalty is concentrated in states that are not the main perpetrators of it. It would seem that the states who practice it the most would want to understand its details. However, this does not appear to be the case.

In addition to the difference in magnitude between the two categories, the frequency is quite telling. Looking first at the reporting states, three of them only report executions in one year while the other four report executions in between two and six years, with none exceeding six years. The non-reporting states look very different. Of these states, Arkansas is the only state that executes in only one year. The others range from two to eleven, with five of the eight states executing in more than six years. This is evidence that not only are these non-reporting states executing the most people, but are executing with the most consistency.

Due to the self-reporting nature of this data, it is a reflection of the state to not make a full effort to understand this program. It would be unlikely that those states would see inverse cost results with the death penalty, meaning they are all expected to see higher costs.

C. Future Research

There is a drought of deterrence studies, as my research revealed no new studies after 2020. This means that policymakers may be using decades-old data that is no longer reflective of the effects of executions on murders. The potential change could come in the form of an adjusted lag period that is more reflective of the actual time between executions. However, that may still lead to old data as the average time in 2020, the most recent report from the Bureau of Justice Statistics, is a 227-month lag, which is 18 years. This lag would look nothing like the 6-year lag present in available deterrence literature, and may impact the perpetrator's expected risk estimation and the fiscal costs incurred by the state. With an average time of 18 years spent on death row, the offender may have an overwhelming present bias (McAdams 2008, 23). The benefits of committing the crime are instant, whereas the punishment would not come for almost two decades. Looking state-by-state, many may not even have a lag to evaluate at all. As of 2024, there are over 2,000 people on death row across 27 states (Criminal Defense Team Baldwin Perry & Wiley 2024). Those 27 states are all of the states that still have the death penalty in place, however, many of them have not practiced it in the past decade. California, Indiana, Kansas, Montana, Nevada, North Carolina, and Oregon have no executions but still have inmates on death row. As spoken about in section II.C., the fiscal costs of housing a maximum security inmate are always higher than medium and minimum. In addition, death row inmates are housed independently (American Civil Liberties Union 2013), meaning each inmate gets their own cell, costing more than shared cells. These states are all spending more money to have death sentences without the theoretical deterrent effect. When there are high costs with seemingly no benefit, the practice becomes inefficient and wasteful of taxpayer dollars.

Another facet of deterrence that would prove effective to study is the difference in deterrence between a LWOP sentence and a death sentence. As Dezhbakhsh et al. conclude, the deterrent effect of death sentences is far less than the deterrent effect of executions. That paper concludes that executions save lives, death sentences do not. In future research, it might be telling to see how the extended time on the death penalty might impact deterrence. As discussed above, the average time in 2020 spent on death row is 18 years. Section II.C. calls attention to how LWOP average sentences in Louisiana and Pennsylvania are 16 years and 20 years, respectively, with the death row time of 18 years sitting right in the middle. Roughly two decades are found to be spent on both death row and LWOP sentences. The only way out of a LWOP sentence is death in prison, so it can be assumed that these averages of 16 and 20 reflect years before death in prison. Should a potential perpetrator have this information when calculating their expected risks, they may calculate that they would spend the same amount of time on death row as serving LWOP, both ending in their death. This may bring these two sentences somewhat even in deterrence, should longevity of life be one of their considerations. There is also more room to study the effect of pure sentencing without follow-through as there is a subset of states that can be used for comparison. However, this study would assume that executions do have a deterrent effect in these states, which is not a proven hypothesis as of 2024.

It may also be an interesting area of future research to see how the most common crime prevention programs compare in lives saved to deterrence studies that conclude a deterrent effect. If we assume the fiscal costs of the death penalty are the fiscal costs of deterrence, those costs can be compared to the costs of prevention programs. The existing number of lives saved by deterrence should be used as the goal of lives saved by prevention programs. Once that number is determined, the fiscal costs for those programs to get there can be compared to the active fiscal cost of the death penalty as it differs from the fiscal cost of LWOP.

For future research to happen, there needs to be more open access to the raw data of incarcerations within states. This will allow for more cost analyses to be run, especially in those states with high execution rates. There also appears to be an urgency for correcting the death sentencing system, as just under one-third of sentences are later overturned. (Johnson and McGunigall-Smith 2008, 344) When somebody's life is on the line, time is of the essence.

Further, 200 inmates have been exonerated resulting from a variety of issues since 1973 (Death Penalty Information Center). When it is clear that so many people sit on death row who are not meant to be there, a sense of urgency fuels research.

D. Putting the Death Penalty on Trial

While this study is equally unable to decisively conclude the existence of deterrence as those that have come before, this inconclusivity is important on its own. For this section, we flip the script and imagine the death penalty is being proposed by policymakers rather than being the status quo. It is hard to imagine this policy being adopted on the scale that it currently exists. The lowest burden of proof in the legal system is a preponderance of the evidence, which does not even come close to the burden for criminal cases. The evidence in support of deterrence does not meet even this low burden. The existing evidence is inconsistent and dated, meaning it is not convincing that it is more likely than not that deterrence exists. However, it is consistently found that the death penalty costs substantially more to the taxpayers than LWOP, with every state showing increased fiscal costs. Further, Maryland shows how just having a death notice filed costs substantially more than no notice being filled. It makes one think if there is so little empirical evidence of the death penalty doing what it is supposed to, why does it still exist? This question is most urgent in states that still have the death penalty but have not executed anybody in over a decade. They continue to waste money housing inmates and even sentencing new people to death row (Death Penalty Information Center). California has over 600 people on death row and in 2023 sentenced one more person, even though the last execution in California took place in 2006 (Death Penalty Information Center). This shows that California is bearing costs that they should not as they maintain death row prisons full of people who do not appear to have execution in their future. Putting this practice on trial means the plaintiff (for the sake of argument, setting the lowest burden of proof would be a civil case, meaning there is a plaintiff) must provide evidence for why the death penalty is good, but here the argument for it is significantly outweighed by the fiscal costs it places on the state. It might be time to look at the perspective of if states would choose to implement the death penalty today. Taking this view would seriously change the discussion of the death penalty when the burden of proof would fall on proving that the death penalty should exist rather than trying to prove that it should not. There may be no perfect answer, as it appears deterrence may exist in some states, but there is value in trying to eliminate excessive spending where it is needed. The main conversation around the death penalty remains its morality, which would add another 20 pages to this paper, but the best way to fully understand an issue is to understand all the areas it touches, so looking at the economics of the death penalty is one piece of a much larger puzzle.

Bibliography

Appleton, Catherine, and Bent Grover. "The Pros and Cons of Life Without Parole." *British Journal of Criminology*, vol. 47, no. 4, July 2007, pp. 597–616. HeinOnline.

Berry, William W., III. "Life-with-Hope Sentencing." *Ohio State Law Journal*, vol. 76, no. 5, 2015, pp. 1051–86. HeinOnline.

"Chapter 6, § 620: Appointment of Counsel in Capital Cases." *United States Courts*, www.uscourts.gov/rules-policies/judiciary-policies/cja-guidelines/chapter-6-ss-620-appointment-counsel-capital-cases#:~:text=(a)%20As%20required%20by%2018,defense%20of%20death%20penalty%20cases. Accessed 6 Dec. 2024.

Cook, Philip J. Oxford University Press on Behalf of the American Law and Economics Association , Durham, NC, 2009, *Potential Savings from Abolition of the Death Penalty in North Carolina*.

"Death Penalty Appeals Process: Capital Punishment in Context." *Death Penalty Appeals Process | Capital Punishment in Context*, capitalpunishmentincontext.org/resources/dpappealsprocess. Accessed 6 Dec. 2024.

Dezhbakhsh, Hashem, Paul H. Rubin, and Joanna M. Shepherd. "Does Capital Punishment Have a Deterrent Effect? - New Evidence from Postmoratorium Panel Data." *American Law and Economics Review* , vol. 5, no. 2, Fall 2003, pp. 344–76. HeinOnline.

"Execution Database." *Death Penalty Information Center*, https://deathpenaltyinfo.org/database/executions

"Executions by State by Year." *Death Penalty Information Center*, https://deathpenaltyinfo.org/executions/executions-overview/executions-by-state-and-year

FEDERAL PRISON SYSTEM PER CAPITA COSTS FY 2017 , www.bop.gov/foia/docs/fy2017_per_capita_costs.pdf. Accessed 16 Dec. 2024.

Gould, Jon B., and Lisa Greenman. 2010, *Report to the Committee on Defender Services Judicial Conference of the United States Update on the Cost and Quality of Defense Representation in Federal Death Penalty Cases*.

Indiana, Office of Fiscal and Management Analysis, and Sen. Boots. *Fiscal Impact Statement*, SB 43, 2010.

"Innocence Database." *Death Penalty Information Center*, deathpenaltyinfo.org/database/innocence. Accessed Dec. 2024.

Jolls, Christine, "On Law Enforcement with Boundedly Rational Actors" (2004). Harvard Law School John M. Olin Center for Law, Economics and Business Discussion Paper Series. Paper 494. http://lsr.nellco.org/harvard_olin/494

Johnson, Robert, and Sandra McGunigall-Smith. “Life without Parole, America's Other Death Penalty: Notes on Life under Sentence of Death by Incarceration.” *Prison Journal*, vol. 88, June 2008, pp. 328–46.

Kansas, Death Penalty Advisory Committee. *Report of the Judicial Council* , 2014.

Kaplan, Aliza B., et al. Seattle , OR, 2016, *Oregon’s Death Penalty: A Cost Analysis*.

Knoll, Melissa A. Z. "The Role of Behavioral Economics and Behavioral Decision Making in Americans' Retirement Savings Decisions." Social Security Bulletin, vol. 70, no. 4, 2010, pp. 1-24. HeinOnline

Kovandzic, Tomislav V., Lynne M. Vieraitis, and Denise Paquette Boots. “Does the death penalty save lives? New evidence from state panel data, 1977 to 2006.” *Criminology & Public Policy*, vol. 8, no. 4, Nov. 2009, pp. 803-39.

Leigey, Margaret E., and Doris Schartmueller. “The Fiscal and Human Costs of Life Without Parole.” *Prison Journal: An International Forum on Incarceration and Alternative Sanctions*, vol. 99, no. 2, Mar. 2019, pp. 241–62. HeinOnline.

“Life without Parole.” *Death Penalty Information Center*, deathpenaltyinfo.org/policy-issues/sentencing-alternatives/life-without-parole. Accessed 6 Dec. 2024.

“A Matter of Life and Death: The Effect of Life-without-Parole Statutes on Capital Punishment.” *Harvard Law Review*, vol. 119, no. 6, Apr. 2006, pp. 1838-54. HeinOnline.

McFarland, Torin. “The Death Penalty vs. Life Incarceration: A Financial Analysis.” *Susquehanna University Political Review*, vol. 7, no. 4, Apr. 2016.

Petersilia, Joan. *Understanding California Corrections*. Oakland, CA: The Public Policy Institute of California, 2006.

Richard H. McAdams & Thomas S. Ulen, "Behavioral Criminal Law and Economics" (John M. Olin Program in Law and Economics Working Paper No. 440, 2008).

Roman, John, et al. *The Cost of the Death Penalty in Maryland*. Washington, DC: Urban Institute Justice Policy Center, Mar. 2008.

Shepherd, Joanna M. “Deterrence versus Brutalization: Capital Punishment's Differing Impacts Among States.” *Michigan Law Review*, vol. 104, no. 2, 2005, pp. 203–49.

"State By State." *Death Penalty Information Center*, https://deathpenaltyinfo.org/states-landing

"State Studies on Monetary Costs." *Death Penalty Information Center*, https://deathpenaltyinfo.org/policy-issues/costs/summary-of-states-death-penalty

"State Summaries." *Death Penalty Information Center*, https://deathpenaltyinfo.org/curriculum/high-school/state-by-state-data/state-summaries

Steele, Tracey, and Norma Wilcox. "A View From the Inside: The Role of Redemption, Deterrence, and Masculinity on Inmate Support for the Death Penalty." *Crime and Delinquency*, vol. 49, no. 2, Apr. 2003, pp. 285–312. HeinOnline.

Tennessee's Death Penalty: Costs and Consequences. Nashville, TN: Tennessee Comptroller of the Treasury, July 2004.

The Criminal Defense Team. "Death Row in the United States: A Statistical Analysis [2024]." *The Criminal Defense Team Baldwin Perry & Wiley P.C.*, 9 Oct. 2024, www.criminaldefenseteam.com/death-row-in-the-united-states/.

"The Death Penalty in 2023: Year End Report." *Death Penalty Information Center*, deathpenaltyinfo.org/research/analysis/reports/year-end-reports/the-death-penalty-in-2023-year-end-report#new-death-sentences. Accessed 6 Dec. 2024.

United States, Congress, Office of Justice Programs , et al. *Prisoners in 2022 - Statistical Tables*, U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics, 2024.

United States, Office of Justice Programs , and Tracy L. Snell. *Capital Punishment, 2020--Statistical Tables*, NCJ 302729. *Statistical Tables*. United States, Congress, Interim Finance Committee, et al. *State of Nevada Performance Audit, Fiscal Costs of the Death Penalty*, Legislative Counsel Bureau, 2014.

United States, Congress, Office of Justice Programs , and Danielle Kaeble. *Time Served in State Prison, 2018*, U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics, 2021.

Utah Death Penalty: A Review of the Data. Salt Lake City, UT: Utah Commission on Criminal and Juvenile Justice, June 2018.

Washington Death Penalty Cost Analysis. Salem, OR: Oregon Criminal Justice Commission, Feb. 2017.

"Wasteful & Inefficient." *Equal Justice USA*, 28 Jan. 2020, ejusa.org/resource/wasteful-inefficient/.

"What Is the Cost?" *SCADP*, www.scadp.org/what-is-the-cost. Accessed 6 Dec. 2024.