

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Applying Google Trends Data to Questions of Gender-Based Discrimination and Violence

Permalink

<https://escholarship.org/uc/item/2696978d>

Author

Maloney, Elizabeth Maureen

Publication Date

2022

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Applying Google Trends Data to Questions of Gender-Based Discrimination and Violence

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Elizabeth Maloney

Dissertation Committee:
Professor David Neumark, Chair
Professor Emily Owens
Associate Professor Yingying Dong

2022

DEDICATION

I dedicate this work to my parents, Richard and Patsy Maloney, who have always encouraged me to observe and ask questions about the world around me.

TABLE OF CONTENTS

	Page
LIST OF FIGURES	v
LIST OF TABLES	vii
LIST OF ALGORITHMS	ix
ACKNOWLEDGMENTS	x
VITA	xi
ABSTRACT OF THE DISSERTATION	xiii
1 The Gender Wage Gap: A Product of Misogyny, Not Just Gender Norms	1
1.1 Introduction	1
1.2 Background	3
1.3 Models of Discrimination	4
1.3.1 The Becker Model of Taste-Based Discrimination	5
1.3.2 Black’s Search Model of Discrimination	7
1.4 Testing Model Predictions on the Gender Wage Gap	8
1.5 Measuring Sexism Against Women	12
1.5.1 Normative Sexism	26
1.5.2 Measures of Sexism and Other Labor Market Measures	33
1.6 Results	34
1.6.1 Measures of Sexism and Labor Market Outcomes	34
1.6.2 Testing Models of Discrimination	38
1.7 Conclusion	41
2 The Effects of Sheltering-in-Place on Domestic Violence Survivor Response Strategies	43
2.1 Introduction	43
2.2 Mechanisms Influencing Domestic Violence	48
2.3 Data and Methods	52
2.4 Results	63
2.5 Some Checks	87
2.6 Conclusion	99

2.7 Mobility Reduction Timelines in Each DMA 100

LIST OF FIGURES

	Page
1.1 Reactionary Search Terms Over Time	19
1.2 Derogatory Search Terms Over Time	20
1.3 Manosphere Search Terms Over Time	20
1.4 Violent Search Terms Over Time	21
1.5 Index of Misogyny by State	24
1.6 Correlation Between Various Measures	25
1.7 Average Normative Sexism Index Over Time by Census Division	28
1.8 Responses to GSS Prejudice Questions Over Time	28
1.9 Normative Sexism by State	29
1.10 Misogyny vs. Normative Sexism Indices	31
1.11 Normative Sexism - Men vs. Women	32
2.1 Timing of Mobility Drops	56
2.2 Mobility by DMA	57
2.3 Week of Lockdown by State	57
2.4 Average Percent Difference in Mobility by State	58
2.5 Search Interest in COVID-19/Coronavirus and Mobility Over Time	61
2.6 Search Interest in Domestic Violence Hotlines and Mobility Over Time	65
2.7 Search Interest in Domestic Violence Help and Mobility Over Time	66
2.8 Search Interest in Shelters and Mobility Over Time	66
2.9 Search Interest in Restraining/Protective Orders and Mobility Over Time	67
2.10 Search Interest in Child Protective Services/Child Abuse and Mobility Over Time	67
2.11 Event Study - Search Interest in Domestic Violence Hotlines	68
2.12 Event Study - Search Interest in Domestic Violence Help	69
2.13 Event Study - Search Interest in Shelters	69
2.14 Event Study - Search Interest in Protective/Restraining Orders.	70
2.15 Event Study - Search Interest in Child Protective Services/Child Abuse	70
2.16 Average Percent Change in Mobility Quartile in Each State	72
2.17 Average Quartile of Post Period Covid Search Interest in Each State	74
2.18 2020 State Average Household Size Quartiles	77
2.19 Pre-Period Female-to-Male State Employment Ratio Quartiles	81
2.20 Quartiles of the Percent Difference in State Unemployment Insurance Claims Rates	84

2.21	Quartiles of Pre-Period Unemployment Insurance Claims Rates	86
2.22	Event Study - Search Interest in Toads	94
2.23	Event Study - Search Interest in Toothaches	94
2.24	Average Quartile within a State of Percent of 2020 Households with 1 Person Living Alone	96
2.25	Average Quartile within a State of Percent of 2020 Households that Include a Couple Living Together	97
2.26	Mobility Reductions by DMA - Part 1	101
2.27	Mobility Reductions by DMA - Part 2	102
2.28	Mobility Reductions by DMA - Part 3	103
2.29	Mobility Reductions by DMA - Part 4	104
2.30	Mobility Reductions by DMA - Part 5	105
2.31	Mobility Reductions by DMA - Part 6	106
2.32	Mobility Reductions by DMA - Part 7	107
2.33	Mobility Reductions by DMA - Part 8	108
2.34	Mobility Reductions by DMA - Part 9	109
2.35	Mobility Reductions by DMA - Part 10	110

LIST OF TABLES

	Page
1.1 GSS Questions on Attitudes Toward Women	10
1.2 Sexist Search Terms	14
1.3 Top 20 Queries and Relative Interest - Derogatory Search Terms	15
1.4 Top 20 Queries and Relative Interest - Violent Search Terms	16
1.5 Top 20 Queries and Relative Interest - Reactionary Search Terms	17
1.6 Top 20 Queries and Relative Interest - Manosphere Words	18
1.7 GSS Responses by Census Division	30
1.8 Relationship between Sexism Indices and Labor Market Outcomes (Part 1) .	35
1.9 Relationship between Sexism Indices and Labor Market Outcomes (Part 2) .	36
1.10 Testing the Predictions of Black and Becker Models Using the Index of Misogyny	39
1.11 Testing the Predictions of Black and Becker Models Using the Index of Nor- mative Sexism	40
2.1 Top Related Queries - Outcomes of Interest	53
2.2 Top Related Queries to COVID-19	60
2.3 Effect of Staying-at-Home on Domestic Violence Related Search Behavior . .	65
2.4 Effects of Sheltering-In-Place by Percent Difference in Mobility Quartiles . .	73
2.5 Effects of Sheltering-In-Place by Post Period COVID Search Interest Quartiles	75
2.6 Effects of Sheltering-In-Place by Quartiles of Average 2020 Household Size .	76
2.7 Relationship Between Domestic Violence Related Search Interest and Potent- tial Mechanisms Directly	78
2.8 Effects of Sheltering-In-Place by Quartiles of Pre-Period Female-to-Male Em- ployment Ratios	80
2.9 Effects of Sheltering-In-Place by Quartiles of the Percent Difference in Unem- ployment Insurance Claims Rates	83
2.10 Effects of Sheltering-In-Place by Quartiles of the Pre-Period Unemployment Insurance Claims Rate	85
2.11 Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Year-to-Year Comparisons	88
2.12 Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Excluding 2019	89
2.13 Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Including a Time Trend	90

2.14	Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Clustering at DMA Level	91
2.15	Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Full Year	92
2.16	Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Placebo Variables	93
2.17	Effects of Sheltering-In-Place by Quartiles of the Percent Living Alone in 2020	95
2.18	Effects of Sheltering-In-Place by Quartiles of the Percent Living with a Partner in 2020	98

LIST OF ALGORITHMS

	Page
1 Detecting Timeline of Reductions in Mobility	55

ACKNOWLEDGMENTS

I am incredibly grateful to Professor David Neumark for his guidance and support over the past four years. I could not have asked for a better advisor and mentor. I would also like to thank Professors Emily Owens and Yingying Dong for their insightful comments about this work.

Thanks to my sister and brother-in-law, Katherine and Chris Pham, who encouraged me to take my first economics course when I realized English was not the major for me, and to the professors at Pacific Lutheran University who fostered my love of economics and inspired me to keep studying it, especially Professor Lynn Hunnicutt, Professor Karen Travis, and Professor Norris Peterson.

I would also like to thank my husband, Alex Parret, who has been a constant source of support and encouragement; my friend and colleague Alexandra Marr, who kept me sane during our daily walks and Friday evening trips to Eureka; and the rest of my friends and family who have championed me along the way. I cannot thank you enough.

VITA

Elizabeth Maloney

EDUCATION

- Doctor of Philosophy in Economics** **2022**
University of California, Irvine *Irvine, California*
- Bachelor of Arts in Economics and Computer Science** **2016**
Pacific Lutheran University *Tacoma, Washington*
- Bachelor of Science in Mathematics with a Minor in Statistics** **2016**
Pacific Lutheran University *Tacoma, Washington*

PUBLICATIONS

Maloney, Katherine N., Ryan T. Botts, Taylor S. Davis, Bethany K. Okada, Elizabeth M. Maloney, Christopher A. Leber, Oscar Alvarado et al. “Cryptic Species Account for the Seemingly Idiosyncratic Secondary Metabolism of Sarcophyton glaucum Specimens Collected in Palau.” *Journal of Natural Products* 83, no. 3 (2020): 693-705.

RESEARCH EXPERIENCE

- Graduate Student Researcher** **2021**
University of California, Irvine *Irvine, California*
- Summer Institute on Corporate Welfare** **2018**
University of California, Irvine *Irvine, California*

TEACHING EXPERIENCE

- Introductory Statistics Course Instructor** **2020**
University of California, Irvine *Irvine, California*
- Teaching Assistant** **2016–2021**
University of California, Irvine *Irvine, California*

ADDITIONAL EXPERIENCE

Economic Consulting
With David Neumark

2019-Present
Irvine, California

FELLOWSHIPS

Merit Fellowship

2022

Dissertation Completion Fellowship

2021

Associate Dean Fellowship

2020

Summer Research Fellowship

2019, 2020

ABSTRACT OF THE DISSERTATION

Applying Google Trends Data to Questions of Gender-Based Discrimination and Violence

By

Elizabeth Maloney

Doctor of Philosophy in Economics

University of California, Irvine, 2022

Professor David Neumark, Chair

In my first chapter, I construct a novel measure of misogyny using Google Trends data on searches that include derogatory terms for women. I show that misogyny is an economically meaningful and statistically significant predictor of the wage gap, and use it to test the predictions of two influential labor market discrimination models. I find that the gender wage gap is inconsistent with the Becker model of taste-based discrimination, but that it fits Black's search model of discrimination which allows for discrimination from even a small group of misogynists to result in a wage gap. In my second chapter, I investigate the impact of COVID-19 lockdowns on domestic violence by leveraging Google Trends data on search interest in domestic violence resources. I find that COVID-19 lockdown orders result in a meaningful and statistically significant increase in search interest in domestic violence hotlines but a decrease in search interest for other resources such as protective orders and domestic violence shelters, which require victims to leave their homes in the midst of a global pandemic. Moreover, I find that increased exposure between victim and perpetrator exacerbate domestic violence and that economic hardship and a lack of independent earning potential make women especially vulnerable. Finally, I find that search interest in terms likely to be searched by third-party reporters drastically decrease as a result of COVID-19 lockdowns, suggesting that external mechanisms for identifying and reporting domestic violence are affected by lockdown orders.

Chapter 1

The Gender Wage Gap: A Product of Misogyny, Not Just Gender Norms

1.1 Introduction

The wage gap between men and women in the United States is long established, and cannot be fully explained by observable wage determinants such as education and experience. Much of the research on the subject attribute all or part of the unexplained portion of the wage gap to discrimination, building off of the Becker model of discrimination (Becker, 1957). Others argue that that discrimination as described by Becker cannot explain the persistence of the wage gap under competitive market conditions, posing instead models of statistical discrimination (Arrow, 1973; Phelps, 1972) or search models of discrimination (Black, 1995), which can allow for a persistent wage gap. But the question of how well these discrimination models actually fit in the gender context requires empirical validation. Becker developed his model of wage discrimination in the context of race, allowing employer utility to depend on their level of distaste toward employing black workers. In the context of racial discrim-

ination, where slavery and segregation are part of our not-so-distant past, the presence of true distaste along racial lines and its relevance to the black-white wage gap is clear and has been empirically validated. Charles and Guryan lay out several testable implications of Becker's model and evaluate how well these fit the data on the wage gap between black and white workers, finding evidence of Becker's predictions. However, the role of distaste for women with regard to the gender wage gap is less obvious. While true distaste for women undoubtedly exists, it is most obvious among fringe groups. The majority of men are romantically entwined with women and men and women have always been joined by reproductive necessity. Accordingly, men and women have never been truly segregated from one another and in fact, have had to coordinate with one another at least at the familial level. This type of interaction and cooperation make it hard to believe that the gender wage gap is driven by employers who receive disutility from interacting with women, as characterized by models of taste-based discrimination. These employers interact with their mothers, wives, and sisters, so it is not obvious why workplace interactions would be substantively different.

In their 2018 working paper, Charles et al. begin to answer this question, but they measure gender prejudice using only General Social Survey data on gender norms. While gender norms likely play an important role in understanding the wage gap, it is not clear that it plays the same role as distaste for women as described in the Becker model. To measure the role of distaste for women (or misogyny) in determining the wage gap, I construct a novel index of misogyny based on Google Trends data on sexist search terms. I find that the gender wage gap is increasing in misogyny and that this relationship is statistically significant, even after accounting for education and experience and controlling for strength of gender norms. Misogyny is also strongly related to other labor market outcomes including the proportion of women who remain unmarried between the ages of 20 and 40, the average age at which women have their first child, labor force participation, and college completion. Because I estimate the gender wage gap conditional on observables such as labor force participation

and education, the fact that misogyny is an important predictor of these things means that the role of misogyny on the gender wage gap is likely underestimated.

I use this measure to test the empirical implications of both the Becker model of discrimination and Black's sequential search framework, and find that the data is more consistent with Black's search model of discrimination. Under this framework, a small minority contingent of misogynistic employers can create a wage gap even if they do not hire women, which could explain why misogyny may be an important determinant of the wage gap despite the fact that most men do not exhibit an aversion to women.

1.2 Background

Using data from the Panel Study of Income Dynamics since the 1950's, Blau and Kahn (2017) document the evolution of the wage gap over time. Prior to 1980, they show that the ratio of womens' to mens' wages hovers around .6. During the 1980's and 1990's the ratio improved until it was just under .8. In recent years however, progress has stagnated.

Efforts to explain the gender wage gap have largely begun by exploring the difference between the observable characteristics of men and women, including education and work experience. Blau and Kahn (2017) note that the education gap between men and women has actually reversed. Since 2011, women have had higher levels of education on average than men. Women have also made substantial gains with regard to labor market experience. In 2011, the gap between the labor market experience of men and women was only 1.4 years despite having been almost seven years in 1981. Experience gaps may reflect differences in use of time outside the workplace (Becker, 1985; Mincer & Polachek, 1974). Time-use surveys show that women participate in non-market labor like household chores and child care responsibilities

at higher rates than men (Bianchi et al., 2000). Faced with the burden of responsibilities at home, women may allocate effort away from the market and toward household labor.

Even accounting for differences in human capital accumulation and other observable differences between men and women, a wage gap persists. Drawing from the Becker model, this residual wage gap is the portion we consider to be potentially due to discrimination. Of course, I cannot assume this residual wage gap entirely reflects discrimination. Unobserved differences between men and women that contribute to productivity will also factor into observed wages. If these unobserved differences, on average, reflect higher productivity among men relative to women, the residual wage gap will overstate the effect of discrimination on wages. Conversely, as Blau and Kahn (2007) point out, discrimination could influence a woman's rate of human capital accumulation as well as her choice of occupation. If this is true, then examining the residual wage gap after accounting for these factors may understate the role of discrimination. For example, if women are pushed into lower paying fields due to discrimination, the residual wage gap after accounting for occupational choice will underestimate the portion of the wage gap caused by discrimination. A widening distribution of the returns to different fields exacerbates this issue as the returns to male-dominated fields advance relatively more quickly and women are left "swimming upstream" (Blau & Kahn, 2007). Despite these drawbacks, the use of residual wage gaps is common in the literature. The Becker model of discrimination provides the theoretical underpinning of a discrimination-based wage gap.

1.3 Models of Discrimination

In this paper, I focus mainly on taste-based models of discrimination to explain the residual wage gap. Alternatively, Arrow (1973) and Phelps (1972) offer models of statistical discrimination in which employers try to infer the productivity of applicants based on a noisy signal

of productivity (e.g. education) as well as beliefs about the average productivity of men and women. Aigner and Cain (1977) point out that if employers correctly identify the average productivity of men and women (which we would expect them to do in the long run), this practice does not constitute economic discrimination against a group in the typical sense since expected productivity equals true productivity on average. But this does not mean equally productive men and women will be paid the same. Because of the weight employers place on the group average, the expected productivity of a woman will be less than that of an equally productive man, and she will be paid accordingly. Altonji and Pierret (2001) point out that under statistical discrimination, we would expect residual wage gaps to decline with age and tenure, as employers are able to better observe true productivity and rely less on beliefs about the average productivity of women. Sin et al. (2020) find no evidence of such declines for the male-female wage gap in the New Zealand labor market and Munasinghe et al. (2008) find that returns to experience actually tend to decrease for women relative to men in the US labor market. Even if statistical discrimination does play a role in the wage gap between men and women, it should have less impact on the residual wage gap (after controlling for experience), which we use here, relative to an unconditional wage gap. Thus, I restrict my focus to taste-based models of discrimination and consider two strains of such models.

1.3.1 The Becker Model of Taste-Based Discrimination

Consider first the taste-based discrimination model described by Becker (1957) with a perfectly competitive labor market and employers who are prejudiced against a less privileged group, which I denote with the letter F. I use M to denote the more privileged group which does not experience discrimination. Employer utility V_i increases with profit π_i and decreases with the interaction between degree of prejudice d and the number of F-types he employs.

That is,

$$V_i = \pi_i - d_i F_i.$$

Assuming F- and M-types are perfect substitutes, $\pi_i = f(F + W) - w_M M - w_F F$ where w_M and w_F are the wages for M and F-types, respectively. The utility maximizing choice of M- and F-types at an interior solution is given by

$$\begin{aligned} f' - w_M &= 0 \\ f' - w_F - d_i &= 0. \end{aligned}$$

The short-run equilibrium involves a completely segregated labor force, with F-types working for the least prejudiced employers and M-types working for the most prejudiced employers. Assuming the distribution of prejudice is fairly smooth, there must exist an employer who is indifferent between hiring F- and M-types. I refer to this employer as the “marginal discriminator” and his degree of prejudice as the “marginal” level of prejudice d^* . This employer hires at equilibrium wages

$$w_M = w_F + d^*.$$

More prejudiced employers hire only M-types at wage $w_M = f'$ and less prejudiced employers hire only F-types at wage $w_F = w_m - d^*$. As a result of this labor market segregation, the equilibrium wage is entirely determined by the marginal level of prejudice, whereas the average level of prejudice will not directly affect the wage gap. To see why this is the case, suppose average prejudice is very high because there is a lot of prejudice at the top of the prejudice distribution. Since these employers do not hire F-type employees, their prejudice is irrelevant to the wage gap. If the employers who actually hire F-types are unprejudiced, there will still be no wage gap. Thus, we expect the relationship between the marginal level of prejudice and the wage gap to be stronger than that of the average level of prejudice.

This also means that prejudice increases will affect the wage gap differently depending on where they come from in the prejudice distribution. For example, suppose prejudice increases at the top of the prejudice distribution. Any increase above the marginal prejudice level should have no effect on the wage gap. If, on the other hand, prejudice increases at the bottom of the prejudice distribution, this is likely to push up the prejudice level of the marginal discriminator, and the wage gap will widen.

Now consider the impact of increasing the number of F-types in the labor market. All else equal, as additional F-types enter, they must work for increasingly prejudiced employers, pushing up the marginal prejudice level, and increasing the wage gap. Thus, the wage gap will increase as the proportion of F-types in the labor force increases. Note that under this framework, a small contingent of very prejudiced employers should not contribute to the wage gap as long as the number of F-types is not so high that they are forced to work for these very prejudiced employers.

1.3.2 Black's Search Model of Discrimination

Black (1995) offers an alternative model of discrimination that incorporates search frictions. In this model, M-types (those not discriminated against) and F-types (who experience discrimination) enter the labor market and search sequentially for a suitable job. In each search period, applicants pay a fixed cost and are matched with a potential employer who chooses whether or not to hire him or her. Prejudiced employers will not hire F-type applicants, which means that on average, F-types must search longer for jobs. Because search costs are higher for this group, their reservation wage is lower and unprejudiced employers can offer them a lower wage relative to M-types. Note that in contrast with the Becker model, a small contingent of very prejudiced employers do contribute to wage gaps under this model, as

they represent higher average search costs for F-types despite the fact that they do not hire them.

In this model, the key drivers of the wage gap are the share of employers that are prejudiced and the share of the workforce made up by F-types. As the share of prejudiced employers increases, F-type applicants face longer searches, on average, and the power of unprejudiced firms increase, resulting in even lower wages for F-type workers. Now consider the impact of additional F-types in the labor force. As more F-types enter the work force, unprejudiced firms are able to hire more low-cost (F-type) employees and the profits of unprejudiced firms increase relative to those of prejudiced firms. As unprejudiced firms become more profitable, some prejudiced firms are driven from the market, reducing the share of prejudiced firms. This reduces search costs for F-types and increases their reservation wage. Now, unprejudiced firms must offer them higher wages to induce them to accept a job offer, and the wage gap shrinks. Note that this prediction about the relationship between the wage gap and the proportion of women in the labor market is the reverse of Becker's, allowing us to sharply distinguish between the empirical predictions of the two models.

In this framework, employers are either prejudiced or not and intensity of discrimination is not a relevant factor. Despite this, we can still expect a relationship between the average level of prejudice and the wage gap since average prejudice increases systematically with the proportion of prejudiced firms.

1.4 Testing Model Predictions on the Gender Wage Gap

I test whether or not the gender wage gap is consistent with either model by examining cross-state variation in the residual gender wage gap, the proportion of women in the labor

market, the proportion of prejudiced employers in the labor market, and various parts of the prejudice distribution using a multi-stage regression model.

In the first stage, I regress log wages on education, a quadratic in potential experience, and gender-specific year and state dummies using data from the CPS May Outgoing Rotation Group between 1979 and 2019, weighted by the census earnings weight. Specifically, I use the cleaned and standardized version of this data provided by the Center for Economic and Policy Research. The coefficients on the state-by-female indicators give the residual wage gap in each state.

To estimate the percent of prejudiced employers, I assume employer prejudice mirrors that of the general population, which I estimate with the proportion of sexist responses to GSS questions on attitudes toward women. These questions are listed in Table 1.1. I assign a numeric scale to responses so that they are centered at zero and increasing in prejudice level. Thus, prejudiced answers are greater than zero and non-prejudiced answers are less than zero. I then average across all questions and compute the percent of this average that is greater than zero within each state. That is, I take the percent of respondents in each state whose responses are sexist, on average.

For now, suppose I have an index of individual prejudice that varies within each state (I will detail exactly how I construct this index in the next section). I can take the 10th and 90th percentiles of the prejudice distribution as measures of its lower and upper tails, respectively. Following Charles and Guryan (2008), I approximate the marginal prejudice level by computing the p^{th} percentile of the prejudice distribution where p is the percent of women in a state's workforce. This reflects Becker's definition of the marginal discriminator as the least prejudiced employer who hires women. Finally, I take the mean or average level of prejudice in the state. I also compute the proportion of women in the labor force in each state, averaging over time.

Table 1.1: GSS Questions on Attitudes Toward Women

fework:	<i>Married Women Shouldn't Work</i> Do you approve or disapprove of a married woman earning money in business or industry if she has a husband capable of supporting her?
fehome:	<i>Women Should Run Homes</i> Do you agree or disagree with this statement? Women should take care of running their homes and leave running the country up to men.
fepol:	<i>Men Better for Politics</i> Tell me if you agree or disagree with this statement: Most men are better suited emotionally for politics than are most women.
fepres:	<i>Would Not Vote for Woman President</i> If your party nominated a woman for President, would you vote for her if she were qualified for the job?
fechld:	<i>Working Mothers Worse Relationship</i> A working mother can establish just as warm and secure a relationship with her children as a mother who does not work.
fepresch:	<i>Preschoolers Suffer When Mom Works</i> A preschool child is likely to suffer if his or her mother works.
fehhelp:	<i>Husband's Career More Important than Wife's</i> It is more important for a wife to help her husband's career than to have one herself.
fefam:	<i>Wife Takes Care of Home and Family</i> It is much better for everyone involved if the man is the achiever outside the home and the women takes care of the home and family.

This table describes questions about gender norms from the GSS. The left hand column gives the variable name used by the GSS. The italicized text in the right hand column is the summary of the question topic I use to identify the question in tables and figures throughout this paper. The non-italicized text in the right hand column is the question as phrased in the GSS.

I then use my estimates of the residual wage gap as the dependent variable in a second set of regressions on the proportion of women in the labor market, the proportion of prejudiced employers in the labor market, and various points of the prejudice distribution (mean, marginal, and each tail), weighted by the precision at which I estimate the residual wage gap in the first step. I can then evaluate the results to see if they are consistent with the predictions of either the Becker or Black model.

Under the Becker model, we expect the wage gap to be increasing in the proportion of women in the labor force. I can also test for evidence of a marginal discriminator by examining which points of the prejudice distribution influence the wage gap more strongly. If the marginal discriminator determines the wage gap, as predicted by the Becker model, we expect the wage gap to be systematically related to the marginal, rather than the average level of prejudice. We would also expect prejudice at the bottom of the prejudice distribution to be a better predictor of the wage gap than prejudice at the top of the distribution. That is, we expect the 10th percentile of the prejudice distribution to be more strongly correlated with the residual wage gap than the 90th percentile. Recall that given the level of gender integration in the labor market, a lack of evidence for the relevance of a marginal discriminator would not be surprising. If the wage gap was increasing in the proportion of women in the labor force, but there was no evidence of a marginal discriminator, this would still be consistent with the modified Becker model proposed by Neumark (1988), which allows employer utility to depend on the relative proportion of women in the labor force and predicts an integrated labor force with no marginal discriminator.

Under Black's model on the other hand, I expect the wage gap to be decreasing, rather than increasing, in the proportion of the labor force made up by women. I also expect the wage gap to be increasing in the proportion of prejudiced employers. This model does not yield predictions about various points of the prejudice distribution as the wage gap here does not

rely on a marginal discriminator, but we still expect the wage gap to be increasing in average prejudice levels.

Other work has demonstrated the relevance of both models to other wage gaps. Charles and Guryan (2008) use a similar procedure to empirically validate the Becker model’s predictions in the context of the wage gap between black and white workers. They find that the black-white wage gap is increasing in the proportion of black individuals in the labor force, that the marginal level of prejudice is more strongly related to the wage gap than the average prejudice level, and that the lower tail of the prejudice distribution is a strong predictor of the wage gap whereas the upper tail is not. They conclude that the Becker model describes the racial wage gap well. Burn (2019) on the other hand, performs some of these tests on the wage gap between gay and straight men and finds evidence that Black’s search model is more consistent with data on the gay-straight wage gap.

1.5 Measuring Sexism Against Women

The tests outlined in the previous section require a measure of distaste for women. To construct such a measure, I use Google Trend data on search interest for sexist terms. This yields some important advantages over survey-based measures of sexism. Survey-based estimates of prejudice receive criticism because of social censoring that may bias estimates of prejudice downward, especially in areas where it is less socially acceptable to express prejudiced views (Berinsky, 1999). In his paper on the role of racial prejudice on vote shares in either of Obama’s presidential runs, Stephens-Davidowitz (2014) uses Google Trends data on searches for “the n-word” as a proxy for racial prejudice. He finds that his proxy strongly correlates with GSS measures of racial prejudice, but argues that it improves upon survey-based measures because it is not tempered by social attitudes toward prejudice. He points out that the vote-share for John Kerry (a democrat) is negatively correlated with more

prejudiced responses to GSS questions, but that it does not have a statistically significant correlation with his Google search proxy. He argues that in more democratic areas, it is socially unacceptable to admit racial prejudice so measures like the GSS underestimate the degree of racial prejudice in such areas. Compared to other surveys, the GSS is especially vulnerable to this critique since it is conducted via face-to-face interviews. Google searches on the other hand, are conducted anonymously. Moreover, the survey questions about women on the GSS do not really get at distaste for women, but rather at beliefs about gender norms. As such, I construct a measure of gender prejudice using Google Trends data on misogynistic search terms.

To construct such a measure, I need to identify appropriate search terms to include. Following Stephens-Davidowitz (2014), I initially consider slurs against women. This is not without basis in the literature. In a paper aimed at describing misogynistic language used on Reddit’s “Manosphere” (a portion of the internet devoted to promoting misogyny), T. Farrell et al. (2019) identify words such as “bitch,” “cunt,” and “whore” as part of a group including “violent verbs, and slurs that are not immediately racist or homophobic.” They find that this group of words is rampant across the manosphere, and is actually the most prominent of all categories they examine. I use these three words as well as a few synonyms for them to construct my index of misogyny. I refer to this group of terms as “Derogatory” throughout this section. I also collect data on search interest for three other categories of words which may also be informative, if not as easily interpreted. The categories are “Violent,” “Manosphere,” and “Reactionary.” The “Violent” search terms describe violent sexual acts that are often (though not always) perpetrated against women (e.g. rape). The “Manosphere” group includes terms often used within the manosphere (e.g. Men’s rights). Finally, the “Reactionary” terms are those one might search in response to a negative incident and include the words “sexual harassment” and “workplace harassment.” Table 1.2 lists all words in each category. Note that Google censors search interest when it falls below some threshold of minimum interest. We exclude search terms that are censored too frequently.

Table 1.2: Sexist Search Terms

Derogatory	Manosphere
bitch	misandry
cunt	men's rights
whore	sjw
thot	
twat	
Violent	Reactionary
rape	sexual harassment
gangbang	workplace harassment

This table gives the search terms, by category, used to explore potential measures of sexism based on Google Trends data.

For example, “femoid” was considered for inclusion in the group of “Manosphere” words, but was excluded because it is censored by Google in all but a handful of states.

Tables 1.3, 1.4, 1.5 and 1.6 list the top 20 related search queries for each of the terms given in Table 1.2. Note that these tables contain a lot of offensive language, but I include them in order to give a sense of the intent behind some of these searches. While some of the related queries are obviously in search of anti-woman content, others pertain to song lyrics or searches for definitions of the terms. This is not overly concerning as noisy search terms have been able to explain regional behavior differences well in other work. For example, Stephens-Davidowitz (2014) shows that variation in search interest for the word “God” strongly predicts regional variation in the the percent of people who believe in God (R^2 of 0.65), despite the fact that the top query related to searches for “god” during the time period he examine is “god of war,” a popular video game released during the period. Interest in “god of war” actually had almost double the search interest of “church of god” during this time period. “Greek god”, “god of war walkthrough”, “oh god”, “egyptian god”, “their eyes were watching god”, “sun god”, and “god bless america,” are also found in the top search queries related to searches for the word “god.”

Table 1.3: Top 20 Queries and Relative Interest - Derogatory Search Terms

<i>Bitch</i>		<i>Cunt</i>		<i>Whore</i>	
Int.	Query	Int.	Query	Int.	Query
100	you bitch	100	fuck	100	slut
49	that bitch	62	definition	51	whore house
45	bitch lyrics	61	cunt definition	44	whore definition
39	bad bitch	57	what is cunt	29	man whore
35	bitch ass	54	cunt mean	27	hoe
26	bitch song	54	what does cunt	26	whore lyrics
22	bitch meme	52	what a cunt	26	whore meaning
22	this bitch	47	what does cunt mean	26	whore meme
21	im a bitch	44	cunt meaning	25	bitch
15	fuck you bitch	44	ass	22	what is a whore
15	bitch nigga	41	what is a cunt	21	whore of babylon
14	lil bitch	33	word cunt	21	boo you whore
13	bitch in spanish	33	anal cunt	18	whore gif
12	bitch face	29	bitch	17	definition of whore
12	son of a bitch	22	definition of cunt	15	whore movie
11	bitch quotes	21	slut	15	attention whore
11	little bitch	19	cunt wars	14	prostitute
10	crazy bitch	18	cunt meme	14	whore mouth
10	bitch gif	18	define cunt	14	whore song
10	fat bitch	16	whore	13	whore in this moment
<i>Thot</i>		<i>Twat</i>			
Int.	Query	Int.	Query		
100	what is thot	100	twat waffle		
74	thot meaning	64	definition twat		
74	what is a thot	61	definition		
69	definition thot	53	twat meaning		
68	thot urban	51	what is twat		
67	definition	45	what is a twat		
67	urban dictionary thot	38	what does twat		
65	she a thot thot	35	what does twat mean		
64	urban dictionary	33	cunt		
57	thot thot lyrics	23	twat urban		
56	thot mean	23	twat define		
56	what does thot	21	twat dictionary		
47	what does thot mean	19	urban dictionary twat		
39	thot meme	16	urban dictionary		
36	begone thot	15	definition of twat		
36	white thot	14	twat meme		
31	thot spot	13	twit		
30	hoe	11	ass		
27	duckie thot	11	twat slang		
26	thots	11	twatt		

This table provides the top 20 related searches that occur in the same search session as each “Derogatory” search term. Interest (Int.) reflects relative search volumes for related queries. Scores of 100 imply that the query is most common and interest in other queries is given relative to it.

Table 1.4: Top 20 Queries and Relative Interest - Violent Search Terms

<i>Rape</i>		<i>Gangbang</i>	
Int.	Query	Int.	Query
100	trump rape	100	gangbang porn
92	gay rape	55	gangbang creampie
89	rape victim	51	black gangbang
81	rape video	49	anal
79	rape videos	49	gangbang anal
78	what is rape	48	wife gangbang
76	statutory	43	teen gangbang
74	gang rape	34	girl gangbang
73	statutory rape	33	bbc gangbang
71	date rape	30	big gangbang
64	rape movie	28	gay gangbang
61	teen rape	27	gangbang sex
56	rape me	25	cum gangbang
46	rape definition	21	interracial gangbang
46	rape statistics	20	gangbang dp
45	ear rape	19	forced gangbang
43	child rape	18	ebony gangbang
42	rape stories	18	gangbang videos
39	real rape	17	free gangbang
37	prison rape	17	gangbang milf

This table provides the top 20 related searches that occur in the same search session as each “Violent” search term. Interest (Int.) reflects relative search volumes for related queries. Scores of 100 imply that the query is most common and interest in other queries is given relative to it.

Table 1.5: Top 20 Queries and Relative Interest - Reactionary Search Terms

<i>Workplace Harassment</i>		<i>Sexual Harassment</i>	
Int.	Query	Int.	Query
100	harassment in workplace	100	sexual harassment definition
89	harassment in the workplace	96	harassment definition
74	workplace sexual harassment	89	sexual harassment workplace
74	sexual harassment	80	sexual harassment training
57	sexual harassment in workplace	78	what is sexual harassment
54	sexual harassment in the workplace	67	sexual harassment in workplace
24	work harassment	60	sexual harassment in the workplace
19	harassment at workplace	53	sexual assault
18	what is harassment	53	sexual harassment law
17	what is workplace harassment	39	sexual harrassment
14	discrimination	39	sexual harassment policy
14	workplace harassment laws	35	sexual harassment news
14	harassment laws	34	sexual harassment laws
14	workplace harassment definition	33	sexual harassment at work
13	harassment definition	33	accused of sexual harassment
13	workplace discrimination	32	definition of sexual harassment
11	harassment at the workplace	31	discrimination
11	what is harassment in the workplace	30	sexual harassment california
10	harassment at work	30	quid pro quo harassment
10	workplace bullying	29	quid pro quo sexual harassment


This table provides the top 20 related searches that occur in the same search session as each “Reactionary” search term. Interest (Int.) reflects relative search volumes for related queries. Scores of 100 imply that the query is most common and interest in other queries is given relative to it.

Table 1.6: Top 20 Queries and Relative Interest - Manosphere Words

<i>Men's Rights</i>		<i>Misandry</i>		<i>SJW</i>	
Int.	Query	Int.	Query	Int.	Query
100	mens rights reddit	100	misogyny	100	sjw meaning
64	mens rights movement	60	feminism	60	what is sjw
36	mens rights activists	48	define misandry	42	sjw meme
36	mens rights activist	34	what is misandry	39	what does sjw
26	mens rea	23	misogynist	31	sjw reddit
25	mra	23	feminist	27	what does sjw mean
		23	misandrism	14	marvel sjw
		20	misogyny definition	14	tumblr sjw
		20	misandry meaning	13	anti sjw
		18	definition of misandry	12	youtube sjw
		13	misogynistic	11	sjw star wars
		13	misandry today	10	what is an sjw
		12	misandry game of thrones	9	sjw hate
		11	misandry pronunciation	9	what is a sjw
		11	misandry meme	9	sjw cringe
		10	feminism definition	8	sjw feminist
		10	misandry gif	8	sjw memes
		10	misandry bubble	8	whats sjw
		10	misandry def	8	sjws
		7	misogyny meaning	7	define sjw

This table provides the top 20 related searches that occur in the same search session as each “Manosphere” search term. Interest (Int.) reflects relative search volumes for related queries. Scores of 100 imply that the query is most common and interest in other queries is given relative to it.

Figure 1.1: Reactionary Search Terms Over Time




This figure gives the Google search interest between 2006 and 2020 for the “Reactionary” search terms. These are normalized by the maximum rate of searches in the time period and scaled by 100.

Figures 1.1, 1.2, 1.3, and 1.4 give interest over time in the words listed in Table 1.2 across the US between 2006 and 2020 for each category. These are normalized by the maximum rate of searches on the topic in a given period and scaled by 100. Interest in the “Derogatory”, “Manosphere”, and “Violent” groups of words have trended downward in recent years.

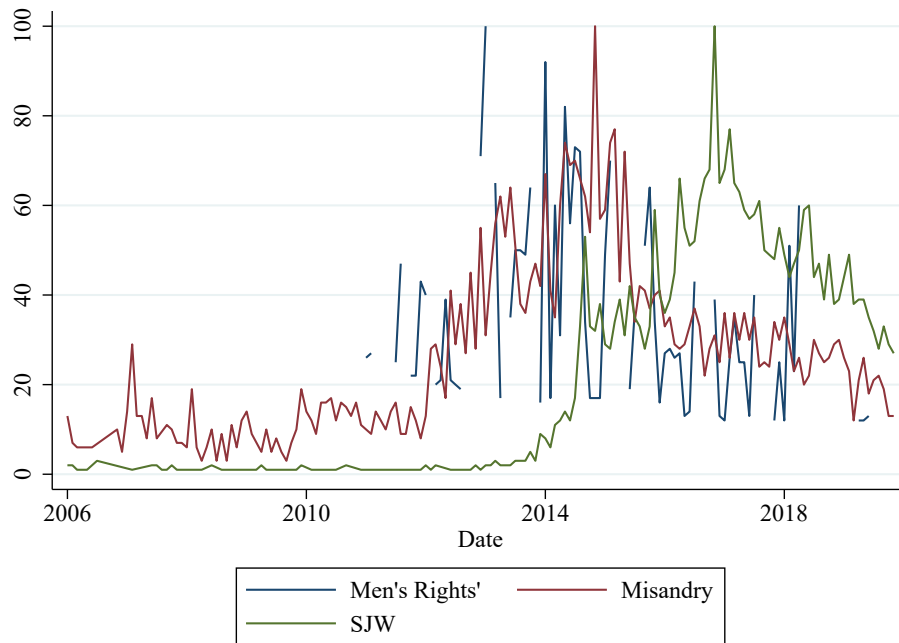
By contrast, search interest in the terms sexual harassment and workplace harassment, seen in Figure 1.1, has been much more stable without an obvious up- or downward trend, but with two notable spikes. The first occurs in search interest for “workplace harassment” in the fall of 2009. This may be due to David Letterman’s October 2009 revelation that he had slept with female members of his staff. The second big spike in both “workplace harassment” and “sexual harassment” occurs toward the end of 2017. This is likely due to the “me too” movement, which exploded in late 2017 after sexual assault allegations against Harvey Weinstein were brought to the public’s attention. These spikes likely reflect the effect of external media buzz rather than personal experience. Moreover, Table 1.5 indicates that

Figure 1.2: Derogatory Search Terms Over Time




This figure gives the Google search interest between 2006 and 2020 for the “Derogatory” search terms. These are normalized by the maximum rate of searches in the time period and scaled by 100.

Figure 1.3: Manosphere Search Terms Over Time



This figure gives the Google search interest between 2006 and 2020 for the “Manosphere” search terms. These are normalized by the maximum rate of searches in the time period and scaled by 100.

Figure 1.4: Violent Search Terms Over Time



This figure gives the Google search interest between 2006 and 2020 for the “Violent” search terms. These are normalized by the maximum rate of searches in the time period and scaled by 100.

much of the search interest in these words reflect searches for laws and policies about the topic. Due to concerns that search interest in these terms may reflect media reports or legal interest rather than personal experiences, I do not include them in my measure of sexism.

Turning to the group of “Manosphere” words shown in Figure 1.3, note the spike in searches for “SJW” during late 2016. “SJW” is an abbreviation of the term “social justice warrior” and was a common critique of the Hillary Clinton coalition during the 2016 presidential election. While this term is commonly used in the manosphere, it is also used by the political right, and may reflect political ideology rather than manosphere activity. The use of the word misandry may also be hard to interpret because it was adopted and used ironically by members of feminist groups. This usage was highlighted in articles published by Slate, Time Magazine, and The Guardian in August and December of 2014 (Begley, 2014; Hess, 2014; Zimmerman, 2014). And we do see a corresponding spike in search interest for “misandry” in late 2014. Such uses of these words make it hard to interpret them as reflective of sexism.

I exclude these words out of concern that they may reflect political ideology rather than misogyny.

Figure 1.4 shows that search interest in the “Violent” group of words grows steadily between 2006 until around 2014, and then begins to decline. This group of words is perhaps most likely to be related to porn interest. Table 1.4 shows that the related search queries for “gangbang” appear almost exclusively to be porn related and none of the search queries for either word reflect obvious anti-female sentiment. As such, I exclude these words from my final index of misogyny.

In contrast, Table 1.3 demonstrates that the group of “Derogatory” search terms reflect negative stereotypes and aggression toward women. For example, the top 20 search queries for the word “bitch” include “fat” and “crazy,” which are commonly used to disparage women. All five of the words in this group include searches for memes or gifs with the given words, indicating that the searches are for content that ridicules or mocks women. After investigating these four groups of potential search terms, I conclude that the group of “Derogatory” search terms most clearly reflects anti-women sentiment, and I use interest in these terms to construct my index of misogyny.

In addition to interest over time, Google Trends also provides data on the rate of Google search interest for a given phrase at the US state or direct market area (DMA) level within a specified period of time based on a random sample of all searches. They report these relative to the maximum rate of searches for the term experienced in any region during that time period, scaled by 100. For example, if Ohio has the highest rate of searches for the term “weather” across the US, it would receive a score of 100. Then a score of 25 for Indiana would reflect that the search rate for “weather” in Indiana is one quarter of Ohio’s during the time period. I can obtain these for various periods of time, but since each is normalized relative to the maximum in that time period, they cannot be readily compared across time.

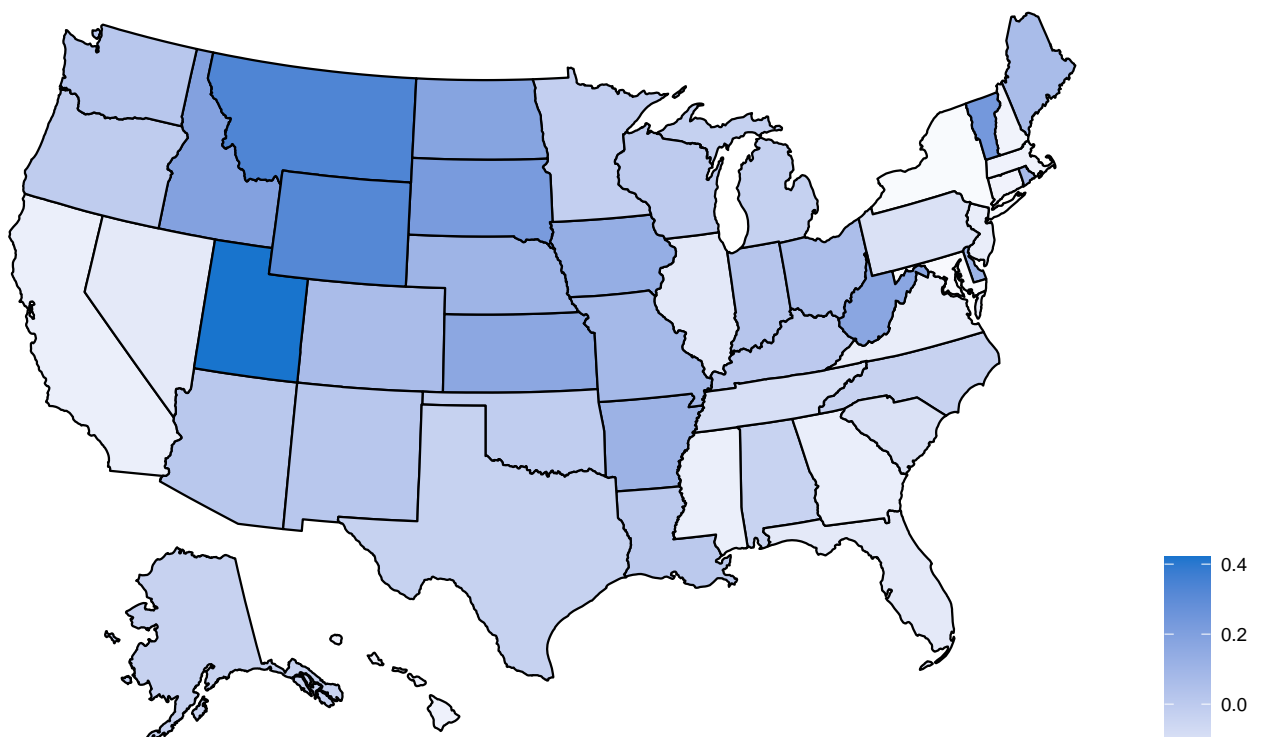
To construct a measure of misogyny or distaste for women, I take each term in the group of “Derogatory” words and collect relative search interest in that term for each region (state or DMA) over the period of a year. I do this for both the state and DMA levels. Because the Google search index is a sample, I collect 10 samples for each year and search term and average over all samples. I normalize the interest in each word by its mean and standard deviation in 2016. For each observation, I take the average over all five “Derogatory” search terms, regress it on year dummies, and use the residuals as my index of misogyny.

To get the average level of misogyny in a state, I simply compute the mean of this index, constructed on state-level data. Figure 1.5 gives the average level of misogyny by state. To obtain other parts of the misogyny distribution, I need additional within-state variation, so I compute this index of misogyny using DMA level data. I then compute the 10th and 90th percentiles of the DMA-based index for each state, weighting each DMA by its population.¹ I also estimate the marginal level of misogyny to be the level of misogyny at the p^{th} percentile of the misogyny distribution where p is the percent of women in the labor force, again weighting by DMA population.

Given that some of the words in the “Derogatory” group are slang for female genitalia, it might be reasonable to worry that the index reflects porn interest rather than sexism. Figure 1.6 provides the correlation between average state misogyny and other measures including Google search interest in the word “porn.” Note that search interest in the the word porn include searches for sites like “Porn Hub” or other sites with porn in the name, so it should be correlated with traffic to porn sites even if it is an imperfect measure. While there is some correlation between searches for porn and my index of misogyny, this correlation is weak and statistically insignificant.

¹I compute DMA populations by summing the populations of all counties within a given DMA, using county populations by year provided by the US Census Bureau, and a mapping between Google trend DMAs and US counties obtained via Kaggle (Pastor, 2021). In cases where DMAs span multiple states, I assign the DMA to the state with fewer censored words in it (due to inadequate search interest).

Figure 1.5: Index of Misogyny by State



This figure gives average misogyny levels by state after residualizing out year effects.

Figure 1.6: Correlation Between Various Measures



This figure gives the correlation between various measures including misogyny, normative sexism, the percent the labor force made up by women, average normalized interest in the manosphere, violent, and reactionary search terms, and porn search interest.

This measure, which captures the use of vulgar and disparaging terms for women, is designed to pick up on true misogyny, not unlike the kind of animus we associate with explicit racism. But it is worth considering whether or not we really believe this kind of distaste makes sense in the context of gender where men and women have always been linked within families. The majority of men fall in love with, marry, and reproduce with women and most were raised by at least one female guardian. While the presence of this kind of animus certainly exists among fringe individuals, it is hard to imagine that enough men are influenced by it to make an impact in the labor market, especially in the Becker framework, where we would not expect men at the top of the prejudice distribution to influence the gender wage gap. Consider how this dynamic differs from that of racial integration between whites and blacks. Though racial segregation has declined over time, integration began less than 100 years ago in many parts of the country and De la Roca et al. (2014) show that segregation between whites and blacks remains quite high even in recent years. As such, it is much easier

to imagine the role racial animus plays in determining the labor market outcomes of black individuals.

1.5.1 Normative Sexism

Now consider the role of gender norms or what I call normative sexism in the labor market. Gender norms may directly dictate a woman's human capital accumulation as well as her decision to remain in or drop out of the labor force after having children. Norms about who is expected to provide for a family could influence wage negotiations both from the employer's perspective as well as from that of a wife conscientious about out-earning her husband. In fact, Fortin (2005, 2008) has shown the relevance of norms for female labor market outcomes including labor force participation in the US and the gender pay gap across countries.

While gender norms do not play directly into either the Becker or Black models of discrimination as originally expressed, we could consider discrimination in these models to describe the cost associated with deviating from traditional gender norms and to depend on how strongly employers value those norms. Under this framing, employers who value more traditional genders norms would incur greater costs from employing women, and the same predictions of the models should work whether we consider a measure of misogyny or a measure of normative sexism.

Thus, I supplement my index of misogyny with a measure of normative sexism using restricted-access GSS responses to questions about the role of men and women in society. Table 1.1 lists the questions I include. They examine opinions about whether a women should work outside the home or focus on caring for her family, whether children of working mothers are worse off, and whether men are better leaders than women. Rather than identifying misogyny or hatred toward women, these questions identify opinions about the appropriate domains of men and women.

Following Charles and Guryan (2008), I recode responses to these questions so that answers reflecting more traditional gender norms have higher numeric value and then normalize responses to each question by the mean and standard deviation in 1977.² For each question k , individual i , and year t , the normalized individual scores are given by

$$\tilde{d}_{it}^k = \frac{d_{it}^k - E[d_{i,77}]}{\sqrt{Var(d_{i,1977}^k)}}.$$

I then average over all questions asked in a given year. That is, I compute


$$D_{it} = \sum_k \tilde{d}_{it}^k / K_t$$

where K_t is the number of these questions asked in year t . Finally, I regress D_{it} on a full set of year dummies to get \tilde{D}_{it} which I aggregate by state to construct the distributional prejudice measures I described earlier.

Figure 1.7 shows the trends in average normative sexism for each census division over time. There is a general decrease in the rigidity of gender norms over time across all regions. Figure 1.8 depicts this index and the questions used to construct it, across the US over time. Respondents give less prejudiced responses in more recent years. Note that sexist responses to the question about how warm and secure a relationship working mothers can have with their children have decreased much more slowly than responses about whether or not a preschooler suffers when his or her mother works. This difference between answers to very similar questions may indicate that responses are sensitive to exactly how a question is phrased, and highlights one concern with survey-based measures like this one. Interestingly, sexist responses to the question about voting for a female president have declined much more slowly than those of other questions.

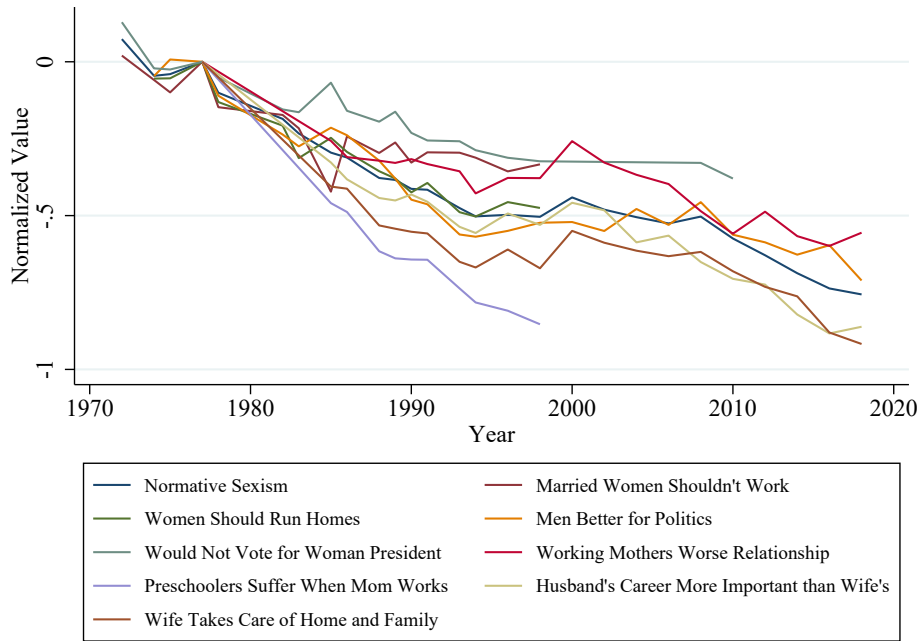
²I normalize by mean and standard deviation in 1977, because this is one of the few years in which all questions are asked.

Figure 1.7: Average Normative Sexism Index Over Time by Census Division



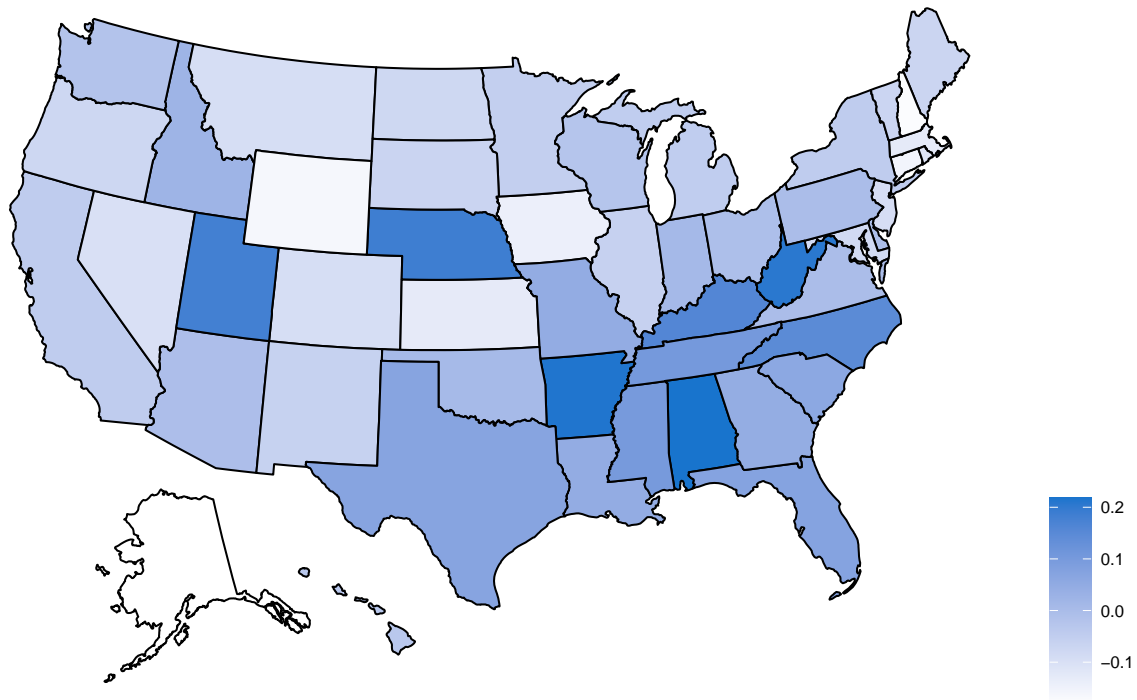
This figure gives the average level of normative sexism over time for each census division.

Figure 1.8: Responses to GSS Prejudice Questions Over Time



This figure gives the time trends in responses to GSS questions about women as well as the normative sexism index.

Figure 1.9: Normative Sexism by State



This figure gives the average level of normative sexism in each state.

Table 1.7 gives the average responses to the GSS questions about women as well as the average prejudice index by census division, after regressing on years to account for changes over time. Prejudice of this kind is highest in the south east of the country and lowest in New England. Most survey questions seem to follow the same geographic trends as the aggregate index of normative sexism, though there is more variation across regions for questions about a working mother's relationship with her children, whether or not married women should work, and whether a wife should put her husband's career first. Figure 1.9 gives the variation in the aggregate index of normative sexism across states. As reflected in Table 1.7, normative sexism is highest in the south east of the country, but there are a few states out west with high levels of normative sexism including Utah and Nebraska. This map differs quite a bit from that of misogyny (Figure 1.5), where we see low misogyny levels across the southeast of the country.


Table 1.7: GSS Responses by Census Division

	Normative Sexism	Married Women Shouldn't Work	Women Should Run Home	Men Better for Politics	Would Not Vote for Woman President
<i>Panel A.</i>					
East South Central	0.147	0.102	0.297	0.234	0.154
West South Central	0.071	0.049	0.065	0.087	0.077
South Atlantic	0.060	0.023	0.126	0.078	0.060
East North Central	-0.027	-0.001	-0.050	-0.034	-0.027
West North Central	-0.035	0.058	-0.049	-0.043	-0.006
Mountain	-0.039	0.033	-0.085	-0.093	-0.044
Pacific	-0.042	-0.076	-0.138	-0.072	-0.058
Middle Atlantic	-0.043	-0.040	-0.040	-0.049	-0.050
New England	-0.128	-0.109	-0.145	-0.138	-0.099
	Working Mothers Worse Re- lationship	Preschool- ers Suffer When Mom Works	Husband's Career More Important Than Wife's	Wife Takes Care of Home and Family	
<i>Panel B.</i>					
East South Central	0.073	0.046	0.145	0.227	
West South Central	0.012	0.033	0.080	0.141	
South Atlantic	0.026	0.007	0.083	0.068	
East North Central	-0.001	-0.024	-0.016	-0.048	
West North Central	-0.040	-0.065	-0.029	-0.090	
Mountain	-0.025	0.033	-0.095	-0.071	
Pacific	0.018	0.040	-0.118	-0.060	
Middle Atlantic	-0.030	-0.021	0.008	-0.041	
New England	-0.096	-0.077	-0.138	-0.176	

This table gives average GSS responses by census division after residualizing out year indicators. Standard errors are given in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 1.10: Misogyny vs. Normative Sexism Indices

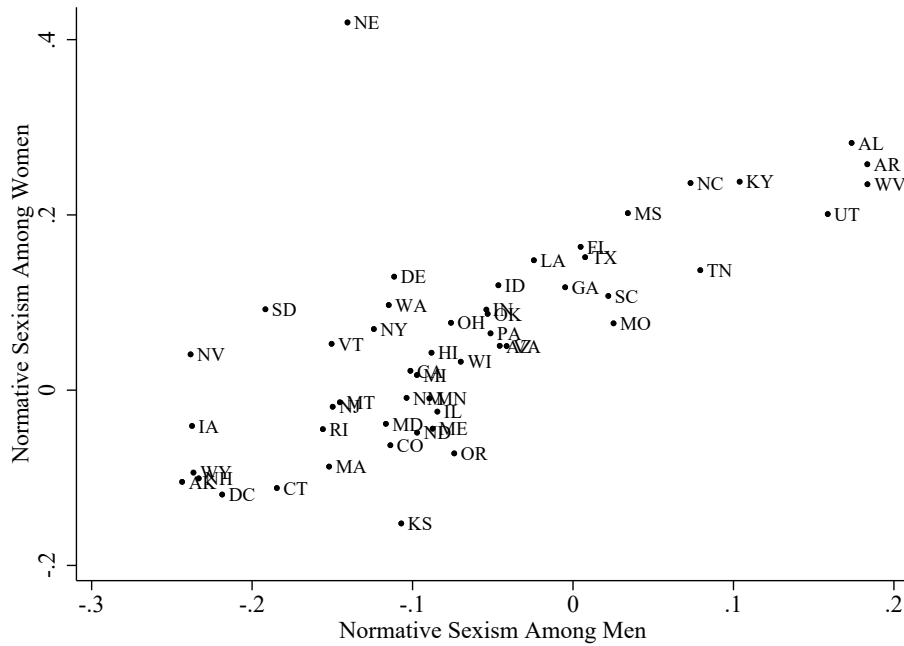


This figure plots average misogyny against average normative sexism for each state. The dashed grey line gives the linear relationship between the two.

Regional differences in the distribution of these two measures suggest that they are not simply two manifestations of the same type of sexism. Figure 1.10, which plots the index of gender norms against the index of misogyny, emphasizes this point. The two measures are positively correlated, but weakly so. Notice states like Alabama which have very strong gender norms, but below average levels of misogyny. Conversely, states like Montana have very high levels of misogyny despite weaker gender norms.

Figure 1.11 gives plots the average normative sexism by state constructed using only responses from men against that of women. Note the strong positive correlation (p-value less than 0.001) between normative sexism levels of men and women. This highlights the fact that this kind of sexism is different from misogyny toward women. It reflects societal expectations about the role of men and women which are shared across genders. It does not reflect distaste for women or discrimination against women which we would expect to be higher among men, on average.

Figure 1.11: Normative Sexism - Men vs. Women



This figure gives the average level of normative sexism constructed from responses by men against the average level of normative sexism constructed from responses by women.

For some, it may be counter-intuitive to disentangle these two types of sexism. But consider a man who has grown up in a community in which women nearly always stay home and care for children. He may think this is the best way for a child to grow up and may want this for his children, but may not have any animosity toward women. In fact, he may value the role women play in raising children and caring for the home. He may have been taught to treat women with respect and the emphasis on gender roles may make him less likely to disparage women in the way that would be picked up by our misogyny measure. And note that women in such a community may share these values as we observe empirically. On the other hand, imagine a man who has been unsuccessful with women romantically and outperformed by them in the workplace. Such a man may have developed a distaste for all women regardless of which roles they occupy since they have spurned or surpassed him in both public and private domains. While both men exhibit sexism, it is not clear that both types of sexism

have the same labor market implications or that they are equally distributed throughout the population.

1.5.2 Measures of Sexism and Other Labor Market Measures

In addition to empirically testing the predictions of the Black and Becker models with regard to the gender wage gap, I also estimate the relationship between each measure of sexism and other labor market outcomes including the proportion of women aged 20 to 40 who remain unmarried, the average age at which women give birth to their first child, labor force participation, and college degree attainment, after conditioning on relevant observables. Recall that if prejudice is influencing decisions about educational attainment and labor market participation, estimates of the relationship between prejudice and the residual wage gap will be underestimates. I regress each outcome on my measure of misogyny, overall normative sexism, and gender-specific normative sexism separately.

I use CDC data on natality between 2016 and 2019 to obtain the average age of women when they give birth to their first child. I take the average in each state after residualizing out year effects. For the remaining outcomes, I use the same Census Outgoing Rotation Group data that I used to compute the residual wage gaps. I estimate college degree attainment controlling for age and year dummies. For the proportion of women never married between the ages of 20 and 40 and full-time labor force participation rates, I control for education, age, and year indicators. In all cases, I weight estimates by the census earnings weight.

These regressions provide evidence about the role of both misogyny and normative sexism in determining labor market outcomes for women. If discrimination reduces the human capital attainment and experience of women relative to men, then examining the relationship between the wage gap and discrimination, controlling for these factors, will understate the effects of discrimination. If these choices are determined not by discrimination, but by gender

norms on the other hand, policies implemented to prevent discrimination will have little effect on these gaps. I distinguish between normative sexism among men and normative sexism among women to consider whether these impacts are driven by norms that are pushed on women by men or that are internalized by women and affect her decisions accordingly.

1.6 Results

1.6.1 Measures of Sexism and Labor Market Outcomes

Tables 1.8 and 1.9 give the average impact of both indices on outcomes related to the labor market. Panel A of Table 1.8 shows that the residual wage gap is increasing in both measures and that this relationship is statistically significant. Here, a negative coefficient indicates an increase in the wage gap. We can also see that the relationship between normative sexism and the wage gap is driven by normative sexism among women, rather than men. Normative sexism among women is a statistically significant predictor of the wage gap whether or not we include misogyny, but normative sexism among men is not.

The strength of the relationship between misogyny and the residual wage gap is striking. A one standard deviation increase in misogyny is associated with with a reduction in women's relative wages by about .0166 log points or approximately 4.3 percent of the average residual wage gap. This effect is on the same order of magnitude (though slightly larger) as that of normative sexism. A one standard deviation increase in normative sexism is associated with a reduction in women's relative wages by about 0.0144 log points or approximately 3.7 percent of the mean residual wage gap. Despite the high levels of interaction and coordination between men and women in private domains, overt misogyny appears to be an important factor in determining the wage gap.

Table 1.8: Relationship between Sexism Indices and Labor Market Outcomes (Part 1)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A. Dependent Variable - Wage Gap</i>							
Misogyny	-0.108** (0.038)				-0.098* (0.038)	-0.100* (0.038)	-0.101** (0.037)
Normative Sexism		-0.134* (0.059)			-0.114* (0.057)		
Normative Sexism - Men			-0.094 (0.054)			-0.075 (0.052)	
Normative Sexism - Women				-0.131* (0.058)			-0.118* (0.055)
<i>N</i>	51	51	51	51	51	51	51
<i>R</i> ²	0.14	0.09	0.06	0.09	0.21	0.17	0.21
<i>Panel B. Dependent Variable - Proportion of Women Aged 20-40 Never Married</i>							
Misogyny	-0.280*** (0.079)				-0.252*** (0.070)	-0.255** (0.072)	-0.266*** (0.070)
Normative Sexism		-0.436*** (0.116)			-0.397*** (0.104)		
Normative Sexism - Men			-0.344** (0.107)			-0.308** (0.096)	
Normative Sexism - Women				-0.403** (0.115)			-0.383*** (0.101)
<i>N</i>	48	48	48	48	48	48	48
<i>R</i> ²	0.21	0.23	0.18	0.21	0.41	0.36	0.40
<i>Panel C. Dependent Variable - Average Age at First Birth</i>							
Misogyny	-4.660*** (1.200)				-3.925*** (0.982)	-4.024*** (1.031)	-4.154*** (1.009)
Normative Sexism		-8.108*** (1.594)			-7.293*** (1.410)		
Normative Sexism - Men			-6.577*** (1.485)			-5.862*** (1.320)	
Normative Sexism - Women				-7.273*** (1.617)			-6.654*** (1.412)
<i>N</i>	51	51	51	51	51	51	51
<i>R</i> ²	0.24	0.35	0.29	0.29	0.51	0.46	0.48

This table gives the relationship between various labor market outcomes and measures of sexism including misogyny, normative sexism, normative sexism among men, and normative sexism among women. The residualized wage gap in each state (Panel A) is computed using log wages, controlling for education, a quadratic in potential experience, and year, and weighting by the census earnings weight. The proportion of women aged 20-40 never married (Panel B) controls for age, education level, and year indicators and weighting by the census earnings weight. The average age at first birth (Panel C) is computed controlling for education level and year indicators. All regressions are weighted by the precision at which the outcome is estimated. Standard errors are given in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1.9: Relationship between Sexism Indices and Labor Market Outcomes (Part 2)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel A. Dependent Variable - Full-Time Labor Force Participation Gap</i>							
Misogyny	-0.112** (0.034)				-0.101** (0.033)	-0.103** (0.034)	-0.105** (0.033)
Normative Sexism		-0.140* (0.054)			-0.118* (0.050)		
Normative Sexism - Men			-0.100* (0.049)			-0.078 (0.046)	
Normative Sexism - Women				-0.137* (0.053)			-0.122* (0.049)
<i>N</i>	51	51	51	51	51	51	51
<i>R</i> ²	0.18	0.12	0.08	0.12	0.26	0.23	0.28
<i>Panel B. Dependent Variable - Full-Time Female Labor Force Participation Rate</i>							
Misogyny	-0.061 (0.033)				-0.055 (0.033)	-0.059 (0.034)	-0.055 (0.032)
Normative Sexism		-0.076 (0.051)			-0.063 (0.051)		
Normative Sexism - Men			-0.031 (0.047)			-0.019 (0.046)	
Normative Sexism - Women				-0.106* (0.049)			-0.099* (0.048)
<i>N</i>	51	51	51	51	51	51	51
<i>R</i> ²	0.06	0.04	0.01	0.09	0.09	0.07	0.14
<i>Panel C. Dependent Variable - Gender Gap in College Degree Attainment</i>							
Misogyny	0.011 (0.014)				0.015 (0.014)	0.015 (0.014)	0.014 (0.014)
Normative Sexism		-0.039 (0.021)			-0.043 (0.021)		
Normative Sexism - Men			-0.029 (0.019)			-0.032 (0.019)	
Normative Sexism - Women				-0.039 (0.021)			-0.041 (0.021)
<i>N</i>	51	51	51	51	51	51	51
<i>R</i> ²	0.01	0.07	0.04	0.07	0.09	0.06	0.09
<i>Panel D. Dependent Variable - Female College Degree Attainment</i>							
Misogyny	-0.366*** (0.100)				-0.322*** (0.092)	-0.322** (0.093)	-0.342*** (0.094)
Normative Sexism		-0.535** (0.153)			-0.465** (0.140)		
Normative Sexism - Men			-0.466** (0.138)			-0.399** (0.126)	
Normative Sexism - Women				-0.433** (0.156)			-0.386** (0.140)
<i>N</i>	51	51	51	51	51	51	51
<i>R</i> ²	0.22	0.20	0.19	0.14	0.36	0.35	0.32

This table gives the relationship between various labor market outcomes and measures of sexism including misogyny, normative sexism, normative sexism among men, and normative sexism among women. The labor force participation gap (Panels A and B) are computed based on full-time labor force participation, controlling for age, education, and year indicators and weighting by the census earnings weight. College completion rates (Panels C and D) are computed, controlling for age and year indicators and weighting by the census earnings weight. All regressions are weighted by the precision at which the outcome is estimated. Standard errors are given in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Panel B of Table 1.8 shows the relationship of each measure with the proportion of women between ages 20 and 40 who have never married after controlling for level of education, age, and year. Both measures are associated with smaller shares of such women, implying that women in areas with stronger gender norms and more misogyny are more likely to marry young, all else equal. Similarly, Panel C shows that both measures are associated with lower average ages at first birth (after controlling for year fixed effects), meaning women tend to start having children earlier in areas with more traditional gender norms and greater misogyny, all else equal. While it is not surprising that gender norms influence the probability of marriage and timing of child bearing, the role of misogyny in these decisions is less obvious. Given the relationship between misogyny and the wage gap, however, it is possible that women respond to labor force discrimination by exiting the labor force in favor of marriage and motherhood.

Panel B of Table 1.9 shows that full-time labor force participation among women is decreasing in normative sexism among women. It is also decreasing in misogyny, and the coefficient is just shy of statistical significance at the 5% level. When we consider the gap between the full-time labor force participation of women relative to men, given in Panel A of Table 1.9, we see that the gap in participation is increasing in both types of sexism and the relationships are statistically significant. That is to say, when we look at women, irrespective of men, norms internalized by women are the strongest predictors of their labor force participation rates. But when we consider labor market participation relative to men, it is clear that both misogyny and norms are important determinants of labor force participation.

Panel C of Table 1.9 shows that normative sexism is associated with an increase in the gender gap in college degree attainment and that this relationship is statistically significant at the 10% level. Misogyny is actually associated with a smaller gap in degree attainment though this relationship is statistically insignificant. Panel D shows that though the gap between men and women is smaller in areas with higher misogyny, women there are still getting less

education (conditional on age and year dummies) than in areas with lower misogyny, all else equal. This is consistent with lower education among both men and women in high misogyny areas. Panel D makes clear that strongly gendered norms and misogyny are both related to reduced educational attainment among women.

Overall, it is clear that both misogyny and gender norms play an important role in determining the residual wage gap as well as in other labor market outcomes. We cannot simply attribute differences in outcome to differences in choices driven by gender norms, though these do play an important role. Moreover, the fact that misogyny strongly correlates with the choices women make regarding human capital accumulation and labor force participation suggests that estimates linking discrimination to the residual wage gap understate the impact of discrimination.

1.6.2 Testing Models of Discrimination

Table 1.10 gives the regressions of the residual wage gap on various parts of the misogyny distribution as well as on the percent of women in the labor force and the percent of the population that is sexist. A negative coefficient indicates an increase in the wage gap between men and women. I do not find evidence that the wage gap is determined by a marginal discriminator as predicted by the Becker model, since marginal prejudice is a weaker predictor of the wage gap than is mean prejudice. Moreover, the wage gap is not more strongly related to the bottom of the prejudice distribution than to the top. Given the high degree of labor market integration between men and women, it is not surprising that a marginal discriminator plays little role in determining the gender wage gap. What is more telling, perhaps, is the fact that the wage gap is decreasing, rather than increasing in the percent of women in the labor force. This is completely inconsistent with Becker's predictions, even when we allow for labor market integration as Neumark (1988) does. This is however,

Table 1.10: Testing the Predictions of Black and Becker Models Using the Index of Misogyny

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mean	-0.108** (0.038)		-0.215* (0.082)	-0.053 (0.059)			-0.014 (0.027)
Marginal		-0.095* (0.043)	0.092 (0.082)	0.039 (0.055)			
Percent Women				2.022*** (0.276)		2.067*** (0.260)	2.025*** (0.249)
Lower Tail					-0.008 (0.043)	-0.004 (0.027)	
Upper Tail					-0.025 (0.021)	-0.003 (0.014)	
Percent Sexist							-0.042 (0.073)
<i>N</i>	51	46	46	46	46	46	51
<i>R</i> ²	0.14	0.10	0.22	0.66	0.14	0.65	0.66

This table gives the relationship between the residual wage gap (controlling for education, a quadratic in potential experience, and year) and the percent of women in the labor force, the average percent of the population that responds with sexism on the GSS, and various parts of the misogyny distribution (the mean, the marginal, and the upper and lower tails). All regressions are weighted by the precision at which the residual wage gap is estimated. Standard errors are given in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

consistent with Black's search model of discrimination. Also consistent with Black, the wage gap is increasing in the proportion of the population that is sexist, though this relationship is statistically insignificant. Given that Black's model allows for the influence of a small contingent of very sexist employers, it follows that this model may fit the data on the gender wage gap better than Becker.

Table 1.11 gives the same results using the measure of gender norms. As in the case of misogyny, the mean level of normative sexism is a stronger predictor of the wage gap than is the marginal level. However, we do see that there is some evidence that the lower tail of the prejudice distribution does impact the wage gap more strongly than does the top of the

Table 1.11: Testing the Predictions of Black and Becker Models Using the Index of Normative Sexism

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Mean	-0.134*		-0.425**	-0.035			-0.150
	(0.059)		(0.125)	(0.095)			(0.075)
Marginal		-0.050	0.301*	-0.021			
		(0.057)	(0.116)	(0.086)			
Percent Women				2.056***		1.963***	2.024***
				(0.255)		(0.245)	(0.221)
Lower Tail					-0.244**	-0.056	
					(0.078)	(0.056)	
Upper Tail					-0.001	-0.015	
					(0.042)	(0.028)	
Percent Sexist							0.206
							(0.143)
<i>N</i>	51	51	51	51	51	51	51
<i>R</i> ²	0.09	0.02	0.21	0.67	0.21	0.67	0.68

This table gives the relationship between the residual wage gap (controlling for education, a quadratic in potential experience, and year) and the percent of women in the labor force, the average percent of the population that responds with sexism on the GSS, and various parts of the normative sexism distribution (the mean, the marginal, and the upper and lower tails). All regressions are weighted by the precision at which the residual wage gap is estimated. Standard errors are given in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

distribution, when we do not include the percent of women in the labor force in our model. Note here that the proportion of sexist individuals is actually decreasing in the wage gap rather than increasing in it, though this estimate is not statistically significant. The sign of this coefficient has switched in this model relative to the model controlling for misogyny. This may be due to the relationship between our normative sexism measure and the percent of sexist individuals since they are both constructed using GSS responses to the same survey questions about gender norms. The strength of the correlation between the two can be seen in Figure 1.6.

The positive and statistically significant relationship between the percent of women in the labor force and the residual wage gap is consistent with Black's search model of discrimination and not with the Becker model. While there is some evidence that the wage gap is influenced by prejudice at the bottom of the normative sexism distribution, there is no evidence of a marginal discriminator when I examine the relationship between the gender wage gap and misogyny. This is consistent with a small group of misogynistic employers influencing women's wages even if they are unlikely to hire women, which the Black model allows for. Under Becker on the other hand, prejudice can only be reflected in a wage gap if it impacts employers who actually hire women. If we believe that normative sexism may be more prevalent than outright misogyny, it makes sense that the Becker framework may fit slightly better in this context.

1.7 Conclusion

In this paper, I construct an entirely novel measure of misogyny based on Google search interest in misogynistic terms used to describe women, and present evidence that this measure captures true distaste for women. In addition to my measure of misogyny, I also construct a measure of gender norms based on survey responses regarding the role of men and women in society. Using both measures, I test the predictions of both the Becker and Black models of discrimination in the context of the gender wage gap and their relationship with various labor market outcomes which contribute to wages.

Quite surprisingly, I find that misogyny is strongly related to the residual wage gap, even when gender norms are accounted for. It is also an important determinant of other labor market indicators including the probability that women aged 20 to 40 remain unmarried, the age at which women have their first child, labor force participation, and college degree attainment. While subtle discrimination against or underestimation of women may be un-

surprising in the workplace, this measure picks up on the use of vulgar words that reflect outright disdain for women. It's relationship with the wage gap implies that a much more overt kind of discrimination is acting against women in the labor market. Moreover, the impact of misogynistic discrimination against women on inputs to the wage equation such as educational attainment and labor market participation imply that estimates of the impact of discrimination on the residual wage gap understate the problem.

I also use my measure of misogyny to test the predictions of two strains of labor market discrimination models: the original Becker model of discrimination (1957) and Black's 1995 search model of discrimination. I find evidence consistent with Black's search model, and show that the data is inconsistent with Becker-type models.

This paper highlights that even accounting for norms, distaste for women is an important predictor of the gender wage gap. This type of sexism reflects more than subtle or implicit bias against women as the measure reflects language that reveals anger and contempt toward women. This paper provides compelling evidence that the gender wage gap cannot be waved away as merely reflecting women's choices, gender norms, or subtle discrimination against women. In fact, true misogyny plays an important and previously unmeasured role in determining the wage gap.

Chapter 2

The Effects of Sheltering-in-Place on Domestic Violence Survivor Response Strategies

2.1 Introduction

According to the CDC, approximately 1 in 4 women experience some form of intimate partner violence during their lifetime. S. Smith et al. (2015) Despite this pervasiveness, much is still unknown about factors that exacerbate domestic violence and the strategies victims use to cope with abuse. Concerns about COVID-19 mitigation policies brought the question of what contributes to domestic violence to the forefront of public attention in the wake of the pandemic. To reduce the spread of the virus and prevent COVID-related deaths, US cities and states began implementing shelter-in-place or lockdown orders one-by-one beginning with the San Francisco Bay Area on March 17th, 2020. By April 7th, 42 states had enacted such orders. Washington D.C., Puerto Rico, and individual cities in Wyoming, Utah, and

Oklahoma were also under stay-at-home orders by this time. Even in areas that did not lock down, school and business closures resulted in massive reductions in mobility. These lockdowns isolated families together in close quarters under stressful conditions, making women and children in abusive homes especially vulnerable. Almost immediately, the media began to report concerns about increased domestic violence (e.g. Economist, 2020; Neuman, 2020; Taub, 2020). While the topic has been explored in other work, evidence has been limited to a smattering of cities and relies on crime or calls-for-service data, which are subject to under-reporting concerns and are not recorded consistently from city to city. Instead, I leverage Google Trends data on search terms associated with incidents of and responses to domestic violence, taking advantage of the timing of these lockdowns and the subsequent decreases in mobility to estimate the effect of staying-at-home on domestic violence. This allows us to better understand the consequences of such policies as well as to improve upon our understanding of domestic violence risk factors and the coping strategies employed by domestic violence survivors.

Several studies document the relationship between external sources of stress and domestic violence. For example, Kirby et al. (2014) finds an increase in domestic violence associated with both wins and losses by the English national team. Moreover, the increase associated with losses is much larger than that of wins. Some research has specifically linked natural disasters to increases in domestic violence. Gearhart et al. (2018) find that exposure to natural disaster results in more assaults in Florida between 1999 and 2007. Parkinson (2019) interviews Australian women after the 2009 bushfires. A majority of these women reported new or increased exposure to domestic violence and identified the disaster as the cause of the additional violence. Enarson (1999) interviewed staff at domestic violence programs in Canada and the US and reveals that they report higher levels of demand for their services in the wake of a disaster. Financial stress has also been linked with domestic violence risk. Renzetti (2009) details the relationship between economic stress and domestic violence. Beland et al. (2020) use survey data to learn about sources of family stress and its associations

with domestic violence, finding that a family's inability to keep up with their financial obligations is associated with greater reported risk of domestic violence. Using survey data collected from Australian women, Morgan and Boxall (2020) find that financial distress nearly doubles the probability of first-time domestic violence in the context of the COVID-19 pandemic specifically. The onset of the COVID-19 pandemic introduced several sources of stress simultaneously. These included concerns about avoiding illness itself, financial stress due to business shutdowns and resultant unemployment, and childcare disruptions which forced many to assume full time responsibility for child care and home schooling regardless of other work obligations. Given the relationship between domestic violence and stress, we might expect an increase in incidents of domestic violence as a result of these COVID-19 lockdowns, and in fact, there is already some evidence that this may have been the case.

Of the papers that have been done on the topic, estimates of the effect vary. In their meta-analysis of 18 studies conducted on domestic violence during the COVID-19 pandemic, Piquero et al. (2021) find that of studies conducted in the US, the average effect size varied from a 22 percent decrease in domestic violence to a 38.15 percent increase in domestic violence. Most studies they examine find some positive effect and the average effect size is positive 8.10 percent. Leslie and Wilson (2020) examine 14 large metropolitan areas and find a 7.5 percent increase in domestic violence calls-for-service as a result of lockdown orders. Nix and Richards (2021) analyze six U.S. jurisdictions and find significant increases in calls for service related to domestic violence in five of the six jurisdictions, though effect sizes vary dramatically. Hsu and Henke (2021) examine 36 US jurisdictions and find an average increase in domestic violence of five percent. Henke and Hsu (2022) look at the relationship between domestic violence incidents and mobility directly for a sample of 32 US cities finding that a 10 percentage point increase in the average percent of people who stayed home all day resulted in an additional 0.052 daily reports of domestic violence per 100,000 people. Not all studies have found an increase in domestic violence. For example, Ashby (2020) used police-recorded crime data on serious assaults in residences (a proxy for domestic violence)

in eight large U.S. cities and found that these assaults were within predicted confidence intervals in the weeks after lockdown orders were implemented. None of these study cover the US as a whole, and as far as I know, no such study has yet been done.

Attempts to measure the relationship between COVID-19 lockdowns and domestic violence thus far have relied almost exclusively on local police calls-for-service or crime data. These data sources have several weaknesses. For one, victims of domestic violence may fear that calling the police will only exacerbate the issue and are likely to do so only if violence reaches a tipping point. Moreover, this data is only available in certain cities and recording methods vary widely from department to department. Very few departments provide a domestic violence indicator and wide variation in crime recording practices make it hard to consistently identify domestic violence across departments. Domestic violence crimes are believed to be widely under-reported (Reaves, 2017) and reporting patterns are likely to change as a result of the pandemic. For example, child abuse is often caught by teachers or other non-family adults who are unable to perform this role during lockdowns. Similarly, friends or relatives who might ordinarily recognize signs of abuse and encourage reporting may be limited in this capacity during lockdowns. On the other hand, with many people at home who would ordinarily not be, neighbors may be more likely to observe signs of domestic violence, which could result in increased third-party reporting as a result of the pandemic. While survey data may address some of these reporting issues, surveys on domestic violence are generally limited to annual accounts of abuse at the state or national level.

To combat these problems, I use data Google Trends data on search prevalence for queries that may be searched by victims of domestic violence including “domestic violence hotline,” “domestic violence help,” “domestic violence shelters,” and “restraining order.” These search terms are logical first steps for victims of domestic violence in search of resources and coping strategies. The use of Google search trends to identify domestic violence patterns is not unprecedented. Koutaniemi and Einiö (2019) demonstrate the relationship between Google

searches for domestic violence and shelters and show that peaks in these searches correspond to the timing of peaks in police reports of domestic violence incidents in Finland. Because Google searches are conducted in private, they are not subject to the same concerns about under-reporting as calls-for-service or crime data. Moreover, they are uniformly available across the United States, do not rely on reporting codes or systems that vary from region to region, and are available at weekly intervals. Given the difficulty in collecting data on domestic violence, this method is particularly useful. Beyond the practical advantages of Google Trends data, search patterns may provide insights into what strategies women use to cope with domestic violence. For example, do women search for shelters in order to exit an abusive home or are they more likely to search for domestic violence hotlines which may provide emotional support and offer passive coping strategies that do not involve leaving their current home? Understanding the tools domestic violence victims are actually searching may be crucial for identifying effective intervention strategies.

This study offers an estimate of the effect of COVID-19 lockdown orders on domestic violence that is not limited to a handful of cities or regions and does not rely on crime or calls-for-service data. I find that search interest for domestic violence hotlines increases significantly as a result of COVID-19 lockdowns, but that searches for other resources (including shelters and restraining orders) actually decrease. I hypothesize that victims of domestic violence may be substituting away from resources that require them to leave their home during a pandemic and toward coping strategies that they can employ while staying at home, at least in the short term. In addition, I explore several potential mechanisms for these effects including: (1) increased exposure between perpetrators and victims due to the need to stay at home together, (2) household earning dynamics that may influence a women's decision to leave, (3) reductions in income due to unemployment that add stress to families and reduce a women's independent earning potential, and (4) reduced probability that perpetrators are caught since contact with potential third-party reporters is limited. I find evidence that increased exposure and financial stress are important channels through which these

lock-downs influence domestic violence search behavior. There is also some evidence that female employment may reduce domestic violence. Finally, I see a dramatic and statistically significant decline in search interest for child abuse and child protective services as a result of COVID-19 lockdowns which suggests that third-party reporters such as teachers and health care professionals are less able to identify and report signs of child abuse during lockdowns. Section 2.2 provides an overview of the relevant mechanisms that may contribute to domestic violence. Section 2.3 details the data and identification strategy I use to measure domestic violence as well as to tease out the potential mechanisms at play. Finally, I present my results and discuss my findings in Section 2.4, and discuss some robustness and other checks in Section 2.5.

2.2 Mechanisms Influencing Domestic Violence

In this section, I will discuss several important factors that influence rates of domestic violence and their relevance in the context of COVID-19 lockdowns. Note that for convenience, I frame violence as being inflicted upon women by men, though this need not be the case.

Exposure Reduction Theory

The exposure reduction theory of domestic violence postulates that the more time a potential victim spends with a perpetrator, the greater the incidence of domestic violence. There is some empirical evidence directly supporting this theory. Chin (2012) examines the relationship between women’s workforce participation and spousal violence in India, finding a reduction in violence for working women even among those participating in unpaid work. Dugan et al. (1999, 2003) investigate the decline in intimate partner homicides in the US between 1976 and 1992 and argue that it is driven by exposure reduction via a decline of domestic partnerships (both marital and non-marital), improvement in women’s economic

opportunities, and access to additional resources which allow women to reduce their exposure to their abusers.

In addition, a number of studies have identified seasonal patterns of domestic violence that are consistent with this theory. G. Farrell and Pease (1994) find that police calls for domestic disputes in Merseyside, England tend to peak in July and in late December and early January. Similarly, Oths and Robertson (2007) analyze call records to a domestic violence shelter in Alabama and also find peaks in calls during the summer and in January. While there may be alternative explanations, these peaks in domestic violence during holidays (when families generally spend more time in close proximity) are consistent with exposure reduction theory.

Reduced mobility during shelter-in-place orders resulted in families spending much more time together. Adults were often unable to leave the house for work and school closures meant children were forced to learn from home. Under exposure reduction theory, we expect such an increase in family time to result in more instances of domestic violence at the margin. Additionally, fear of the virus may have made some options for reducing exposure (e.g. staying at a shelter or with friends/relatives) less appealing. To test the role of exposure reduction, I assume that exposure is inversely related to mobility. That is, in areas where there is a larger decrease in mobility, I assume within-family exposure increases. I examine the effects of sheltering-in-place on domestic violence related search interest in places with variation in the size of mobility reduction to see if effects are stronger in areas where mobility was reduced by more. I also look at how the effect varies with COVID-19 anxiety which I proxy for using Google search interest in “COVID-19” and “coronavirus,” since it’s plausible that in areas where COVID-19 anxiety is higher, the exposure reduction channel may be more intense.

Bargaining vs. Male Backlash

Farmer and Tiefenthaler (1997) model domestic violence as an outcome in a non-cooperative household bargaining game. They assume that perpetrator utility is increasing in violence

via feelings of power, control, and self-esteem, which are bolstered by violent acts against his spouse. An abuser's ability to inflict violence is reduced by his spouse's prospective utility upon exiting the relationship. A potential victim's ability to earn money and the presence of outside resources available to her (e.g. shelter, family support, etc.) bolster her threat-point and reduce the probability of violence against her under this framework.

Alternatively, male backlash theory postulates that an increase in women's earnings can actually threaten men, particularly if her earnings begin to outpace his own, and that men will respond with violence to re-establish their dominance. There is empirical evidence for both theories. For example, Dhanaraj and Mahambare (2022) find that women in urban India who participate in paid work experience domestic violence at higher rates than those who do not, while Aizer (2010) finds that increases in women's wages relative to men does in fact reduce domestic violence. Macmillan and Gartner (1999) find that the impact of a woman's labor force participation on rates of domestic violence depends upon the employment status of her husband, reducing the probability of domestic violence if her husband also works and increasing it if he does not.

To proxy for these bargaining dynamics, I examine how the effects of sheltering-in-place on domestic violence vary in areas with different female-to-male employment ratios in the weeks leading up to shelter-in-place orders. Under the household bargaining framework, I would expect increases in domestic violence search interest to be weaker in areas where women are more likely to work relative to men. Under male backlash on the other hand, I would expect increases in domestic violence search interest to be stronger in areas where women are more likely to work relative to men.

Financial Stress

As discussed in Section 2.1, financial stress has also been linked empirically to increases in domestic violence. The COVID-19 pandemic resulted in massive unemployment as businesses responded to a huge decrease in demand from consumers who were staying home. I explore

how the effect of sheltering-in-place on domestic violence related search interest varies in areas with larger increases in unemployment to see if areas that are hit hardest financially also experience the greatest increases in domestic violence search interest.

Probability of Detection

The Becker (1968) model of crime offers the insight that crime rates should increase as the probability of getting caught decreases. Lockdowns reduce the interactions between victims of domestic violence and potential reporters including teachers, doctors, friends, and extended family. This means that any evidence of violence (e.g. bruises) are much less likely to be noticed and reported by someone outside the household. Knowing that acts of violence are less likely to be reported by an external observer may reduce a perpetrator's fear of being caught and result in increased violence. There is already some evidence to support this notion. Morgan and Boxall (2020) find that decreased contact with family and friends outside of the household is associated with more domestic violence in the COVID-19 context based on their survey of Australian women.

In order to determine whether or not the probability of detection is playing a role in domestic violence behavior during the COVID-19 lockdowns, I can attempt to identify patterns associated with child abuse reporting. Because child abuse is much more likely to be reported by third-parties such as teachers or health care workers, I would expect violence toward school-aged children to be restrained when children are in school. In fact, Iyengar (2009) finds evidence that this kind of reporting may reduce domestic homicide rates among school-aged children. To gauge whether there is any evidence of reduced reporting during COVID-19 lock downs, I look at search interest in terms that might be used by a third party to find out how to identify or report the behavior (e.g. searches for child protective services) and compare this with searches that are likely to be searched by victims directly (e.g. domestic violence hotline). While this does not directly measure whether abusers respond to decreased supervision, it may suggest a potential channel via decreased reporting behavior.

2.3 Data and Methods

I use Google Trends data on the search prevalence of terms related to domestic violence resources (including “domestic violence hotline,” “domestic violence help,” “domestic violence shelter,” “womens shelter,” “restraining order,” and “protective order”) to proxy changes in domestic violence levels at the direct market area (DMA) level. For each DMA, Google Trends provides the search rate for a given search term relative to the total number of searches within that region and time period, normalized by the maximum relative search interest within that region and a given time frame. I collect this information between 2016 and 2020 in weekly intervals for each of 210 DMAs in the United States. Table 2.1 gives the top 10 search queries related to the search terms “domestic violence hotline,” “domestic violence help,” “domestic violence shelter,” “womens shelter,” “restraining order,” and “protective order.” The related searches for “domestic violence hotline,” “domestic violence help,” “domestic violence shelter,” and “womens shelter,” are consistent with searches available to victims of domestic violence. For example, the top related search to “domestic violence shelter” is “domestic violence shelter near me.” Likewise, related searches to “restraining order” and “protective order” are consistent with information-seeking about how to obtain such an order. For instance, the top search for “restraining order” is “get a restraining order,” and domestic violence is one of the most common reasons for obtaining such orders. For this analysis, I average together search interest in “domestic violence shelter” and “womens shelter.” I also average together interest in “protective order” and “restraining order.”

In order to measure the level of victim-perpetrator exposure, I use Apple mobility trends data at the county level which gives me daily data on the volume of requests for driving directions in each county. While imperfect, this allows me to observe how well individuals in a given area comply with stay-at-home orders. That is, how likely a family in a given area is to remain together at home. Thus, as mobility levels fall, I assume that exposure risk is greater. To merge this mobility data with Google Trends data, which is given at the DMA

Table 2.1: Top Related Queries - Outcomes of Interest

Top Related Queries	Interest	Top Related Queries	Interest
<i>“domestic violence hotline”</i>		<i>“domestic violence help”</i>	
national domestic violence hotline	100	help for domestic violence	100
national domestic violence	100	domestic abuse	17
national domestic hotline	100	help for domestic violence victims	16
domestic abuse	20	domestic violence help near me	15
domestic abuse hotline	19	what is domestic violence	13
abuse hotline	19	domestic violence housing	11
domestic violence number	14	domestic violence hotline	9
domestic violence hotline number	13	domestic violence shelter	8
domestic violence shelter	10	domestic violence support	7
suicide hotline	9	how to help domestic violence victims	7
<i>“domestic violence shelter”</i>		<i>“womens shelter”</i>	
domestic violence shelter near me	100	womens shelter near me	100
shelter near me	100	shelter near me	100
domestic violence shelters	81	womens homeless shelter	37
women shelter	53	homeless shelter	37
domestic violence center	48	womens shelters	30
homeless shelter	46	womens shelter donation	30
domestic abuse shelter	35	womens center	27
domestic violence hotline	35	women shelter	25
domestic violence shelters near me	30	womens shelter donations	23
shelters near me	30	womens and childrens shelter	19
<i>“restraining order”</i>		<i>“protective order”</i>	
get a restraining order	100	protective order texas	100
file restraining order	75	restraining order	95
file a restraining order	64	what is a protective order	78
how to get restraining order	54	emergency protective order	42
how to get a restraining order	49	protective order violation	38
how to file restraining order	38	indiana protective order	37
how to file a restraining order	34	protection order	35
restraining order california	30	violation of protective order	25
what is restraining order	28	protective order virginia	23
temporary restraining order	27	motion for protective order	23
<i>“child protective services”</i>		<i>“child abuse”</i>	
cps	100	child sexual abuse	100
cps child protective services	100	sexual abuse	100
child services number	72	what is child abuse	86
child protective services number	72	report child abuse	78
child protective services texas	40	child neglect	71
child abuse	23	neglect	71
report child protective services	20	child abuse and neglect	47
child protective services jobs	20	child abuse pa	47
department of child protective services	20	reporting child abuse	45
department of child services	19	child abuse hotline	40

This table gives the top 10 related queries that occur in the same search session as each of the italicized search terms. These were collected based on search interest in the U.S. in the years 2019 and 2020. Scores of 100 imply that the query is most common and interest in other queries is given relative to it.

rather than county level. I use a cross walk obtained via Kaggle (Pastor, 2021) that maps between U.S. counties and Google Trend DMAs.¹ There are a handful of counties which are not assigned to DMAs by the cross walk and I fill these in by hand, where possible.

I use this data on Apple maps requests to identify the start, mid-, and end points of the reductions in mobility within a given geographic region according to Algorithm 1. The start points represents the last peak before a substantial decrease in mobility and the end points represents a valley by which point the reduction in mobility has begun to level out (at least temporarily). The mid points are simply the half way points between the start and end dates. Figures 2.26 through 2.35 in Section 2.7 show the trends in mobility and the identified start, mid-, and end points of the mobility drops for each DMA.


Because Google Trends data is given at weekly intervals, I average the mobility levels by week. Figure 2.1 gives the frequency distribution of drop points for all DMAs by week-of-year. Though the last official stay-at-home order was implemented on April 7th (during Week 14 of the year), mobility drops had already occurred in every DMA by April 2nd (Week 13). The vast majority of mobility drops were at their midpoints in Week 11 of 2020.

Figure 2.2 plots the weekly mobility of each DMA as well as the average mobility across all DMAs, both including and excluding Week 11. It shows that regardless of lockdown timing, nearly all DMAs had begun to reduce mobility by Week 11. Week 11 is also the week in which shelter-in-place orders were first implemented in the US. In fact, for about 92 percent of DMAs, the mobility drop midpoint occurs before the eventual lockdown order and in 65 percent of cases, the mobility drop ends before the lockdown. Figure 2.3 gives the week of lockdown by state. These range from Week 10 to Week 13. In addition to mobility drops preceding lockdowns, it does not appear that week of shut down is strongly related to the size of the mobility drop. Figure 2.4 gives the average percent change in mobility levels between

¹There are two DMAs (Palm Springs, CA and Glendive, MT) which do not map to any county. In both cases, they lie within a single county that is almost entirely within another DMA and for that reason, gets assigned to the other DMA. As a result, I have mobility data for 208 of the 210 DMAs in the U.S.

Algorithm 1 Detecting Timeline of Reductions in Mobility

1. I assume that any reduction in mobility will occur between March 3rd and April 12th, and restrict to this time period for the purpose of identifying mobility reductions. Note that this period includes the dates in which shelter-in-place orders were imposed in the U.S., the WHO's global pandemic declaration, and President Trump's national emergency declaration.
2. Because this data is high-frequency, I smooth it with a Savitzky-Golay filter. I use the `savgol.filter()` function implemented in the Scipy Python package with a window length of 15, a third-order polynomial, and all other parameters set to their default values.
3. I use the peak finding function `find_peaks()` from the Scipy Python package to determine the local maxima and minima in the smoothed data using a distance parameter of 6 and all other parameters set to their default values.



4. I take the mobility associated with the maximum and minimum of the local extrema in this time period, and split the range between them into thirds to get an upper, mid- and lower tercile of mobility.
5. I find the first local minima to occur within the lower tercile and take this to be the end point of the mobility drop.
6. I then find the last local maxima in the upper mobility tercile that occurs before the mobility drop endpoint and take this to be the start point of the mobility drop.
7. Finally, I compute the midpoint of the drop halfway between the start and end points.





Figure 2.1: Timing of Mobility Drops




This figure gives the frequency distribution of start, mid-, and end drop points for all DMAs as determined by Algorithm 1 by week-of-year.

the start and end points of the mobility drops. Most states experience similar magnitudes of mobility drops despite differences in lockdown timing.

Since mobility declined prior to lockdowns and in general these declines happened almost simultaneously (at least with respect to week-of-year) across geographies, I will not leverage the staggered lockdowns, but rather take Week 11 to be the week in which mobility declines occur and compare trends in search interest before and after Week 11 in 2020 to that of prior years. That is, I use a differences-in-differences approach to compare weekly search interest in domestic-violence-related queries before and after Week 11 of 2020 relative to that of years 2016 through 2019, controlling for the year, location, and month-of-year fixed

Figure 2.2: Mobility by DMA




(a) Including Week 11

(b) Excluding Week 11

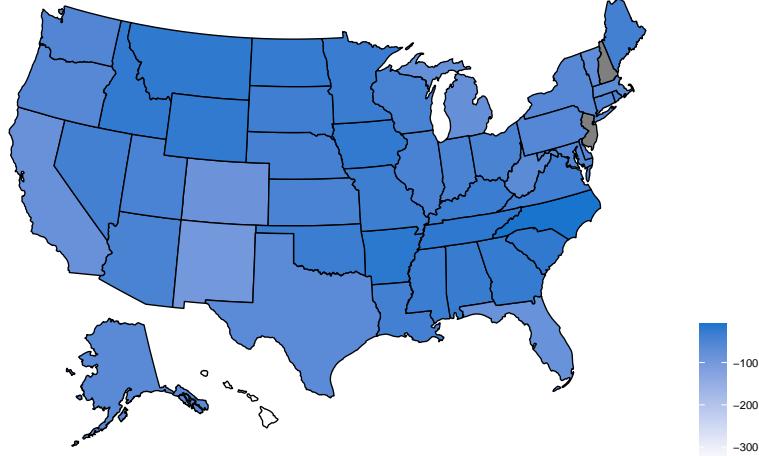
This figure shows the trends in average weekly mobility for each DMA as well as the mean weekly mobility across all DMAs, both including and excluding Week 11.

Figure 2.3: Week of Lockdown by State



This figures gives the week-of-year in which state-mandated lockdowns or shelter-in-place orders took effect in each state.

Figure 2.4: Average Percent Difference in Mobility by State



This gives the state average percent difference between the mobility level at the end and start points of the mobility drop for each DMA.

effects. Specifically, I estimate

$$\text{DV Searches}_{gwm_y} = \beta \text{PostWeek11}_w \times \text{Year2020}_y + \text{PostWeek11}_w + \theta_y + \alpha_g + \delta_m + \varepsilon_{gwy} \quad (2.1)$$

for DMA g , year y , week-of-year w , and month m . I cluster standard errors at the state level and weight observations by DMA population.

In addition, I use an event-study framework to estimate the effect over time according to the following specification:

$$\text{DV Searches}_{gwm_y} = \sum_{\tau=0}^T \beta_{\tau} \text{Week-of-Year } \tau_w \times \text{Year2020}_y + \alpha_g + \theta_y + \varepsilon_{gwy}. \quad (2.2)$$

In order to tease out mechanisms, I also examine how the effect of sheltering-in-place on domestic violence varies by levels of other variables associated with increased victim-perpetrator exposure, household bargaining, and economic distress. Specifically, I take advantage of variation in mobility reduction, COVID-19 anxiety, household size, gender employment ratios, and unemployment insurance claims rates.

I analyze how effects change by these mechanisms in two ways. First, I split the sample into quartiles of each mechanism variable of interest and estimate Equation 2.1 separately for each quartile. Where possible, I also directly interact mobility with each mechanism variable. Because I do not have mobility prior to 2020, I cannot compare pre- and post- Week 11 mobility for 2020 relative to earlier years as I do in the differences-in-differences and event study frameworks. Instead, I regress 2020 interest in each domestic violence search term on the additive inverse of mobility in 2020 (to capture exposure), the mechanism variable in 2020, and the interaction between the two, controlling for domestic violence relate search interest in 2019 to account for seasonal trends in search interest. That is, I estimate

$$\begin{aligned} [\text{2020 DV}]_{gwm} [\text{Searches}]_{gwm} &= \beta_1 [\text{Negative}]_{gwm} [\text{Mobility}]_{gwm} + \beta_2 X_{gwm} + \beta_3 [\text{Negative}]_{gwm} \times X_{gwm} \\ &+ \lambda [\text{2019 DV}]_{gwm} [\text{Searches}]_{gwm} + \alpha_g + \delta_m + \varepsilon_{gw} \end{aligned} \quad (2.3)$$

where X is a variable measuring a possible mechanism through which domestic violence might be impacted. As before, I control for location and month fixed effects, cluster standard errors at the state level, and weight observations by DMA population.

Exposure Reduction Theory

To measure the role of exposure reduction theory, I use Apple maps mobility data on the volume of requests for driving directions. I also hypothesize that greater anxiety about COVID-19 could result in an inability to reduce exposure by exiting a violent home situation due to fear of virus exposure in alternative lodgings such as a crowded shelter. I attempt to measure the salience of COVID-19 anxiety using Google Trends search interest

Table 2.2: Top Related Queries to COVID-19


<i>“COVID-19 + coronavirus”</i>	
Top Related Queries	Interest
coronavirus cases	100
coronavirus update	83
coronavirus us	59
coronavirus symptoms	55
thank you coronavirus helpers	48
coronavirus usa	46
coronavirus tips	45
coronavirus map	43
coronavirus news	41
corona	35

This table gives the top 10 related queries that occur in the same search session as “COVID-19” and “coronavirus.” These were collected based on search interest in the U.S. in the years 2019 and 2020. Scores of 100 imply that the query is most common and interest in other queries is given relative to it.

in “COVID-19” and “coronavirus.” The connection between Google searches and salience of media coverage has been validated for other viruses. Towers et al. (2015) show that more than 65 percent of the variation in Ebola-related Google searches from the US can be explained by variation in Ebola-related news media coverage, for example. This should allow me to capture the role that fear of COVID-19 exposure plays in preventing women from leaving abusive homes or seeking alternative resources to cope with domestic violence. Table 2.2 gives the top 10 searches associated with either COVID-19 or coronavirus. These include searches for case numbers, updates, and symptom information. If COVID-19 intensifies exposure, I would expect stronger impacts on domestic violence related search interest in areas with more COVID-19-related searches. Figure 2.5 plots weekly mobility across the U.S. in 2020 against search interest in “COVID-19” and/or “coronavirus” in both 2020 and the average of earlier years by week-of-year. Unsurprisingly, there is essentially no search interest in either “COVID-19” or “coronavirus” prior to 2020. Note that search interest in these words peaks in Week 11 just as mobility is mid-drop in nearly all DMAs.

I also look at how the effect of sheltering-in-place varies for regions with different average household densities. To do this, I use data on household structure by state from the 2020

Figure 2.5: Search Interest in COVID-19/Coronavirus and Mobility Over Time



This figure depicts weekly relative search interest over time in “COVID-19” or “coronavirus” for the US in 2019 and 2020 alongside the average weekly mobility in 2020. The red vertical line (Week 11) is the week in which the first official lockdown occurred, at which point most DMAs had already begun to reduce their mobility.

Annual Social and Economic Supplement of the Current Population Study, obtained through the Integrated Public Use Microdata Series platform. I use this to compute the average number of individuals per household in each state. In theory, greater household density could intensify the exposure between victim and perpetrator due to crowded conditions and lead to more violence.

Bargaining vs. Male Backlash

In order to tease out the impact of household bargaining via gender-specific earning potential, I use US Census Bureau data from the Basic Monthly Current Population Survey. I take the state-level female employment-to-population ratio and divide it by the male employment-to-population ratio to evaluate the employment of women relative to men in a each state. If the effect of COVID-19 lockdowns on domestic violence is stronger in areas where women are less likely to be employed relative to men, then we may be seeing evidence of decreased

bargaining power. Conversely, if effects are stronger where women are more likely to be employed, this would be consistent with male backlash.

Financial Stress

To determine how these effects vary by levels of financial distress, I will use data on the weekly unemployment insurance claims rate by state provided by the United States Department of Labor.

Probability of Detection

To get a sense of how the pandemic may effect probability of detecting abuse via third-party reporting, I also evaluate the effect of sheltering-in-place on search interest in terms that indicate concern about child abuse. Specifically, I look at Google Trends search interest in the terms “child protective services” and “child abuse.’ Table 2.1 shows the top 10 search queries related to these terms. These searches are consistent with third-party reporting behavior. They include queries like “report child abuse” and “child services number.” I will compare changes in search interest in these terms with that of the other domestic violence related terms, which are more likely to be searched directly by a victim. If third-party reporters such as school teachers are less able to observe signs of abuse, they may have less reason to search for the relevant agency or reporting mechanism. This would be consistent with the notion that lockdowns interfere with abuse identification and prevention through third parties.

Other Data

Finally, I use county-level data from the 2020 American Community Survey on whether the primary member of the household is living alone or with a spouse or romantic partner to compute the percent of households in a DMA in which the householder is living alone and the percent living with a romantic partner. I use this to check that changes in domestic violence search patterns are stronger in areas where individuals are more likely to live with

a romantic partner and less likely to live alone as we would expect if these effects are driven by domestic violence.

Because some DMAs span multiple states, for merges with state-level data I assign the DMA to the state that houses the largest share of the DMA population. To do this, I use 2020 estimates of county population provided by the U.S. Census bureau, aggregated to the DMA-level.

2.4 Results

Table 2.3 shows the estimates of Equation 2.1 for the five outcomes of interest in the first 21 weeks of the year. While there is a large and statistically significant increase in search interest for domestic violence hotlines, searches for domestic violence help, domestic violence shelters, and protective orders actually decrease. This may reflect the different motivations behind these searches. Consider the decision to call a domestic violence hotline. This may be the first step that a victim takes and may come before she is ready for more disruptive action. Hotlines provide tips and strategies for staying safe even for victims who remain in violent situations. Compare this to the decision to seek refuge in a domestic violence shelter or to obtain a protective order. These actions preclude remaining in one's home with an abuser. If concerns about COVID-19 raise the barriers to leaving an abusive home, we may see a decrease in these more direct responses to violence, at least in the short term.

Figure 2.9 shows that as time goes by post-lockdowns, search interest in restraining and protective orders do catch up to pre-2020 levels. Moreover, protective order interest exceeds pre-2020 levels midway through 2020 and there is a notable spike in search interest in restraining orders in the later part of 2020 as well. Given that restraining orders involve legal intervention and require victims to leave their abusers, it make take longer before domestic

violence survivors are ready to pursue them. It may also not be surprising that searches for shelters are less appealing in the midst of a viral pandemic when individuals are concerned about the public health implications of living in a crowded shelter. Note also that related queries to “domestic violence help” include “domestic violence housing,” and “domestic violence shelter.” It’s possible that this too is more likely to be searched by those ready for next steps when violence has already escalated. Figure 2.7 shows that while searches for domestic violence help fall below pre-2020 interest levels in the initial weeks after sheltering-in-place begins, they too rise above pre-2020 levels after about 10 weeks. In contrast, while search interest in domestic violence hotlines remains above pre-2020 levels well into the year, the difference between 2020 and pre-2020 levels is much smaller after the first 10 weeks post Week 11. Thus, these results could be consistent with an immediate increase in more passive strategies to survive domestic violence without leaving one’s home and an eventual substitution toward measures that involve more direct action such as seeking a restraining order.

The magnitudes of the effects are quite large. Specifically, I estimate a 30 percent increase in domestic violence hotline search interest, an 19 percent decrease in domestic violence help search interest, a 16 percent decrease in shelter interest, a 15 percent decrease in restraining/protective order search interest, and a 32 percent decrease in searches indicating concern for children. The results are not sensitive to whether I look at 2020 search interest relative to search interest in years 2016 through 2019 or just relative to 2019.


Figures 2.11, 2.12, 2.13, and 2.14 showcase the event-study estimates of Equation 2.2 on domestic violence search terms in 2020 relative to the four years prior. I have included the average value of the relative difference in search interest between 2020 and earlier years in the weeks prior to Week 11, for reference. Figure 2.11 shows that differences for domestic violence hotline interest hover around this average in the early weeks of the year and then

Table 2.3: Effect of Staying-at-Home on Domestic Violence Related Search Behavior

	(1) DV Hotline	(2) DV Help	(3) Shelter	(4) Order	(5) Child Concern
A. 2020 Relative to All Other Years					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2020	5.3554*** (1.8619)	-2.0438* (1.0393)	-2.3772*** (0.6994)	-3.4231*** (1.0484)	-9.6689*** (1.2967)
<i>Inverse Hyperbolic Sine of Search Interest</i>					
Post Week 11 x Year 2020	0.3017*** (0.0853)	-0.1863* (0.1005)	-0.1630** (0.0622)	-0.1538*** (0.0382)	-0.3255*** (0.0271)
B. 2020 Relative to 2019					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2020	7.3204*** (2.4890)	-2.8023* (1.5436)	-3.3024*** (0.8925)	-4.4102*** (1.2938)	-8.8114*** (1.2637)
<i>Inverse Hyperbolic Sine of Search Interest</i>					
Post Week 11 x Year 2020	0.3399*** (0.1097)	-0.2509 (0.1699)	-0.1962*** (0.0637)	-0.2044*** (0.0372)	-0.2979*** (0.0299)

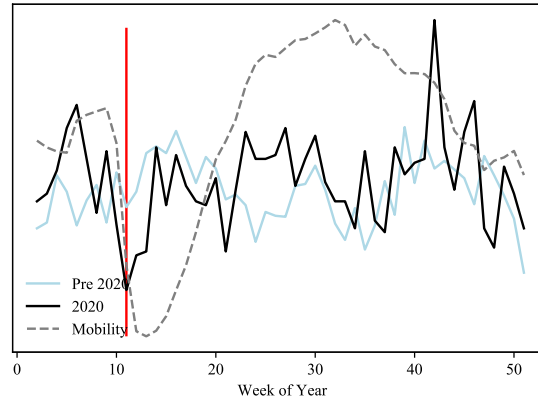
The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), the average of search interest in “protective order” and “restraining order” in column (4), and the average of “child protective services” and “child abuse” in column (5). Observations are at the DMA-week level and are weighted by DMA populations. These regressions include Weeks 1 through 21 of 2016 through 2020 in Panel A and Weeks 1 through 21 of 2019 and 2020 in Panel B. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.
 $*p < 0.1$, $**p < 0.05$, $***p < 0.01$

Figure 2.6: Search Interest in Domestic Violence Hotlines and Mobility Over Time



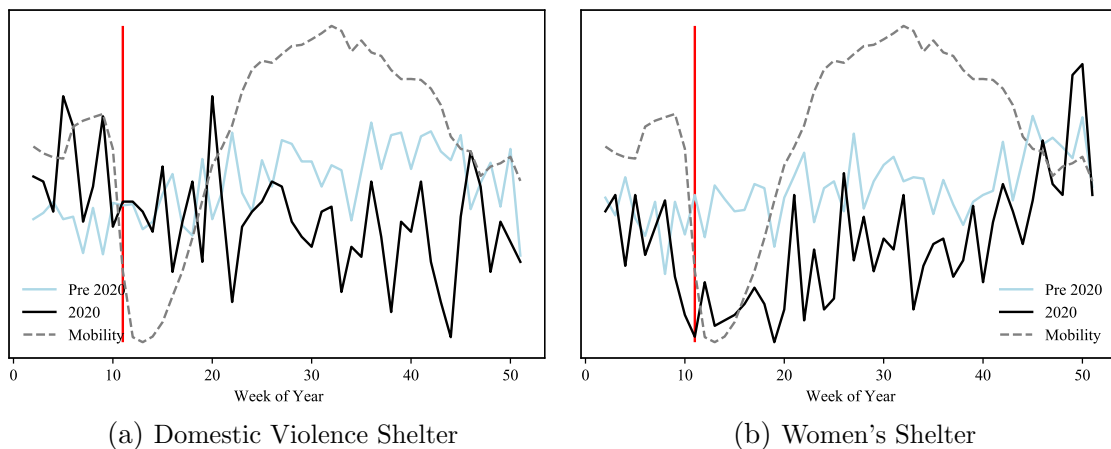
This figure depicts average relative search interest in “domestic violence hotline” over week-of-year for the US in 2016 through 2019 (Pre 2020) and 2020 alongside the average weekly mobility in 2020. The red vertical line (Week 11) is the week in which the first official lockdown occurred, at which point most DMAs had already begun to reduce their mobility.

Figure 2.7: Search Interest in Domestic Violence Help and Mobility Over Time




This figure depicts average relative search interest in “domestic violence help” over week-of-year for the US in 2016 through 2019 (Pre 2020) and 2020 alongside the average weekly mobility in 2020. The red vertical line (Week 11) is the week in which the first official lockdown occurred, at which point most DMAs had already begun to reduce their mobility.

Figure 2.8: Search Interest in Shelters and Mobility Over Time



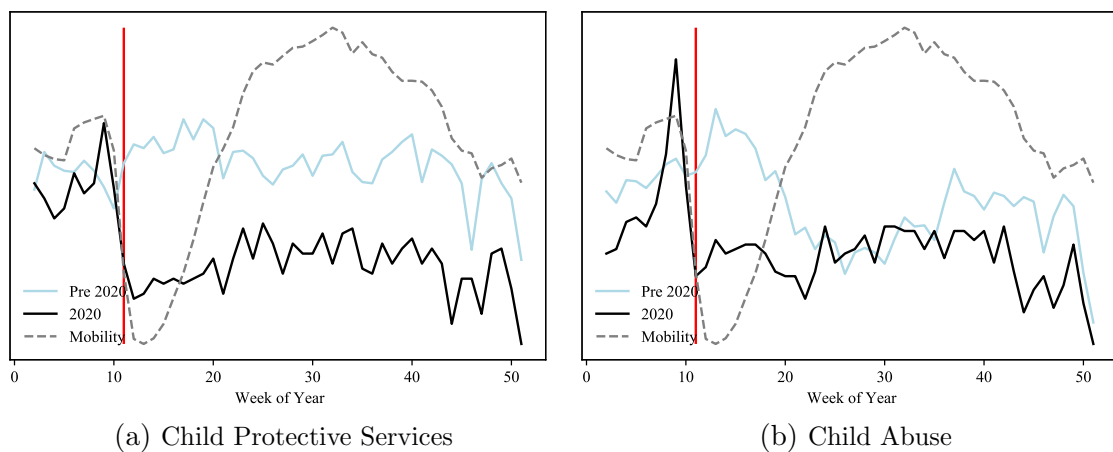
This figure depicts average relative search interest in “domestic violence shelter” and “womens shelter” over week-of-year for the US in 2016 through 2019 (Pre 2020) and 2020 alongside the average weekly mobility in 2020. The red vertical line (Week 11) is the week in which the first official lockdown occurred, at which point most DMAs had already begun to reduce their mobility.

Figure 2.9: Search Interest in Restraining/Protective Orders and Mobility Over Time




This figure depicts average relative search interest in “restraining order” and “protective order” over week-of-year for the US in 2016 through 2019 (Pre 2020) and 2020 alongside the average weekly mobility in 2020. The red vertical line (Week 11) is the week in which the first official lockdown occurred, at which point most DMAs had already begun to reduce their mobility.

Figure 2.10: Search Interest in Child Protective Services/Child Abuse and Mobility Over Time



This figure depicts average relative search interest in “child protective services” and “child abuse” over week-of-year for the US in 2016 through 2019 (Pre 2020) and 2020 alongside the average weekly mobility in 2020. The red vertical line (Week 11) is the week in which the first official lockdown occurred, at which point most DMAs had already begun to reduce their mobility.

Figure 2.11: Event Study - Search Interest in Domestic Violence Hotlines




This figure gives the estimates of the event study model described in Equation 2.2 for the first 21 weeks of the year where the outcome variable is search interest in “domestic violence hotline.” The red line (Week 11) indicates the week by which most DMAs had already begun to decrease their mobility. The black horizontal line indicates the average level of the week-of-year coefficients prior to Week 11.

jump up during the initial weeks of sheltering in place. The opposite is true for search interest in shelters and protective/restraining orders.


Perhaps the most striking results are for search interest indicating concern about potential child abuse. Table 2.3 shows that search interest in “child abuse” and “child protective services” plummets with the onset of sheltering-in-place. Figure 2.15, which depicts event study estimates for this outcome underlines this point. While search interest in these terms appears to be trending upward in the early part of 2020 relative to earlier years, the onset of sheltering-in-place results in a dramatic drop in search interest and the effect continues through Week 21 of 2020. Figures 2.10a and 2.10b show that U.S. search interest in these terms remains largely below earlier levels through the rest of 2020. Since many schools did not reopen until 2021, this is entirely consistent with lower search rates by third-party reporters who were not able to identify potential signs of abuse while students were learning remotely.

Figure 2.12: Event Study - Search Interest in Domestic Violence Help




This figure gives the estimates of the event study model described in Equation 2.2 for the first 21 weeks of the year where the outcome variable is search interest in “domestic violence help.” The red line (Week 11) indicates the week by which most DMAs had already begun to decrease their mobility. The black horizontal line indicates the average level of the week-of-year coefficients prior to Week 11.

Figure 2.13: Event Study - Search Interest in Shelters




This figure gives the estimates of the event study model described in Equation 2.2 for the first 21 weeks of the year where the outcome of interest is the average of search interest in “women’s shelter” and “domestic violence shelter.” The red line (Week 11) indicates the week by which most DMAs had already begun to decrease their mobility. The black horizontal line indicates the average level of the week-of-year coefficient prior to Week 11.

Figure 2.14: Event Study - Search Interest in Protective/Restraining Orders.



This figure gives the estimates of the event study model described in Equation 2.2 for the first 21 weeks of the year where the outcome of interest is the average of search interest in “restraining order” and “protective order.” The red line (Week 11) indicates the week by which most DMAs had already begun to decrease their mobility. The black horizontal line indicates the average level of the week-of-year coefficient prior to Week 11.

Figure 2.15: Event Study - Search Interest in Child Protective Services/Child Abuse



This figure gives the estimates of the event study model described in Equation 2.2 for the first 21 weeks of the year where the outcome of interest is the average of search interest in “child protective services” and “child abuse.” The red line (Week 11) indicates the week by which most DMAs had already begun to decrease their mobility. The black horizontal line indicates the average level of the week-of-year coefficient prior to Week 11.

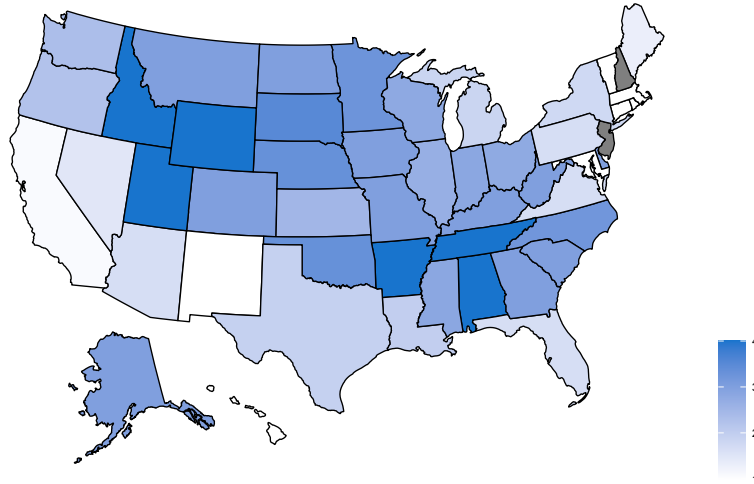
To shed light on the other mechanisms that may be at play here, I examine how these effects vary in places with differing levels of mobility reduction, COVID-19 anxiety, average number of individuals per household, relative female-to-male employment ratios, and unemployment insurance claims rates. I do this both by examining the effects of sheltering-in-place on domestic violence related search interest separately for each quartile of a given mechanism variable, and by directly interacting mobility with the mechanisms variables that vary in 2020 (COVID-19 related search interest, female-to-male employment ratios, and unemployment insurance claims rates).

Exposure Reduction Theory

Table 2.4 gives the effect of sheltering-in-place across quartiles of mobility drops. The size of the mobility drops were computed by taking the percent difference between average of mobility in the 10 weeks after Week 11 and average mobility in the 10 weeks prior to Week 11 for each DMA. The percent difference in mobility ranges from a 70.13 percent decrease in mobility to an 8.90 percent increase in mobility. Figure 2.16 gives the average quartile assigned to DMAs in each state. Effects are generally largest (farthest from zero) in the lowest quartile where mobility drops are largest (most negative). This is consistent with exposure reduction theory. In areas where households reduce their mobility more dramatically and therefore spend much more time together than usual, increases in domestic violence hotline search interest are more pronounced. Reductions in search interest for other domestic violence resources are also more pronounced in these areas.

Columns (1), (3), (5), and (7) of Table 2.7 show that search interest in domestic violence hotlines is increasing in exposure (negative mobility) as we would expect under exposure reduction theory. Search interest in domestic violence help, shelters, and restraining/protective orders however, is decreasing in increased exposure. Again, this may be consistent with a delay between initial lockdowns (when mobility dramatically decreases) and when domestic violence escalates to the point where survivors are ready for these more disruptive resources,

Figure 2.16: Average Percent Change in Mobility Quartile in Each State



This figure reflects the average quartile assigned within a state based on the percent difference in mobility in the state’s DMAs. The percent difference in mobility is calculated by taking the percent difference between average mobility in the 10 weeks before and after Week 11 for each DMA. DMAs are assigned a value between 1 and 4 reflecting the mobility difference quartile they fall in with 1 describing DMAs that fall below the 25th percentile and 4 describing DMAs that fall above the 75th percentile.

by which point mobility has begun to increase again. It also may reflect that victims are reluctant to leave their homes during lockdowns and instead wait until lockdowns begin to lift.

Table 2.5 gives the effect of sheltering-in-place across quartiles of average search interest in “COVID-19” or “coronavirus” in the 10 weeks after Week 11. Figure 2.17 depicts the average quartile assigned to DMAs within each state. Effects are generally largest in areas with the highest level of COVID salience indicating that fear of the virus may intensify victim-perpetrator exposure. Panel A of Table 2.7 gives estimates of Equation 2.3 for COVID related search interest. Though small, the interaction term between exposure (negative mobility) and COVID related search interest on search interest in domestic violence hotlines is positive and


Table 2.4: Effects of Sheltering-In-Place by Percent Difference in Mobility Quartiles

	(1) DV Hotline	(2) DV Help	(3) Shelters	(4) Orders
Quartile 1 % Difference in Mobility of -70.13 to -21.16				
Post Week 11 x Year 2020	9.2661*** (3.1580)	-4.9254*** (1.5652)	-2.3323 (1.3574)	-4.9901** (1.9605)
Mean of Outcome	5.7804	3.5707	6.2821	17.5754
Quartile 2 % Difference in Mobility of -21.16 to -12.77				
Post Week 11 x Year 2020	2.4860*** (0.8067)	1.9132* (0.9364)	-1.7752 (1.1483)	-2.6362** (0.9479)
Mean of Outcome	1.4339	0.7612	2.8119	9.7206
Quartile 3 % Difference in Mobility of -12.77 to -8.80				
Post Week 11 x Year 2020	1.6489 (1.4173)	-1.0397 (0.7749)	-2.8960** (1.3766)	-0.8170 (1.4100)
Mean of Outcome	0.8960	0.4487	1.9620	8.3185
Quartile 4 % Difference in Mobility of -8.80 to 8.90				
Post Week 11 x Year 2020	1.6539 (1.1000)		-2.8196** (1.1707)	-2.8793*** (0.8123)
Mean of Outcome	0.2511	0.0000	0.9767	6.5054

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for each quartile of the distribution of the percent difference in mobility. The percent difference here is taken between average mobility in the 10 weeks before and after Week 11 for each DMA. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), and the average of search interest in “protective order” and “restraining order” in column (4). Observations are at the DMA-week level and are weighted by DMA populations. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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

Figure 2.17: Average Quartile of Post Period Covid Search Interest in Each State



This figure reflects the average quartile of search interest in “COVID-19” or “coronavirus” within each state. I take the average level of COVID related search interest in the 10 weeks before and after Week 11 for each DMA and assign a value between 1 and 4 with 1 describing DMAs that fall below the 25th percentile and 4 describing DMAs that fall above the 75th percentile.

statistically significant indicating that the impact is greater in areas where COVID anxiety is higher. There appears to be little impact on the other outcomes of interest.

Table 2.6 gives the effect of sheltering-in-place across quartiles of the state-level average number of individuals per household in 2020. Figure 2.18 depicts the quartile each state falls in. Effects are generally larger in areas with more individuals per household. This is consistent with exposure reduction theory as it is harder to avoid a violent perpetrator in a more crowded household.

There appears to be strong evidence that exposure reduction contributes to the effects of sheltering-in-place on domestic violence search interest as effects are strongest in areas with

Table 2.5: Effects of Sheltering-In-Place by Post Period COVID Search Interest Quartiles

	(1) DV Hotline	(2) DV Help	(3) Shelters	(4) Orders
Quartile 1 <i>Post Period Covid Search Interest of 0 to 39.91</i>				
Post Week 11 x Year 2020	0.7509 (1.2249)	-0.5265 (0.5029)	-2.3011** (0.9232)	-4.4746*** (1.0353)
Mean of Outcome	0.9524	0.1910	1.6251	8.9696
Quartile 2 <i>Post Period Covid Search Interest of 39.91 to 42.09</i>				
Post Week 11 x Year 2020	1.6074 (1.1756)	-0.5399 (0.5681)	-0.6839 (0.6527)	-2.3088*** (0.6905)
Mean of Outcome	1.8007	0.4474	2.6471	10.1950
Quartile 3 <i>Post Period Covid Search Interest of 42.09 to 44.45</i>				
Post Week 11 x Year 2020	4.8631 (3.5727)	-2.0059 (2.5858)	-3.2820*** (1.1085)	-2.3929* (1.2633)
Mean of Outcome	2.6823	2.0079	4.4164	12.3468
Quartile 4 <i>Post Period Covid Search Interest of 44.45 to 53.36</i>				
Post Week 11 x Year 2020	11.9670*** (3.6837)	-4.2928** (1.6882)	-2.7259 (1.7236)	-4.8526* (2.3762)
Mean of Outcome	2.8041	2.0512	3.1784	10.2226

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for each quartile of the distribution of average COVID search interest in the 10 weeks after Week 11 at the DMA level. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), and the average of search interest in “protective order” and “restraining order” in column (4). Observations are at the DMA-week level and are weighted by DMA populations. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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


Table 2.6: Effects of Sheltering-In-Place by Quartiles of Average 2020 Household Size

	(1) DV Hotline	(2) DV Help	(3) Shelters	(4) Orders
Quartile 1 <i>Average 2020 Household Size of 2.17 to 2.34</i>				
Post Week 11 x Year 2020	-0.3290 (1.1017)	-1.0089* (0.5398)	-4.0830*** (1.0101)	-3.1502* (1.5686)
Mean of Outcome	1.4109	0.4817	2.2461	8.4195
Quartile 2 <i>Average 2020 Household Size of 2.34 to 2.39</i>				
Post Week 11 x Year 2020	1.1430 (0.9177)	-3.4542 (3.0384)	-1.1294 (0.9041)	-3.1104*** (0.9312)
Mean of Outcome	1.2998	0.2995	2.6485	9.2200
Quartile 3 <i>Average 2020 Household Size of 2.39 to 2.49</i>				
Post Week 11 x Year 2020	6.9439* (3.7225)	-1.1217 (1.3228)	-2.7594** (0.9852)	-0.8687 (1.0948)
Mean of Outcome	2.1409	1.6580	2.7027	9.8196
Quartile 4 <i>Average 2020 Household Size of 2.49 to 3.00</i>				
Post Week 11 x Year 2020	9.7941** (3.3672)	-2.4355 (2.1353)	-1.9814 (1.5840)	-5.9345** (2.0267)
Mean of Outcome	3.3744	2.2464	4.2329	14.1016

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for each quartile of the distribution of average 2020 household size. Average 2020 household size is provided at the state level from the Annual Social and Economic Supplement of the CPS. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), and the average of search interest in “protective order” and “restraining order” in column (4). Observations are at the DMA-week level and are weighted by DMA populations. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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

Figure 2.18: 2020 State Average Household Size Quartiles



This figure reflects the quartiles of the distribution of the average number of individuals per household in 2020 at the state level. States are assigned a value between 1 and 4 where 1 describes states that fall below the 25th percentile and 4 describes states that fall above the 75th percentile.

Table 2.7: Relationship Between Domestic Violence Related Search Interest and Potential Mechanisms Directly

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	DV Hotline	DV Hotline	DV Help	DV Help	Shelters	Shelters	Orders	Orders
<i>A. Covid Search Interest</i>								
Negative Mobility	0.0291*	-0.0116	-0.0066	0.0029	-0.0114**	-0.0138	-0.0129	0.0033
	(0.0172)	(0.0141)	(0.0145)	(0.0098)	(0.0045)	(0.0085)	(0.0081)	(0.0125)
Covid Searches		0.2300***		-0.0596		0.0088		-0.1125***
		(0.0571)		(0.0622)		(0.0547)		(0.0385)
Negative Mobility X Covid Searches		0.0024***		-0.0005		0.0002		-0.0008*
		(0.0006)		(0.0005)		(0.0005)		(0.0004)
<i>B. Female/Male Employment Ratio</i>								
Negative Mobility	0.0291*	0.0251	-0.0066	-0.1596	-0.0114**	-0.0102	-0.0129	-0.2162***
	(0.0172)	(0.2189)	(0.0145)	(0.1127)	(0.0045)	(0.0656)	(0.0081)	(0.0741)
Female/Male Emp		10.9943		34.5375		-1.3884		40.3154***
		(42.2679)		(26.3806)		(12.0100)		(13.6595)
Negative Mobility X Female/Male UE		0.0074		0.1743		-0.0005		0.2319***
		(0.2370)		(0.1313)		(0.0732)		(0.0856)
<i>C. UI Claims Rate</i>								
Negative Mobility	0.0291*	-0.0317*	-0.0066	-0.0270*	-0.0114**	-0.0083	-0.0129	-0.0135
	(0.0172)	(0.0178)	(0.0145)	(0.0146)	(0.0045)	(0.0087)	(0.0081)	(0.0146)
UI Claims Rate		1.0888**		0.2742		-0.0821		0.0886
		(0.5237)		(0.2281)		(0.1053)		(0.1954)
Negative Mobility X UI Claims Rate		0.0064*		0.0025		0.0001		-0.0006
		(0.0032)		(0.0020)		(0.0006)		(0.0011)

This table gives the coefficients of interest from Equation 2.3. The outcome variable is search interest in “domestic violence hotline” in columns (1) and (2), “domestic violence help” in columns (3) and (4), the average of search interest in “domestic violence shelter” and “women’s shelter” in columns (5) and (6), and the average of search interest in “protective order” and “restraining order” in column (7) and (8). Observations are at the DMA-week level and are weighted by DMA populations. They include DMA and month-of-year fixed effects. Standard errors are clustered at the state level.

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

the largest declines in mobility. Anxiety about the virus and more densely populated households exacerbate this effect.

Bargaining vs. Male Backlash

Table 2.8 gives the effect of sheltering-in-place across quartiles of female-to-male employment-to-population ratios in the 10 weeks prior to Week 11. These ratios vary from 0.81 to 0.99. Figure 2.19 gives the quartile each state falls in. Effects tend to be slightly larger for DMAs in states with lower female-to-male employment ratios which is consistent with household bargaining theory in which women with better employment prospects are able to negotiate less violence via a more credible threat of leaving. Panel B of Table 2.7 gives estimates of Equation 2.3 for female-to-male employment ratios. The interaction term between exposure (negative mobility) and female-to-male employment is associated with a statistically significant increase in search interest in restraining/protective orders. An increase in abuse for these women is consistent with male backlash in which an abuser, emasculated by a woman's independent earning potential, seeks to establish dominance through violence. However, given that there is not a statistically significant impact on other outcomes, I wonder if restraining orders are a tool that are more likely to be used by empowered women who are also more likely to work. Consider that they are typically enacted via court order and obtaining one requires navigation of the legal system. There is some evidence that in areas where women who are more likely to work, the effects of COVID-19 on domestic violence related search interest was smaller, which is consistent with household bargaining theory.

Financial Stress

Table 2.9 gives the effect of sheltering-in-place across quartiles of the percent difference in state unemployment insurance claims rates in the 10 weeks before and after Week 11. Figure 2.20 gives the quartile each state falls in. There is not a clear pattern with regard to effect sizes across different quartiles. Panel C of Table 2.7 gives estimates of Equation 2.3 for unemployment insurance claims rates. There is a marginally significant but small positive

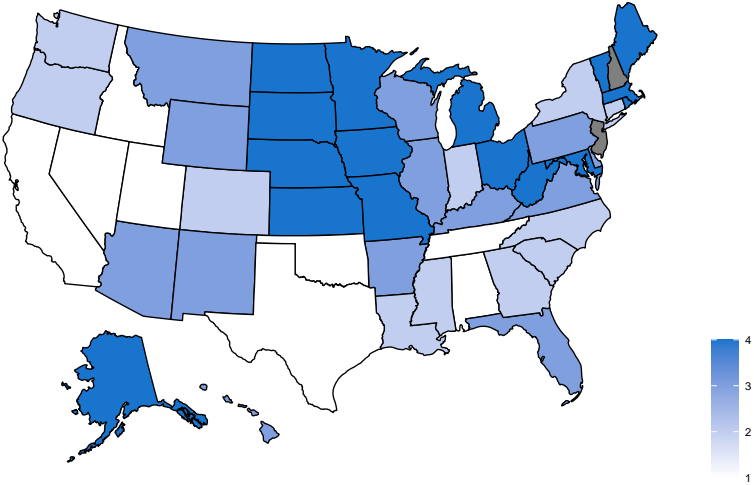
Table 2.8: Effects of Sheltering-In-Place by Quartiles of Pre-Period Female-to-Male Employment Ratios

	(1)	(2)	(3)	(4)
	DV Hotline	DV Help	Shelters	Orders
Quartile 1 <i>Pre-Period Ratio of Female-to-Male Employment of 0.81 to 0.86</i>				
Post Week 11 x Year 2020	7.6452*	-2.0868	-1.2948	-6.1607**
	(3.9737)	(2.4045)	(1.3822)	(2.1669)
Mean of Outcome	2.4416	1.6229	3.4786	13.6313
Quartile 2 <i>Pre-Period Ratio of Female-to-Male Employment of 0.86 to 0.89</i>				
Post Week 11 x Year 2020	5.4802	-1.4969	-2.8674***	-1.8912
	(4.3822)	(0.8466)	(0.7867)	(1.1427)
Mean of Outcome	2.0127	1.2335	2.8101	10.2438
Quartile 3 <i>Pre-Period Ratio of Female-to-Male Employment of 0.89 to 0.93</i>				
Post Week 11 x Year 2020	3.0670**	-2.8656	-2.9741*	-0.7917
	(1.2811)	(2.8664)	(1.3848)	(1.0203)
Mean of Outcome	2.2156	1.5057	3.1254	10.7986
Quartile 4 <i>Pre-Period Ratio of Female-to-Male Employment of 0.93 to 0.99</i>				
Post Week 11 x Year 2020	4.8897	-1.6035	-2.4969	-5.1067***
	(3.6648)	(1.3581)	(1.6476)	(1.2062)
Mean of Outcome	1.6592	0.4393	2.5633	7.4787

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for each quartile of the distribution of female-to-male employment ratios in the 10 weeks prior to Week 11. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), and the average of search interest in “protective order” and “restraining order” in column (4). Observations are at the DMA-week level and are weighted by DMA populations. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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

Figure 2.19: Pre-Period Female-to-Male State Employment Ratio Quartiles



This figure reflects the quartiles of the distribution of the percent difference in female-to-male employment ratios in the 10 weeks before and after Week 11 in each state. Values of 1 describe states that fall below the 25th percentile and values of 4 describe states that fall above the 75th percentile.

relationship between domestic violence hotline search interest and the interaction of negative mobility and unemployment insurance claims rates. The apparently small relationship between changes in unemployment insurance claims and COVID related search interest may be due to the fact that the Coronavirus Aid, Relief, and Economic Security (CARES) Act was passed in March 2020 and provided exceptionally generous unemployment benefits during this period. Because of this, I also examine how the effects vary by pre-period unemployment insurance claims rates to see if the effects are stronger in areas that were more economically vulnerable prior to the pandemic. Table 2.10 gives the effect of sheltering-in-place across quartiles of the unemployment insurance claims rates in the 10 weeks prior to Week 11. Figure 2.21 depicts the quartile each state falls in. In states where unemployment was higher in the weeks leading up to the pandemic, the effect of COVID on domestic violence hotline search interest is strongest, indicating that greater financial distress leaves individuals more vulnerable to domestic violence.

Probability of Detection

As discussed above, Table 2.3 and Figure 2.15 show a dramatic decrease in search interest in child abuse and child protective services at the onset of the COVID-19 pandemic. This is consistent with a reduction in third-party reporting behavior due to decreased contact between children and potential third-party reporters. It is plausible that increases in domestic violence may be related to reduced detection from third-party reporters.

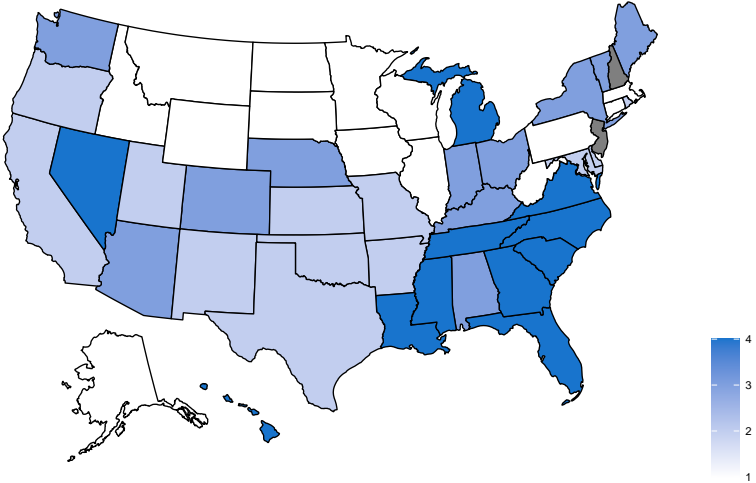
Table 2.9: Effects of Sheltering-In-Place by Quartiles of the Percent Difference in Unemployment Insurance Claims Rates

	(1) DV Hotline	(2) DV Help	(3) Shelters	(4) Orders
Quartile 1 <i>% Difference in UI Claims Rates of 308 to 539</i>				
Post Week 11 x Year 2020	5.4236 (3.6522)	-4.2463 (4.1396)	0.3237 (0.9991)	-2.7645 (1.6730)
Mean of Outcome	1.5731	0.9501	2.1237	7.0298
Quartile 2 <i>% Difference in UI Claims Rates of 539 to 752</i>				
Post Week 11 x Year 2020	7.5715* (3.5980)	-1.9486 (2.2162)	-1.6283 (1.4131)	-6.0729** (1.9930)
Mean of Outcome	2.8742	1.5004	3.7098	12.9887
Quartile 3 <i>% Difference in UI Claims Rates of 752 to 886</i>				
Post Week 11 x Year 2020	7.5021 (4.1635)	-1.2734 (0.8348)	-3.0963** (1.0193)	-1.0711 (1.3578)
Mean of Outcome	2.0738	1.6454	3.1656	10.2657
Quartile 4 <i>% Difference in UI Claims Rates of 886 to 2,238</i>				
Post Week 11 x Year 2020	0.9434 (1.3470)	-1.3751** (0.5725)	-4.3639*** (0.9014)	-2.9653** (1.0718)
Mean of Outcome	1.7671	0.7217	2.8974	11.2087

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for each quartile of the distribution of the percent difference in unemployment insurance claims rates. The percent difference here is given between the average unemployment insurance claims rates in the 10 weeks before and after Week 11. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), and the average of search interest in “protective order” and “restraining order” in column (4). Observations are at the DMA-week level and are weighted by DMA populations. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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

Figure 2.20: Quartiles of the Percent Difference in State Unemployment Insurance Claims Rates



This figure reflects the quartiles of the percent difference in unemployment insurance claims rates assigned to each state. I take the percent difference in the average unemployment insurance claims rates in the 10 weeks before and after Week 11. States are assigned a value between 1 and 4 where 1 describes states that fall below the 25th percentile and 4 describes states that fall above the 75th percentile.

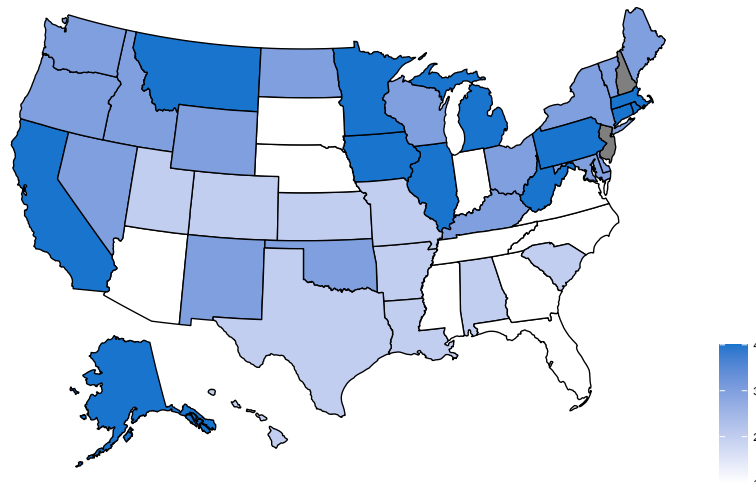
Table 2.10: Effects of Sheltering-In-Place by Quartiles of the Pre-Period Unemployment Insurance Claims Rate

	(1) DV Hotline	(2) DV Help	(3) Shelters	(4) Orders
Quartile 1 <i>Pre-Period UI Claims Rate of 0.39 to 0.78</i>				
Post Week 11 x Year 2020	2.6358 (1.5021)	-1.6334** (0.5729)	-4.9424*** (0.9106)	-1.6748* (0.9083)
Mean of Outcome	2.1011	1.0440	3.4074	11.8788
Quartile 2 <i>Pre-Period UI Claims Rate of 0.78 to 1.09</i>				
Post Week 11 x Year 2020	1.7477 (1.2510)	0.3880 (0.8766)	-3.0492*** (0.8465)	-3.2594*** (0.5289)
Mean of Outcome	1.3711	0.5972	2.3442	10.1008
Quartile 3 <i>Pre-Period UI Claims Rate of 1.09 to 1.86</i>				
Post Week 11 x Year 2020	7.0323 (4.1337)	-0.7692 (0.8593)	-2.3410** (0.8351)	-1.7102 (1.3763)
Mean of Outcome	1.5832	1.1064	2.5756	9.4966
Quartile 4 <i>Pre-Period UI Claims Rate of 1.86 to 2.95</i>				
Post Week 11 x Year 2020	8.7559** (3.8628)	-5.0478* (2.5379)	0.1495 (1.0846)	-6.2946** (2.2841)
Mean of Outcome	3.1506	1.9330	3.5460	10.3144

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for each quartile of the distribution of the average unemployment insurance claims rate in the 10 weeks prior to Week 11. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), and the average of search interest in “protective order” and “restraining order” in column (4). Observations are at the DMA-week level and are weighted by DMA populations. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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

Figure 2.21: Quartiles of Pre-Period Unemployment Insurance Claims Rates



This figure reflects the quartiles of the unemployment insurance claims rate in the 10 weeks prior to Week 11 for each state. States are assigned a value between 1 and 4 where 1 describes states that fall below the 25th percentile and 4 describes states that fall above the 75th percentile.

2.5 Some Checks

A crucial assumption of my differences-in-differences identification strategy is that I can use search interest in the years prior to 2020 as controls for 2020 search patterns. This requires that there were no major changes in search interest before and after Week 11 of prior years. To test this assumption, I estimate Equation 2.1 for each year in turn, comparing it only with the prior year. For example, I estimate the difference in search interest before and after Week 11 of 2019 relative to 2018. Table 2.11 gives the results of these regressions. There are some strongly significant effects in 2019 relative to 2018, but not in other years. To ensure that my results are not being driven by whatever may be happening in 2019, I estimate Equation 2.1 relative to all years except 2019. The results are provided in Table 2.12, and are very similar to estimates that include 2019 as a control year.

Table 2.13 shows that results are not sensitive to adding a time trend to Equation 2.1.

Table 2.14 shows that results are not sensitive to clustering at the DMA rather than the state level.

Table 2.15 provides results including the entirety of each year, rather than just the first 21 weeks as I do in my main results. Effects are slightly attenuated in this case as the effect appears to diminish over time, but do not change substantially.

Table 2.16 shows the impacts of sheltering-in-place on two placebo outcomes: search interest in “toad” and “toothache.” While I do not expect to find an increase in toad search interest, I chose the word toothache because at the onset of the COVID-19 pandemic, dental offices were initially closed and it is likely that individuals would need to search for information about how to deal with any tooth pain at home. This represents substitution away from in-person resources available before sheltering-in-place toward resources available from home. If we believe that women were substituting toward domestic violence hotlines to replace social

Table 2.11: Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Year-to-Year Comparisons

	(1) DV Hotline	(2) DV Help	(3) Shelter	(4) Order	(5) Child Concern
2019 Relative to 2018					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2019	-3.2539** (1.5379)	-0.2463 (0.8791)	1.3197 (0.9354)	1.8347** (0.8602)	-0.3530 (0.4987)
2018 Relative to 2017					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2018	1.9816 (1.6436)	1.3299* (0.6823)	0.2673 (0.5671)	-0.6401 (0.8555)	-0.9941 (1.0436)
2017 Relative to 2016					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2017	-2.4888* (1.4584)	0.9108 (1.0276)	-0.7417 (0.9493)	-0.2016 (0.8646)	-0.4091 (0.9247)

This table gives the coefficient on the Post Week 11 x Year dummy variable of Equation 2.1 for each year relative to the previous year. The outcome variable is untransformed search interest in “domestic violence hotline” in column (1), “domestic violence help” in column(2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), the average of search interest in “protective order” and “restraining order” in column (4), and the average of “child protective services” and “child abuse” in column (5). Observations are at the DMA-week level and are weighted by DMA populations. These regressions include weeks 1 through 21 of the relevant years and include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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

Table 2.12: Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Excluding 2019

	(1)	(2)	(3)	(4)	(5)
	DV Hotline	DV Help	Shelter	Order	Child Concern
2020 Relative to All Years Except 2019					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2020	4.6724** (1.7515)	-1.8179* (0.9518)	-2.0699*** (0.7292)	-3.1246*** (1.0075)	-9.9115*** (1.3285)
<i>Inverse Hyperbolic Sine of Search Interest</i>					
Post Week 11 x Year 2018	0.2881*** (0.0890)	-0.1670* (0.0848)	-0.1519** (0.0680)	-0.1372*** (0.0423)	-0.3329*** (0.0275)

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for 2020 relative to 2016 through 2018. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column(2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), the average of search interest in “protective order” and “restraining order” in column (4), and the average of “child protective services” and “child abuse” in column (5). Observations are at the DMA-week level and are weighted by DMA populations. These regressions include weeks 1 through 21 of the relevant years and include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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

networks and other resources available to them pre-pandemic, then this may give us an idea of the expected magnitude of such a substitution.

Column (1) shows estimates of the effect of sheltering-in-place on search interest in the word toad. It is not obvious to me how to interpret the statistical significance of the increased search interest in this term, but it is worth noting that in percentage terms, the magnitude of the increase is small relative to that of the main outcomes of interest. I estimate an eight percent increase in search interest in the word toad before and after Week 11 in 2020 relative to other years whereas the outcomes of interest were all impacted by 15 percent or more.

Column (2) shows estimates of the effect of sheltering-in-place on search interest in toothaches. As expected, there is a positive and significant increase of about 16 percent as a result of lockdowns. This is about the size of the observed effects on search interest in domestic violence help, shelters, and protective orders, but it is about half the size of the effect on domestic violence hotline search interest and interest in child abuse or child protective ser-

Table 2.13: Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Including a Time Trend

	(1)	(2)	(3)	(4)	(5)
	DV Hotline	DV Help	Shelter	Order	Child Concern
A. 2020 Relative to All Other Years					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2020	5.3971*** (1.8460)	-2.0192* (1.0853)	-2.3995*** (0.7036)	-3.4550*** (1.0331)	-9.6074*** (1.2999)
<i>Inverse Hyperbolic Sine of Search Interest</i>					
Post Week 11 x Year 2020	0.3023*** (0.0828)	-0.1877* (0.1055)	-0.1617** (0.0629)	-0.1528*** (0.0378)	-0.3237*** (0.0269)
B. 2020 Relative to 2019					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2020	7.3204*** (2.4891)	-2.8023* (1.5437)	-3.3024*** (0.8926)	-4.4102*** (1.2939)	-8.8114*** (1.2638)
<i>Inverse Hyperbolic Sine of Search Interest</i>					
Post Week 11 x Year 2020	0.3399*** (0.1098)	-0.2509 (0.1699)	-0.1962*** (0.0637)	-0.2044*** (0.0372)	-0.2979*** (0.0299)

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1, but including a time trend. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), the average of search interest in “protective order” and “restraining order” in column (4), and the average of “child protective services” and “child abuse” in column (5). Observations are at the DMA-week level and are weighted by DMA populations. These regressions include Weeks 1 through 21 of 2016 through 2020 in Panel A and Weeks 1 through 21 of 2019 and 2020 in Panel B. They include year, DMA, month-of-year, and post Week 11 fixed effects as well as a time trend. Standard errors are clustered at the state level.

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

Table 2.14: Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Clustering at DMA Level

	(1)	(2)	(3)	(4)	(5)
	DV Hotline	DV Help	Shelter	Order	Child Concern
A. 2020 Relative to All Other Years					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2020	5.3554*** (1.6983)	-2.0438** (0.9755)	-2.3772*** (0.5966)	-3.4231*** (0.9381)	-9.6689*** (1.0298)
<i>Inverse Hyperbolic Sine of Search Interest</i>					
Post Week 11 x Year 2020	0.3017*** (0.0899)	-0.1863** (0.0871)	-0.1630*** (0.0527)	-0.1538*** (0.0375)	-0.3255*** (0.0248)
B. 2020 Relative to 2019					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2020	7.3204*** (2.2676)	-2.8023** (1.4178)	-3.3024*** (0.9263)	-4.4102*** (1.3360)	-8.8114*** (1.1657)
<i>Inverse Hyperbolic Sine of Search Interest</i>					
Post Week 11 x Year 2020	0.3399*** (0.1013)	-0.2509* (0.1406)	-0.1962*** (0.0654)	-0.2044*** (0.0456)	-0.2979*** (0.0316)

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 clustering standard errors at the DMA level. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), the average of search interest in “protective order” and “restraining order” in column (4), and the average of “child protective services” and “child abuse” in column (5). Observations are at the DMA-week level and are weighted by DMA populations. These regressions include weeks 1 through 21 of the relevant years. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the DMA level.

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

Table 2.15: Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Full Year

	(1)	(2)	(3)	(4)	(5)
	DV Hotline	DV Help	Shelter	Order	Child Concern
A. 2020 Relative to All Other Years					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2020	2.0146** (0.9937)	-0.9472 (0.8709)	-2.7343*** (0.6023)	-1.2209 (0.7550)	-5.5757*** (1.1533)
<i>Inverse Hyperbolic Sine of Search Interest</i>					
Post Week 11 x Year 2020	0.0996 (0.0613)	-0.0780 (0.0716)	-0.1668*** (0.0480)	-0.0692** (0.0284)	-0.1839*** (0.0258)
B. 2020 Relative to 2019					
<i>Untransformed Search Interest</i>					
Post Week 11 x Year 2020	4.5385*** (1.6522)	-1.4367 (0.9142)	-3.1371*** (0.8387)	-1.5830 (1.0714)	-5.4043*** (1.2507)
<i>Inverse Hyperbolic Sine of Search Interest</i>					
Post Week 11 x Year 2020	0.1982** (0.0872)	-0.1481* (0.0797)	-0.1578*** (0.0561)	-0.1112*** (0.0358)	-0.1789*** (0.0326)

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for the entire year. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), the average of search interest in “protective order” and “restraining order” in column (4), and the average of “child protective services” and “child abuse” in column (5). Observations are at the DMA-week level and are weighted by DMA populations. These regressions include all weeks of the relevant years. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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

Table 2.16: Effect of Staying-at-Home on Domestic Violence Related Search Behavior - Placebo Variables

	(1)	(2)
	Toad	Toothache
A. 2020 Relative to All Other Years		
<i>Untransformed Search Interest</i>		
Post Week 11 x Year 2020	2.8451***	6.1627***
	(0.5624)	(1.7555)
<i>Inverse Hyperbolic Sine of Search Interest</i>		
Post Week 11 x Year 2020	0.0806***	0.1642***
	(0.0207)	(0.0305)
B. 2020 Relative to 2019		
<i>Untransformed Search Interest</i>		
Post Week 11 x Year 2020	2.6786***	8.5473***
	(0.7193)	(2.0914)
<i>Inverse Hyperbolic Sine of Search Interest</i>		
Post Week 11 x Year 2020	0.0823***	0.2303***
	(0.0235)	(0.0433)


This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for placebo outcome variables. The outcome variable is search interest in “toad” in column (1) and “toothache” in column (2). Observations are at the DMA-week level and are weighted by DMA populations. These regressions include weeks 1 through 21 of 2016 through 2020 in Panel A and weeks 1 through 21 of 2019 and 2020 in Panel B. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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

vices, so the effects appear to be more than just a result of substitution between in-person and remotely available resources. Figures 2.22 and 2.23 show the event study graphs for the placebo variables. Figure 2.22 shows that search interest in the word toad was exceptionally low in the weeks before Week 11 in 2020 relative to other years.


I also check to see if the results are stronger in states where individuals are less likely to live alone as we should not find domestic violence in single-person households. Table 2.17 shows the effects of sheltering-in-place on domestic violence search interest by quartiles of the percent of households with a householder living alone in each state in 2020. Figure 2.24 depicts the average quartile assigned to each state. As we expect, results are strongest in regions where individuals are least likely to live alone.

Figure 2.22: Event Study - Search Interest in Toads



This figure gives the estimates of the event study model described in Equation 2.2 for the first 21 weeks of the year. The red line (Week 11) indicates the week by which most DMAs had already begun to decrease their mobility. The black horizontal line indicates the average level of the week-of-year coefficient prior to Week 11.

Figure 2.23: Event Study - Search Interest in Toothaches



This figure gives the estimates of the event study model described in Equation 2.2 for the first 21 weeks of the year. The red line (Week 11) indicates the week by which most DMAs had already begun to decrease their mobility. The black horizontal line indicates the average level of the week-of-year coefficient prior to Week 11.


Table 2.17: Effects of Sheltering-In-Place by Quartiles of the Percent Living Alone in 2020

	(1) DV Hotline	(2) DV Help	(3) Shelters	(4) Orders
Quartile 1 <i>Percent Living Alone in 2020 of 0.178 to 0.277</i>				
Post Week 11 x Year 2020	7.5286** (2.7353)	-1.9659 (1.7054)	-3.5273* (1.7591)	-5.4931*** (1.7024)
Mean of Outcome	4.2685	2.7245	5.8681	16.7864
Quartile 2 <i>Percent Living Alone in 2020 of 0.277 to 0.292</i>				
Post Week 11 x Year 2020	6.4625 (3.8939)	-3.4630 (2.5686)	-1.6112* (0.9281)	-1.1793 (1.0256)
Mean of Outcome	2.3207	1.6207	3.6603	11.3922
Quartile 3 <i>Percent Living Alone in 2020 of 0.292 to 0.305</i>				
Post Week 11 x Year 2020	1.0816 (1.0819)	-1.1191 (0.7670)	-1.2681* (0.7080)	-2.2291* (1.1025)
Mean of Outcome	1.0864	0.4353	1.2037	7.7124
Quartile 4 <i>Percent Living Alone in 2020 of 0.305 to 0.352</i>				
Post Week 11 x Year 2020	1.2181 (1.6598)		-2.0369 (1.7218)	-3.9159** (1.6559)
Mean of Outcome	0.6728	0.0000	1.2761	6.1114

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for each quartile of the distribution of percent of households where the householder lives alone in 2020. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), and the average of search interest in “protective order” and “restraining order” in column (4). Observations are at the DMA-week level and are weighted by DMA populations. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.


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

Figure 2.24: Average Quartile within a State of Percent of 2020 Households with 1 Person Living Alone



This figure reflects the average quartile among DMAs within a state of the DMA-level percent of households in which the householder lives alone in 2020. DMAs are assigned a value between 1 and 4 where 1 describes DMAs that fall below the 25th percentile and 4 describes DMAs that fall above the 75th percentile.

Figure 2.25: Average Quartile within a State of Percent of 2020 Households that Include a Couple Living Together



This figure reflects the average quartile among DMAs within a state of the DMA-level percent of households in which the householder lives with a romantic partner in 2020. DMAs are assigned a value between 1 and 4 where 1 describes DMAs that fall below the 25th percentile and 4 describes states that fall above the 75th percentile.

Likewise, I check to see if the results are stronger in states where individuals are more likely to live with a partner as we would expect. Table 2.18 shows the effects of sheltering-in-place on domestic violence search interest by quartiles of the percent of households in which the householder lives with a romantic partner in each state in 2020. Figure 2.25 depicts the quartiles assigned to each state. Results tend to be stronger in the upper quartiles where the heads of household are more likely to be living with a romantic partner.

Table 2.18: Effects of Sheltering-In-Place by Quartiles of the Percent Living with a Partner in 2020

	(1) DV Hotline	(2) DV Help	(3) Shelters	(4) Orders
Quartile 1 <i>Percent Living with a Partner in 2020 of 0.399 to 0.529</i>				
Post Week 11 x Year 2020	6.1067 (4.4251)	-0.9477 (0.6942)	-2.8936*** (1.0357)	-2.2330 (1.4646)
Mean of Outcome	1.8071	0.5198	2.3956	9.6001
Quartile 2 <i>Percent Living with a Partner in 2020 of 0.529 to 0.547</i>				
Post Week 11 x Year 2020	0.8616 (0.8857)	-3.0914 (3.0466)	-1.6548* (0.9292)	-0.9962 (0.9262)
Mean of Outcome	1.8110	1.1079	2.9711	12.0445
Quartile 3 <i>Percent Living with a Partner in 2020 of 0.547 to 0.566</i>				
Post Week 11 x Year 2020	11.0355*** (3.5327)	-3.1888** (1.3375)	-2.9111* (1.6072)	-6.8046** (2.5066)
Mean of Outcome	2.5038	1.2678	3.3673	9.0948
Quartile 4 <i>Percent Living with a Partner in 2020 of 0.566 to 0.655</i>				
Post Week 11 x Year 2020	3.7823** (1.7410)	-0.9097 (1.5845)	-2.1215 (1.5110)	-3.7419*** (0.5568)
Mean of Outcome	2.1971	1.8496	3.2365	11.1658

This table gives the coefficient on the Post Week 11 x Year 2020 dummy variable of Equation 2.1 for each quartile of the distribution of households where the householder lives with a spouse or partner in 2020. The outcome variable is search interest in “domestic violence hotline” in column (1), “domestic violence help” in column (2), the average of search interest in “domestic violence shelter” and “women’s shelter” in column (3), and the average of search interest in “protective order” and “restraining order” in column (4). Observations are at the DMA-week level and are weighted by DMA populations. They include year, DMA, month-of-year, and post Week 11 fixed effects. Standard errors are clustered at the state level.

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

2.6 Conclusion

In this paper, I use Google Trends data to measure the prevalence of search terms related to domestic violence. Despite its prevalence, domestic violence is hard to measure. Crime data is inconsistently coded across police departments and domestic violence is particularly vulnerable to under-reporting concerns. Survey data is generally only available at annual frequencies, and there is a lag between when a survey is conducted and when it is available to researchers. Google Trends data is particularly useful in this context since it provides consistent nationwide data reflecting anonymous searches at high-frequency intervals with essentially no lag time.

I find that COVID-19 lockdowns result in a large and statistically significant increase in search interest for domestic violence hotlines, but a decrease in search interest for other resources related to domestic violence such as protective orders and domestic violence shelters. I hypothesize that the difference between these effects reflects substitution by victims of domestic violence away from more disruptive coping strategies and toward more passive strategies in response to external conditions that make leaving ones home especially costly. That is, when women feel unable to leave their abusers, they are more likely to search for solutions that allow them to stay home. Based on the timing of search interest in these resources, there is some evidence that it may take longer for domestic violence to escalate to the point where women are ready for more disruptive options such as restraining orders.

In addition, I find evidence suggesting that increased exposure between victim and abuser is a key component of this relationship, and that COVID-19 anxiety and increased household density appear to intensify this exposure. The increase in hotline search interest appears to be mitigated in areas where women are more likely to work, suggesting that female employment may be important for preventing abuse. I also find evidence that effects on domestic violence search interest are stronger in areas that were already experiencing increased un-

employment prior to the pandemic, supporting the idea that financial strain contributes to domestic violence. Finally, the relatively large decrease in search interest related to concern for child welfare suggests that the pandemic disrupted third-party reporting channels. As a result, estimates of abuse that rely on reporting are likely to be even more dramatically underestimated during this time than usual.

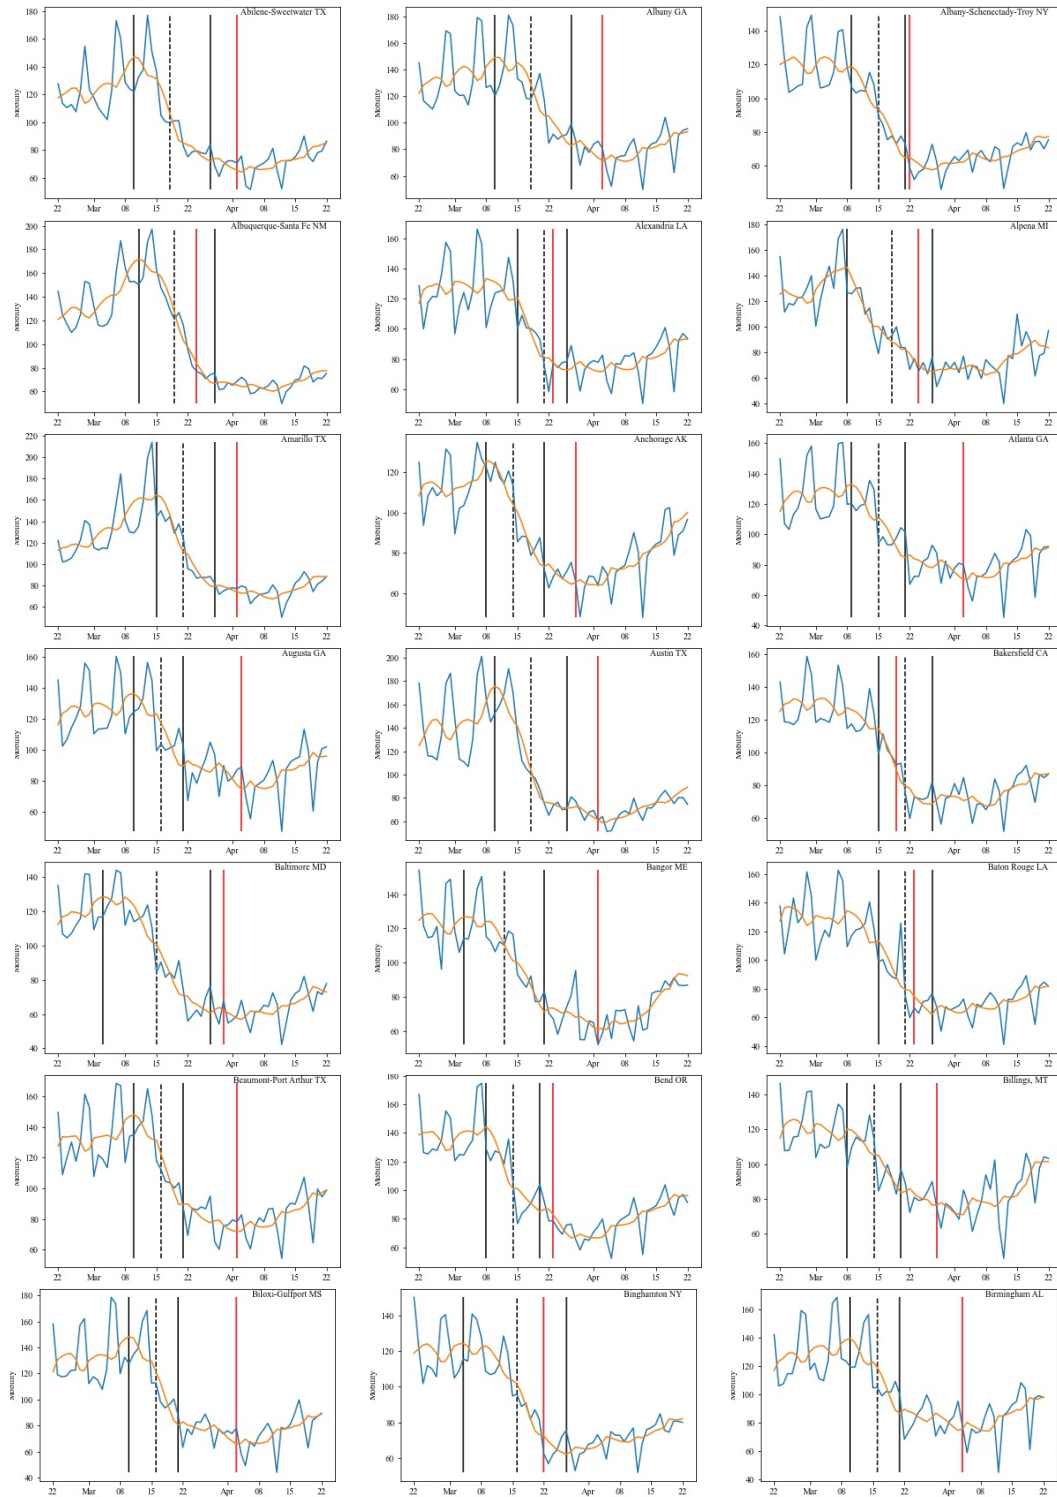
While other papers document a relationship between lockdowns and domestic violence via police calls-for-service, this paper adds to our understanding of how victims cope with domestic violence and shows that these strategies may be nuanced and context-dependent. Victims appear to substitute away from more direct abuse-avoidance tactics (e.g protective orders) and toward more passive measures when they are less able to make a clean break from their abusers.

These findings have implications beyond the context of a global pandemic. They highlight the role of victim-perpetrator exposure and financial stress in exacerbating domestic violence. They also suggest that female economic independence can insulate women from domestic violence and that third-party reporting is an important channel for identifying abuse. Finally, I find that victims seek out domestic violence hotlines, especially when leaving one's home is an impractical option, which suggests that these hotlines are valuable resources for domestic violence survivors.

2.7 Mobility Reduction Timelines in Each DMA

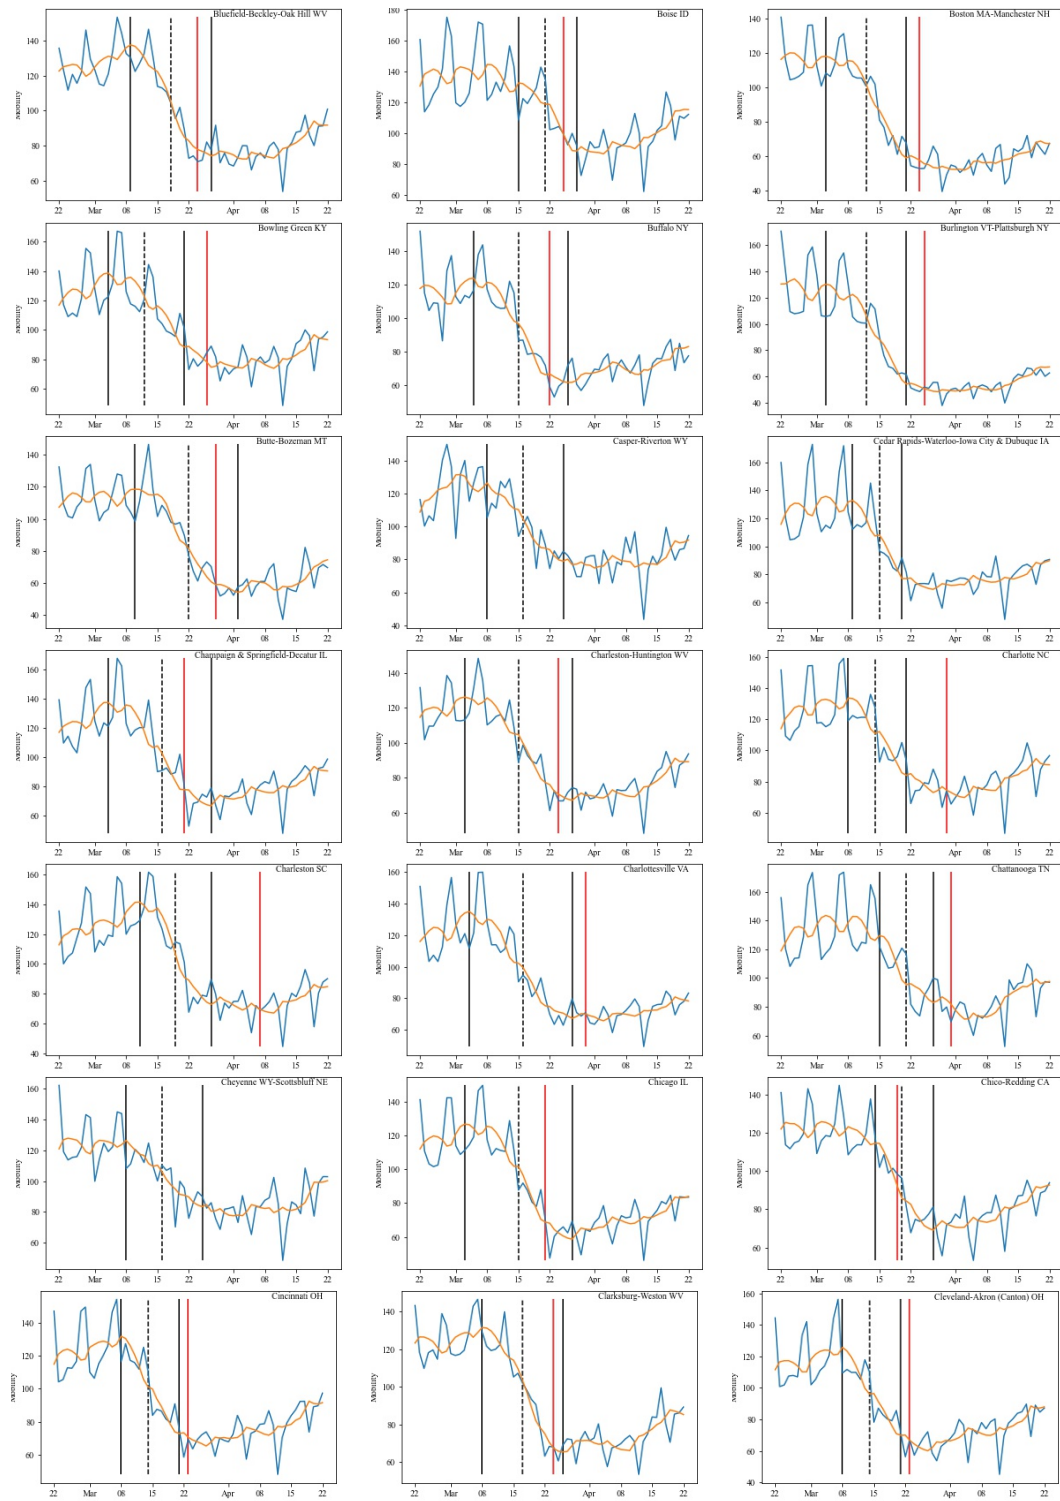
Figures 2.26 through 2.35 demonstrate the timelines of mobility reduction identified for each DMA using Algorithm 1.

Figure 2.26: Mobility Reductions by DMA - Part 1



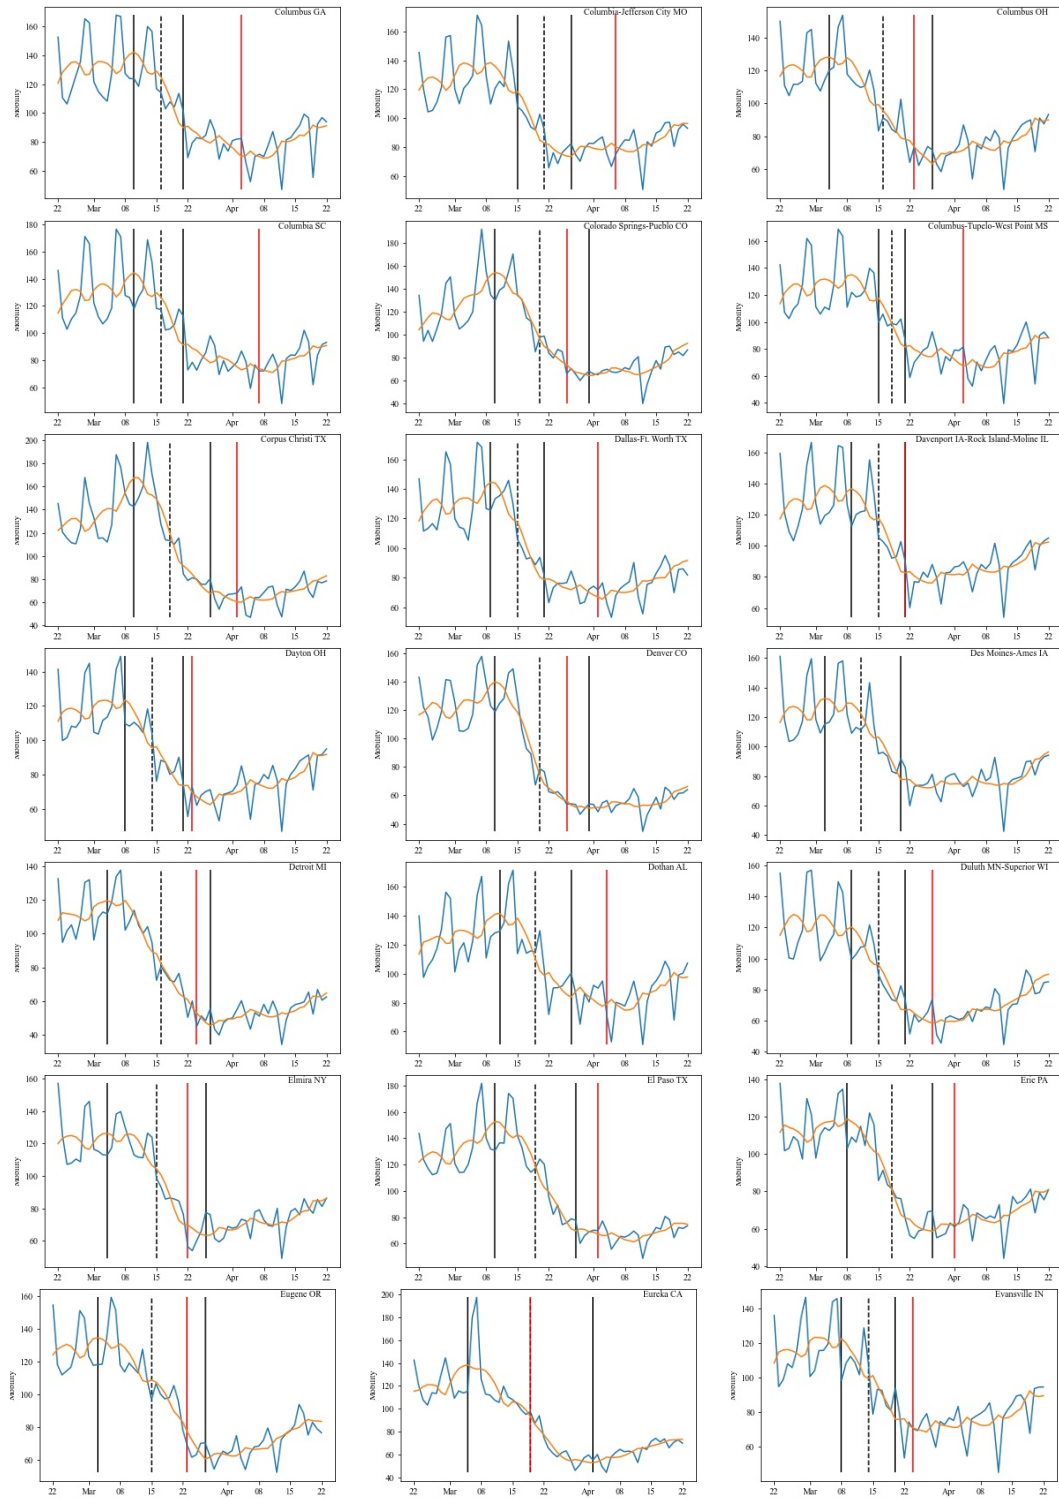
This figure shows show the timeline of mobility reduction for each DMA. The three vertical black lines represent the start, midpoint, and end of the drop in mobility. The vertical red line is the date a state-mandated lockdown order went into effect, where applicable.

Figure 2.27: Mobility Reductions by DMA - Part 2



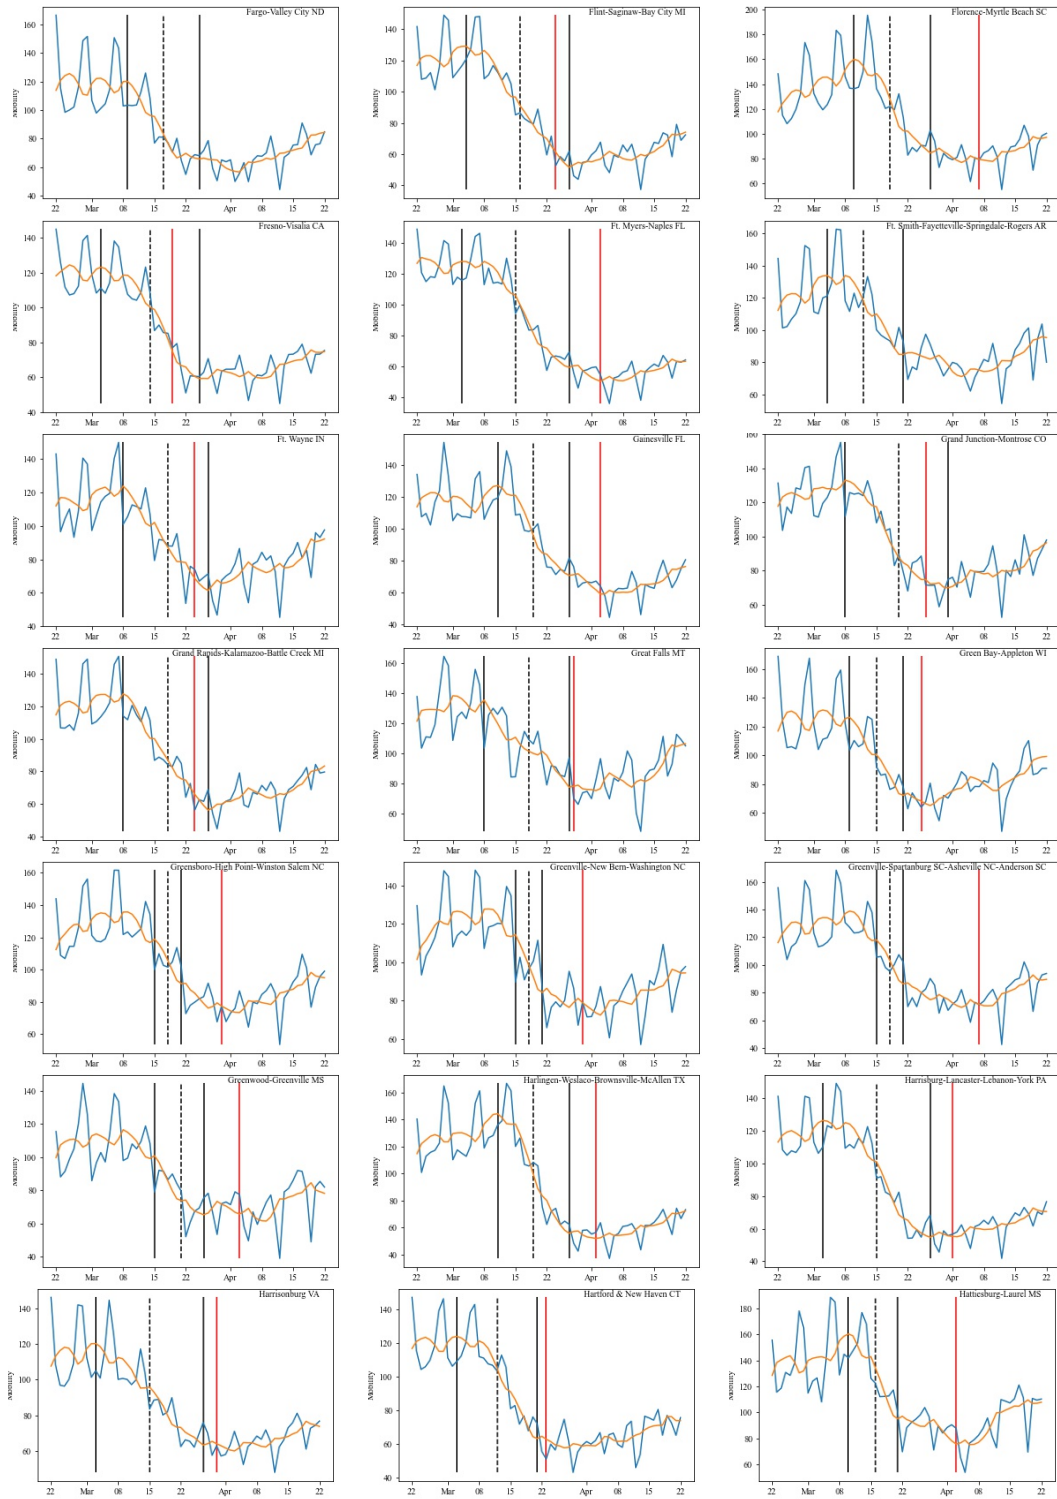
This figure shows show the timeline of mobility reduction for each DMA. The three vertical black lines represent the start, midpoint, and end of the drop in mobility. The vertical red line is the date a state-mandated lockdown order went into effect, where applicable.

Figure 2.28: Mobility Reductions by DMA - Part 3



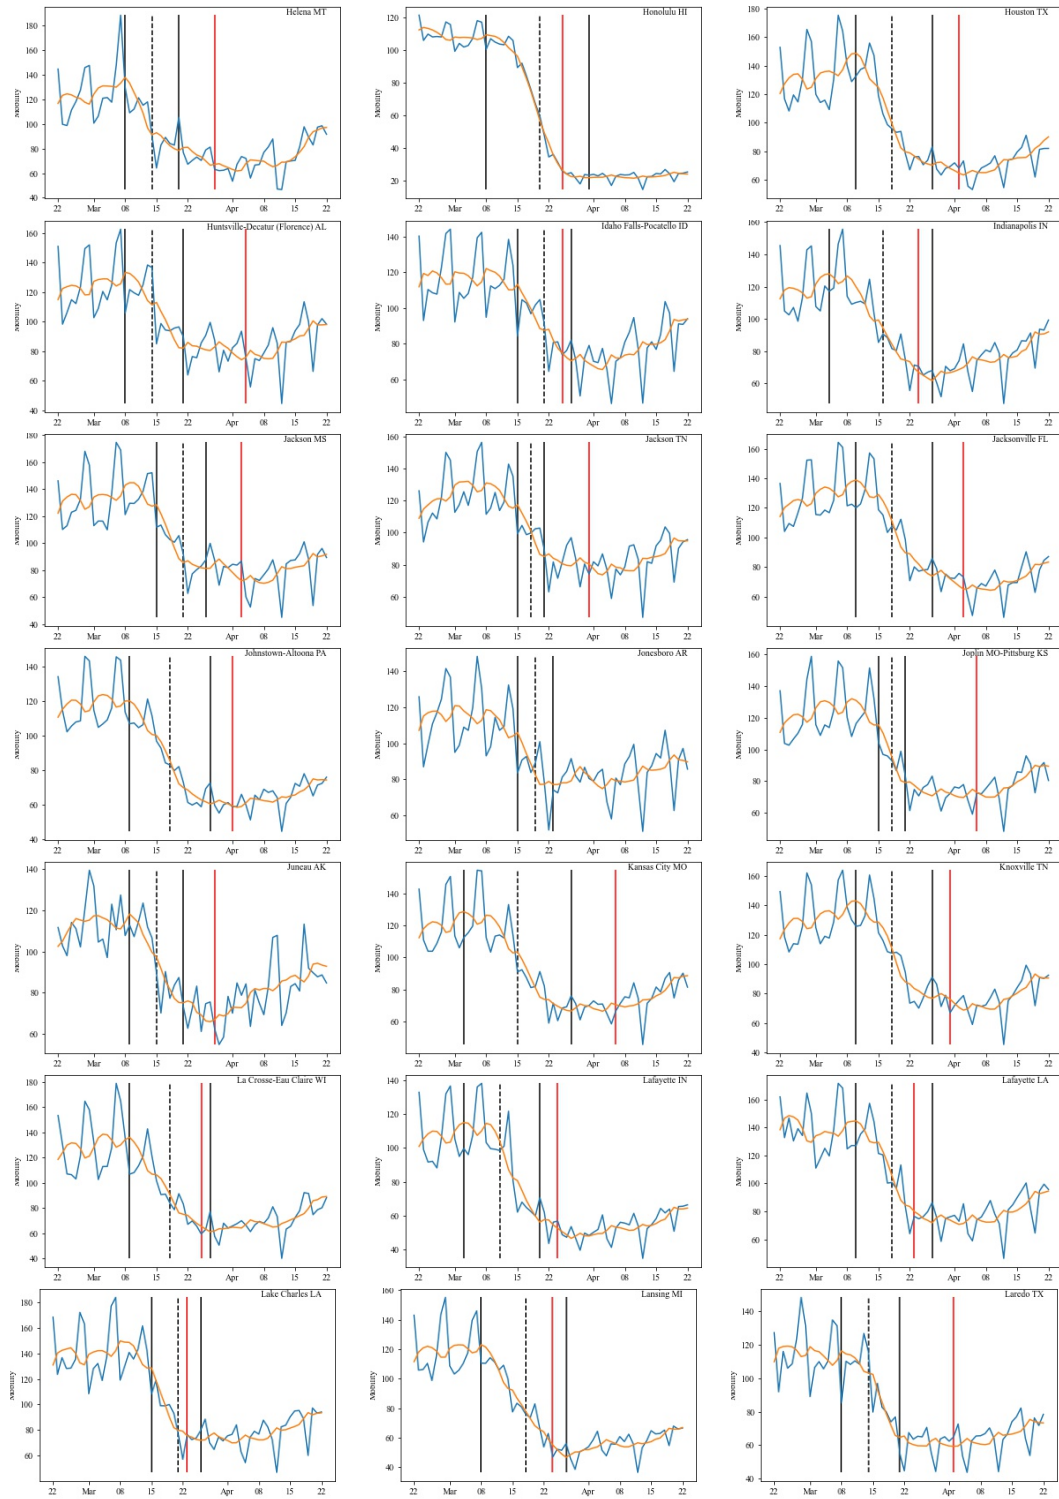
This figure shows show the timeline of mobility reduction for each DMA. The three vertical black lines represent the start, midpoint, and end of the drop in mobility. The vertical red line is the date a state-mandated lockdown order went into effect, where applicable.

Figure 2.29: Mobility Reductions by DMA - Part 4



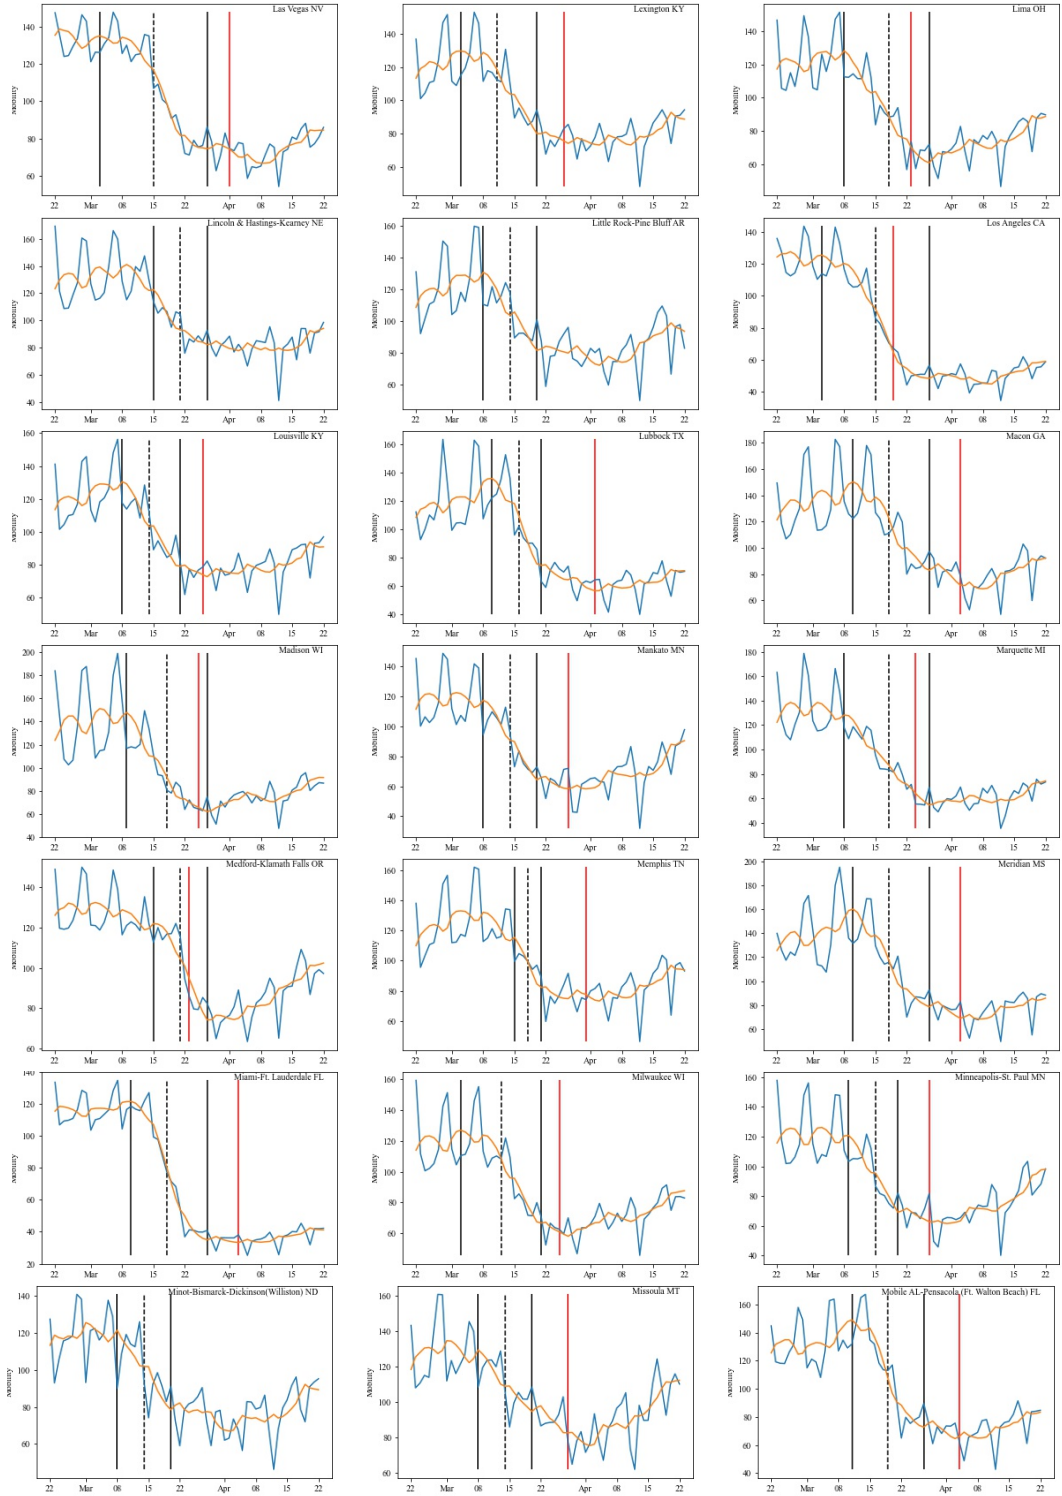
This figure shows show the timeline of mobility reduction for each DMA. The three vertical black lines represent the start, midpoint, and end of the drop in mobility. The vertical red line is the date a state-mandated lockdown order went into effect, where applicable.

Figure 2.30: Mobility Reductions by DMA - Part 5



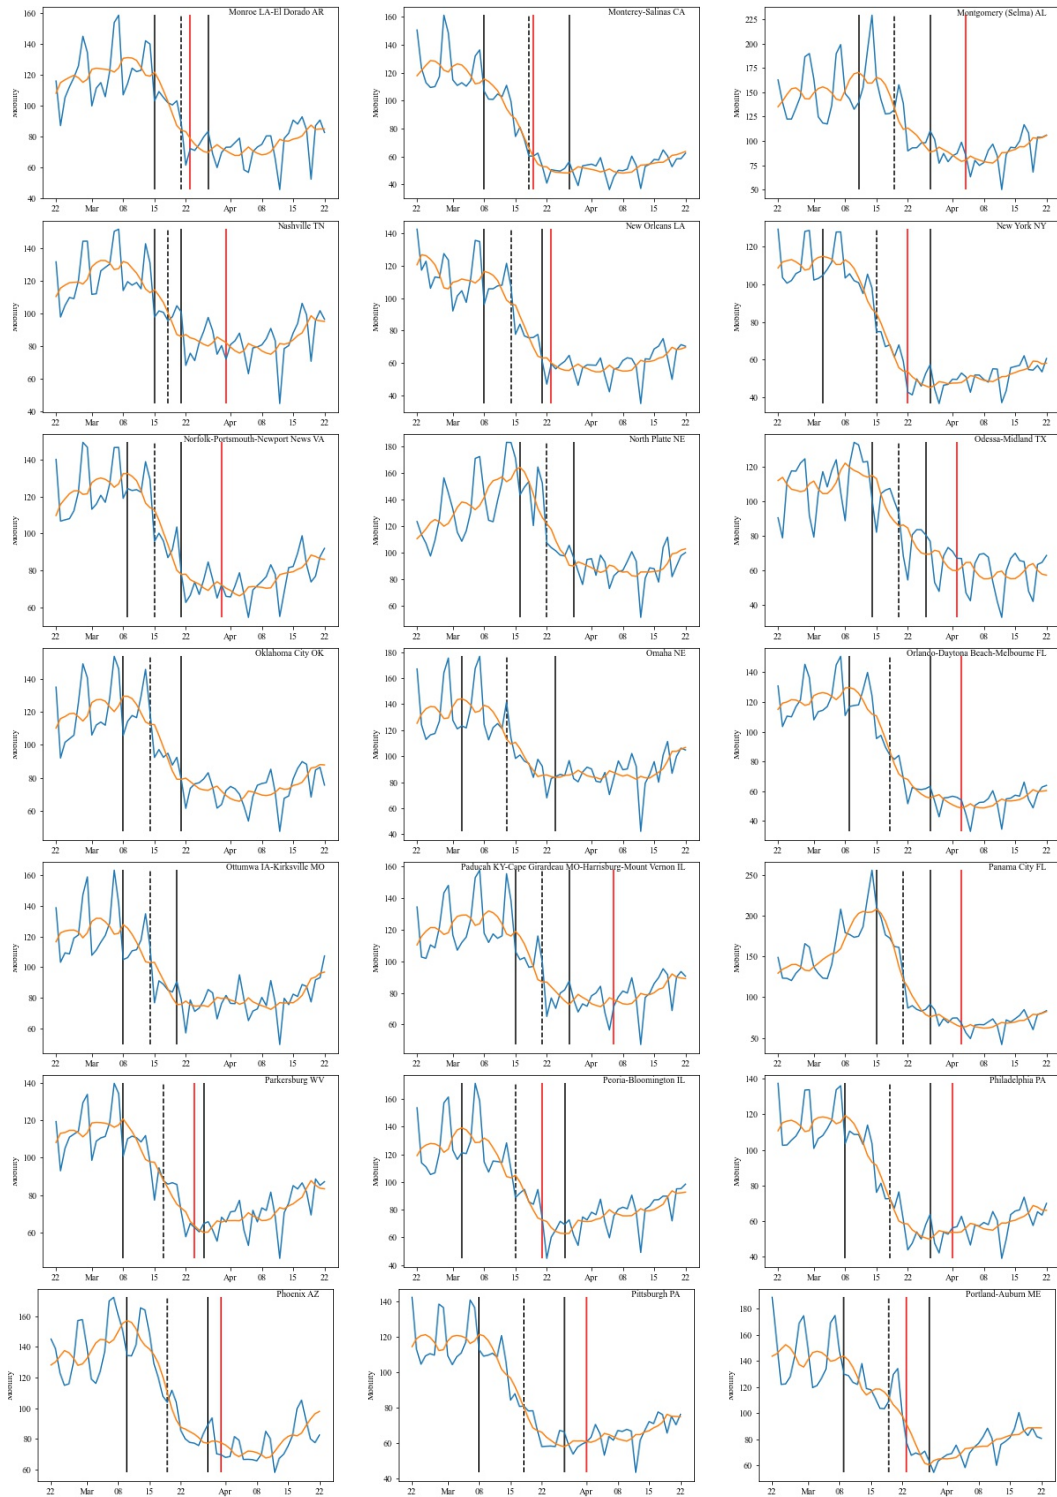
This figure shows show the timeline of mobility reduction for each DMA. The three vertical black lines represent the start, midpoint, and end of the drop in mobility. The vertical red line is the date a state-mandated lockdown order went into effect, where applicable.

Figure 2.31: Mobility Reductions by DMA - Part 6



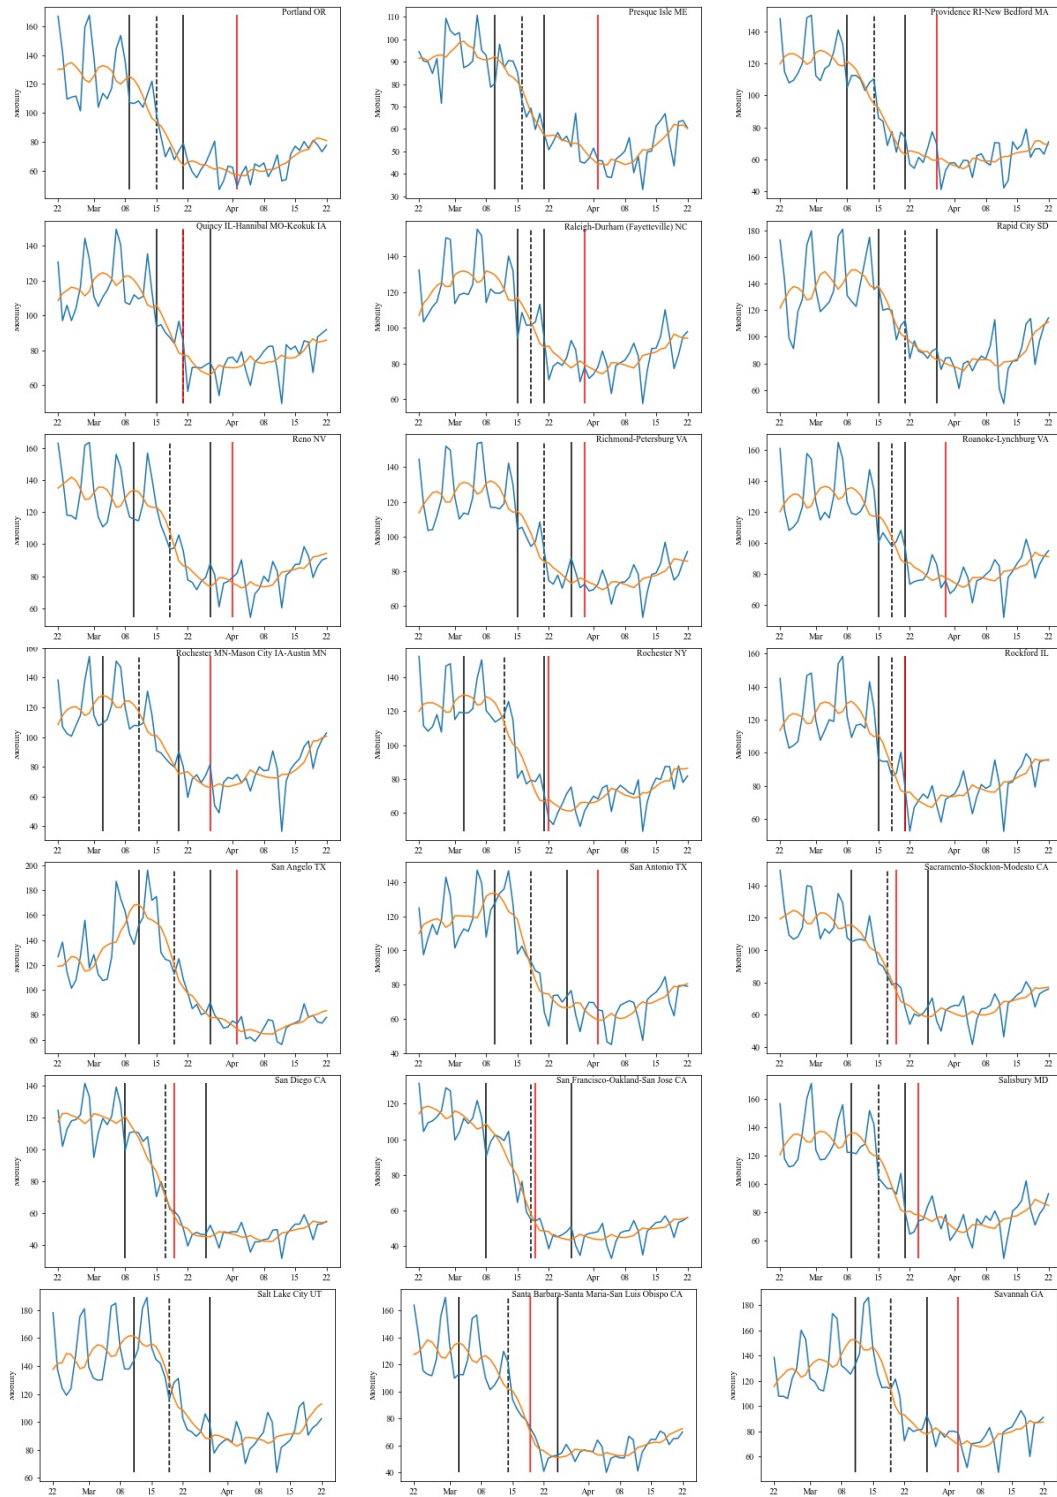
This figure shows show the timeline of mobility reduction for each DMA. The three vertical black lines represent the start, midpoint, and end of the drop in mobility. The vertical red line is the date a state-mandated lockdown order went into effect, where applicable.

Figure 2.32: Mobility Reductions by DMA - Part 7



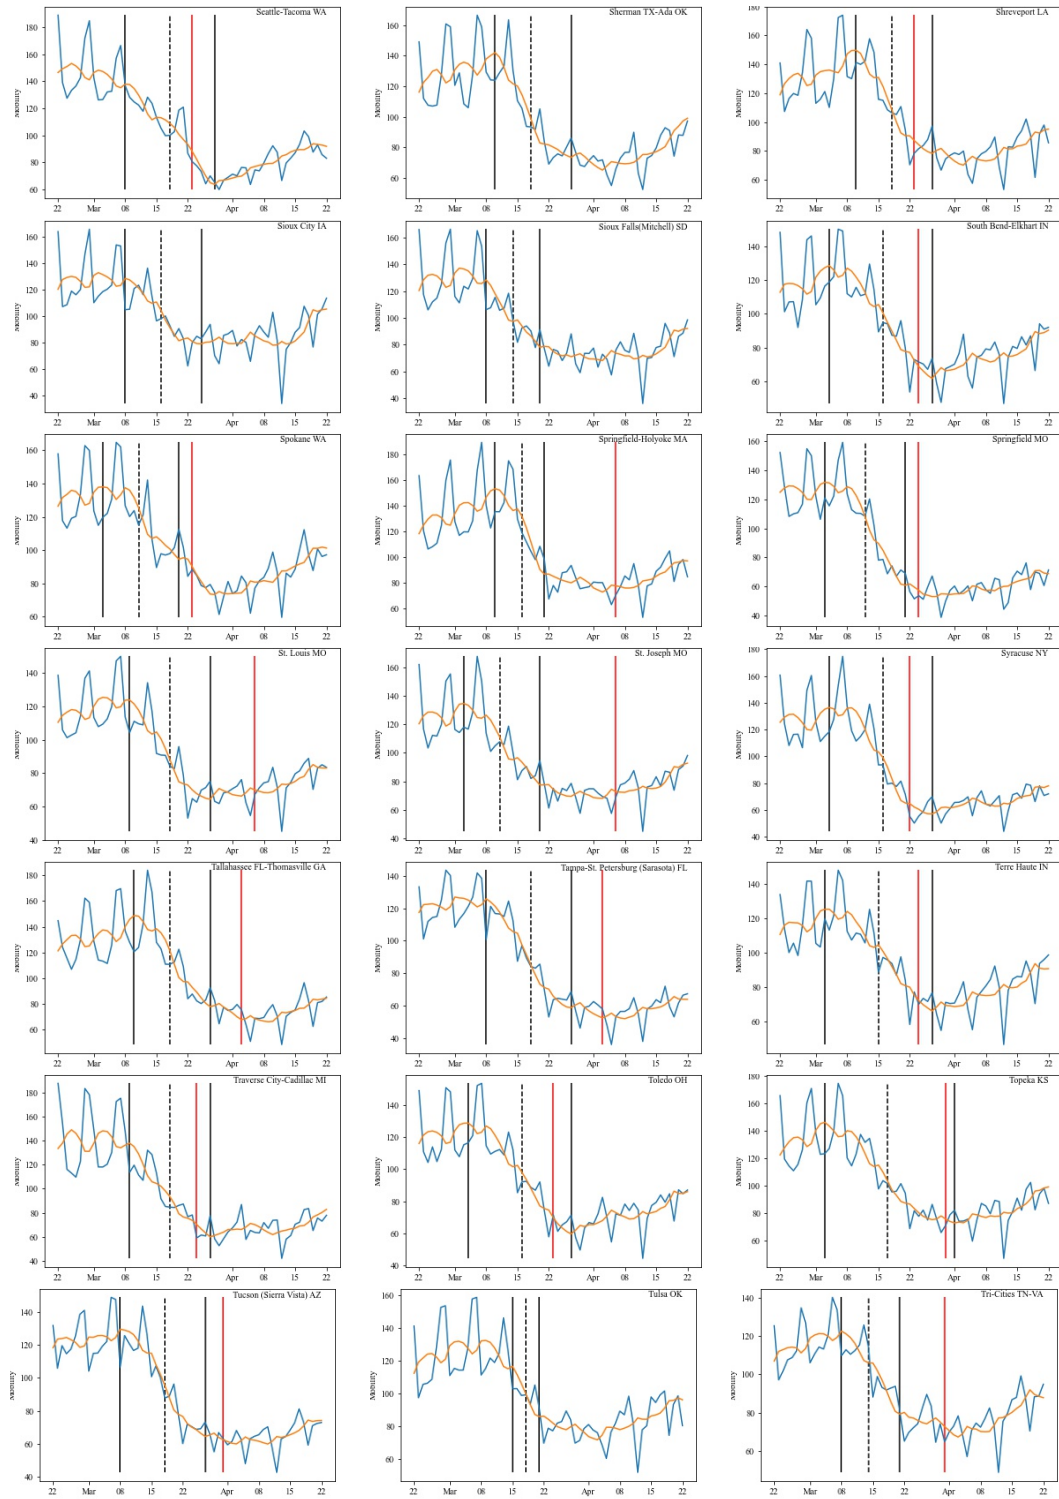
This figure shows show the timeline of mobility reduction for each DMA. The three vertical black lines represent the start, midpoint, and end of the drop in mobility. The vertical red line is the date a state-mandated lockdown order went into effect, where applicable.

Figure 2.33: Mobility Reductions by DMA - Part 8



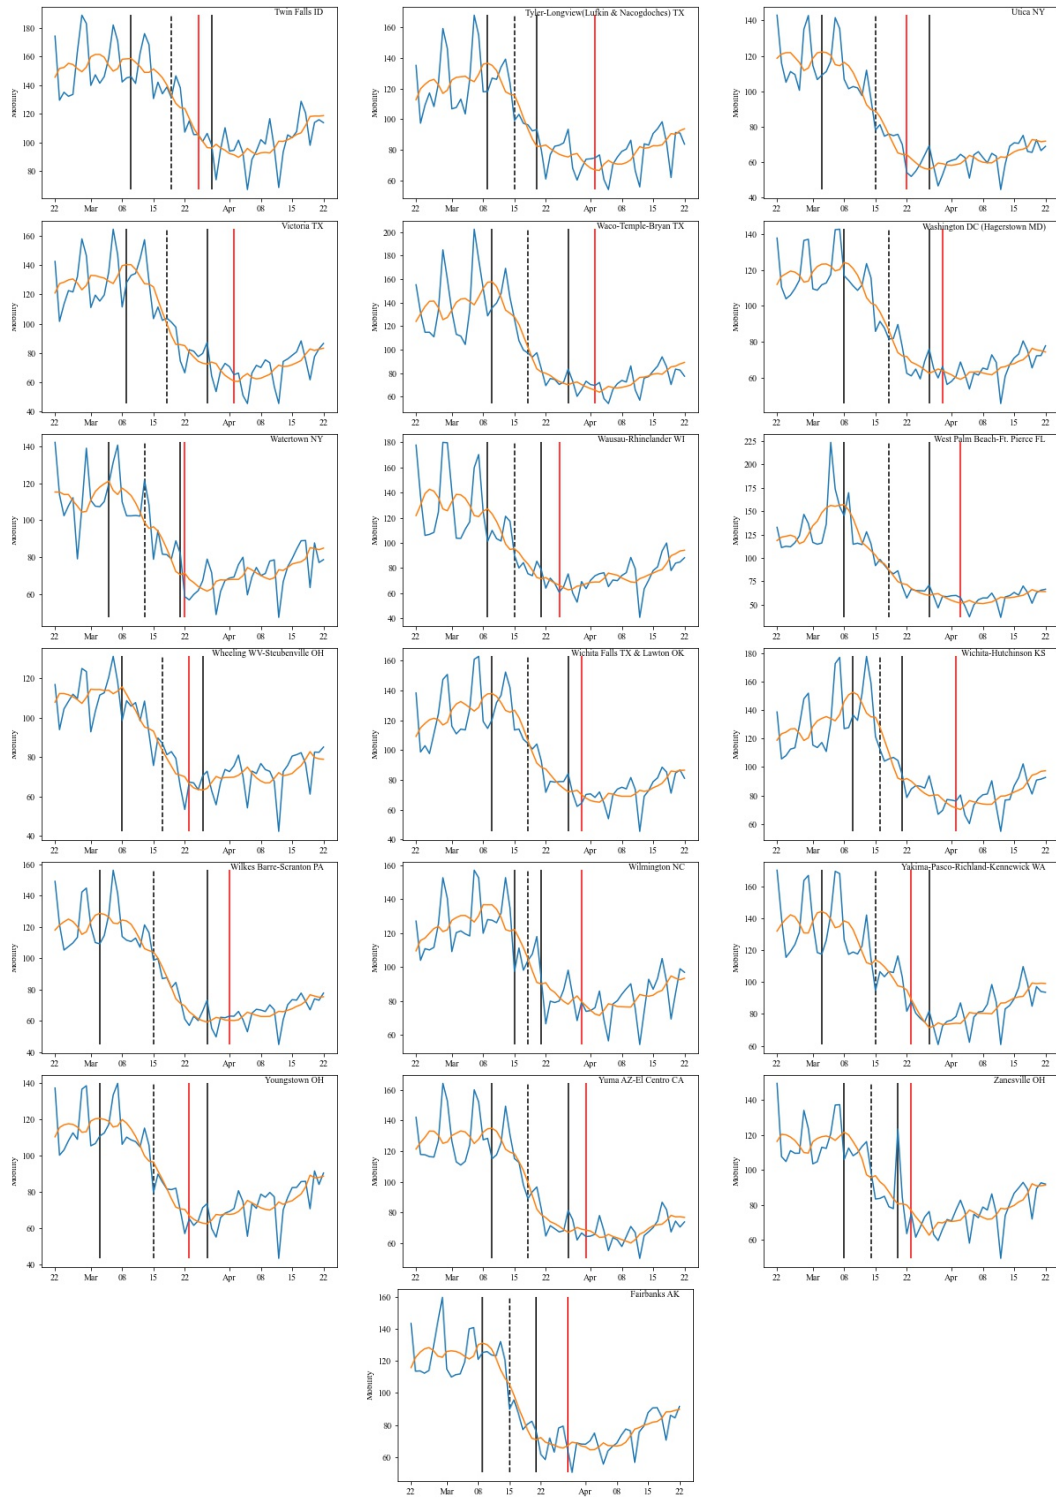
This figure shows show the timeline of mobility reduction for each DMA. The three vertical black lines represent the start, midpoint, and end of the drop in mobility. The vertical red line is the date a state-mandated lockdown order went into effect, where applicable.

Figure 2.34: Mobility Reductions by DMA - Part 9



This figure shows show the timeline of mobility reduction for each DMA. The three vertical black lines represent the start, midpoint, and end of the drop in mobility. The vertical red line is the date a state-mandated lockdown order went into effect, where applicable.

Figure 2.35: Mobility Reductions by DMA - Part 10



This figure shows show the timeline of mobility reduction for each DMA. The three vertical black lines represent the start, midpoint, and end of the drop in mobility. The vertical red line is the date a state-mandated lockdown order went into effect, where applicable.

Bibliography

- Aigner, D. J., & Cain, G. G. (1977). Statistical theories of discrimination in labor markets. *Ilr Review*, 30(2), 175–187.
- Aizer, A. (2010). The gender wage gap and domestic violence. *American Economic Review*, 100(4), 1847–59.
- Altonji, J. G., & Pierret, C. R. (2001). Employer learning and statistical discrimination. *The quarterly journal of economics*, 116(1), 313–350.
- American community survey. (n.d.). <https://data.census.gov/cedsci/>.
- Apple. (2020). Mobility trends reports [Accessed: 2022-01-19]. <https://covid19.apple.com/mobility>.
- Arrow, K. J. (1973). *The theory of discrimination*. Princeton University Press.
- Ashby, M. P. (2020). Initial evidence on the relationship between the coronavirus pandemic and crime in the united states. *Crime Science*, 9(1), 1–16.
- Basic monthly cps. (n.d.). <https://www.census.gov/data/datasets/time-series/demo/cps/cps-basic.2016.html>.
- Becker, G. S. (1957). *The economics of discrimination*. University of Chicago press.
- Becker, G. S. (1968). Crime and punishment: An economic approach. *The economic dimensions of crime* (pp. 13–68). Springer.
- Becker, G. S. (1985). Human capital, effort, and the sexual division of labor. *Journal of labor economics*, 3(1, Part 2), S33–S58.

- Begley, S. (2014). Ironic misandry: Why feminists pretending to hate men isn't funny. *Time Magazine*.
- Beland, L.-P., Brodeur, A., Haddad, J., & Mikola, D. (2020). Covid-19, family stress and domestic violence: Remote work, isolation and bargaining power. *IZA Discussion Papers*.
- Berinsky, A. J. (1999). The two faces of public opinion. *American Journal of Political Science*, 1209–1230.
- Bianchi, S. M., Milkie, M. A., Sayer, L. C., & Robinson, J. P. (2000). Is anyone doing the housework? trends in the gender division of household labor. *Social forces*, 79(1), 191–228.
- Black, D. A. (1995). Discrimination in an equilibrium search model. *Journal of labor Economics*, 13(2), 309–334.
- Blau, F. D., & Kahn, L. M. (2007). The gender pay gap: Have women gone as far as they can? *Academy of Management Perspectives*, 21(1), 7–23.
- Blau, F. D., & Kahn, L. M. (2017). The gender wage gap: Extent, trends, and explanations. *Journal of Economic Literature*, 55(3), 789–865.
- Burn, I. (2019). The relationship between prejudice and wage penalties for gay men in the united states. *ILR Review*, 0019793919864891.
- Charles, K. K., & Guryan, J. (2008). Prejudice and wages: An empirical assessment of becker's the economics of discrimination. *Journal of political economy*, 116(5), 773–809.
- Charles, K. K., Guryan, J., & Pan, J. (2018). *The effects of sexism on american women: The role of norms vs. discrimination* (tech. rep.). National Bureau of Economic Research.
- Chin, Y.-M. (2012). Male backlash, bargaining, or exposure reduction?: Women's working status and physical spousal violence in india. *Journal of population Economics*, 25(1), 175–200.
- Cps org uniform extracts, version 2.5. (2020). *Center for Economic and Policy Research*.

- De la Roca, J., Ellen, I. G., & O'Regan, K. M. (2014). Race and neighborhoods in the 21st century: What does segregation mean today? *Regional Science and Urban Economics*, *47*, 138–151.
- Dhanaraj, S., & Mahambare, V. (2022). Male backlash and female guilt: Women's employment and intimate partner violence in urban india. *Feminist Economics*, *28*(1), 170–198.
- Dugan, L., Nagin, D. S., & Rosenfeld, R. (1999). Explaining the decline in intimate partner homicide: The effects of changing domesticity, women's status, and domestic violence resources. *Homicide Studies*, *3*(3), 187–214.
- Dugan, L., Nagin, D. S., & Rosenfeld, R. (2003). Exposure reduction or retaliation? the effects of domestic violence resources on intimate-partner homicide. *Law & society review*, *37*(1), 169–198.
- Economist. (2020). Domestic violence has increased during coronavirus lockdowns. *The Economist*. <https://www.economist.com/graphic-detail/2020/04/22/domestic-violence-has-increased-during-coronavirus-lockdowns>
- Enarson, E. (1999). Violence against women in disasters: A study of domestic violence programs in the united states and canada. *Violence against women*, *5*(7), 742–768.
- Farmer, A., & Tiefenthaler, J. (1997). An economic analysis of domestic violence. *Review of social Economy*, *55*(3), 337–358.
- Farrell, G., & Pease, K. (1994). Crime seasonality: Domestic disputes and residential burglary in merseyside 1988–90. *British Journal of Criminology*, *34*(4), 487–498.
- Farrell, T., Fernandez, M., Novotny, J., & Alani, H. (2019). Exploring misogyny across the manosphere in reddit. *Proceedings of the 10th ACM Conference on Web Science*, 87–96.
- Flood, S., King, M., Rodgers, R., Ruggles, S., Warren, J. R., & Westberry, M. (2021). Integrated public use microdata series, current population survey: Version 9.0 [dataset] [IPUMS.]. <https://doi.org/10.18128/D030.V9.0>.

- Fortin, N. (2005). Gender role attitudes and the labour-market outcomes of women across oecd countries. *oxford review of Economic Policy*, 21(3), 416–438.
- Fortin, N. (2008). Gender role attitudes and women’s labor market participation: Opting-out and the persistent appeal of housewifery. *Unpublished manuscript, University of British Columbia*.
- Gearhart, S., Perez-Patron, M., Hammond, T. A., Goldberg, D. W., Klein, A., & Horney, J. A. (2018). The impact of natural disasters on domestic violence: An analysis of reports of simple assault in florida (1999–2007). *Violence and gender*, 5(2), 87–92.
- Google trends [<https://www.google.com/trends>]. (2021).
- Henke, A., & Hsu, L. (2022). Covid-19 and domestic violence: Economics or isolation? *Journal of Family and Economic Issues*, 1–14.
- Hess, A. (2014). The rise of the ironic man-hater. *Slate Magazine*.
- Hsu, L.-C., & Henke, A. (2021). Covid-19, staying at home, and domestic violence. *Review of Economics of the Household*, 19(1), 145–155.
- Iyengar, R. (2009). Does the certainty of arrest reduce domestic violence? evidence from mandatory and recommended arrest laws. *Journal of public Economics*, 93(1-2), 85–98.
- Kirby, S., Francis, B., & O’Flaherty, R. (2014). Can the fifa world cup football (soccer) tournament be associated with an increase in domestic abuse? *Journal of research in crime and delinquency*, 51(3), 259–276.
- Koutaniemi, E. M., & Einiö, E. (2019). Seasonal variation in seeking help for domestic violence based on google search data and finnish police calls in 2017. *Scandinavian journal of public health*, 1403494819834098.
- Leslie, E., & Wilson, R. (2020). Sheltering in place and domestic violence: Evidence from calls for service during covid-19. *Journal of Public Economics*, 189, 104241.

- Macmillan, R., & Gartner, R. (1999). When she brings home the bacon: Labor-force participation and the risk of spousal violence against women. *Journal of Marriage and the Family*, 947–958.
- Mervosh, S., Lu, D., & Vanessa, S. (2020). See which states and cities have told residents to stay at home.
- Mincer, J., & Polachek, S. (1974). Family investments in human capital: Earnings of women. *Journal of political Economy*, 82(2, Part 2), S76–S108.
- Morgan, A., & Boxall, H. (2020). Social isolation, time spent at home, financial stress and domestic violence during the covid-19 pandemic. *Trends and Issues in crime and criminal justice*, (609), 1–18.
- Munasinghe, L., Reif, T., & Henriques, A. (2008). Gender gap in wage returns to job tenure and experience. *Labour economics*, 15(6), 1296–1316.
- National vital statistics system, natality on cdc wonder online database [Data are from the Natality Records 2016-2020, as compiled from data provided by the 57 vital statistics jurisdictions through the Vital Statistics Cooperative Program]. (2020).
- Neuman, S. (2020). Global lockdowns resulting in ‘horrifying surge’ in domestic violence, u.n. warns. *NPR*. <https://www.npr.org/sections/coronavirus-live-updates/2020/04/06/827908402/global-lockdowns-resulting-in-horrifying-surge-in-domestic-violence-u-n-warns>
- Neumark, D. (1988). Employers’ discriminatory behavior and the estimation of wage discrimination. *Journal of Human resources*, 279–295.
- Nix, J., & Richards, T. N. (2021). The immediate and long-term effects of covid-19 stay-at-home orders on domestic violence calls for service across six us jurisdictions. *Police practice and research*, 22(4), 1443–1451.
- Oths, K., & Robertson, T. (2007). Give me shelter: Temporal patterns of women fleeing domestic abuse. *Human Organization*, 66(3), 249–260.

- Parkinson, D. (2019). Investigating the increase in domestic violence post disaster: An Australian case study. *Journal of interpersonal violence*, *34*(11), 2333–2362.
- Pastor, K. (2021). Google trends county-dma mapping. <https://www.kaggle.com/kapastor/google-trends-countydma-mapping>.
- Phelps, E. S. (1972). The statistical theory of racism and sexism. *The American Economic Review*, *62*(4), 659–661.
- Piquero, A. R., Jennings, W. G., Jemison, E., Kaukinen, C., & Knaul, F. M. (2021). Domestic violence during the covid-19 pandemic-evidence from a systematic review and meta-analysis. *Journal of Criminal Justice*, *74*, 101806.
- Population estimates 2010-2019. (2019). <https://census.hawaii.gov/wp-content/uploads/2021/05/co-est2020-alldata.xlsx>.
- Population estimates 2020. (2019). <https://www2.census.gov/programs-surveys/popest/datasets/2010-2019/counties/totals/>.
- Reaves, B. (2017). *Police Response to Domestic Violence, 2006-2015* (tech. rep.). Bureau of Justice Statistics, U.S. Department of Justice.
- Renzetti, C. M. (2009). Economic stress and domestic violence. *CRVAW Faculty Research Reports and Papers*.
- Sin, I., Stillman, S., & Fabling, R. (2020). What drives the gender wage gap? examining the roles of sorting, productivity differences, bargaining and discrimination. *The Review of Economics and Statistics*, 1–44.
- Smith, S., Zhang, X., Basile, K., Merrick, M., Wang, J., Kresnow, M.-j., & Chen, J. (2015). National intimate partner and sexual violence survey: 2015 data brief.
- Smith, T. W., Marsden, P., Hout, M., & Kim, J. (2020). General social surveys, 1972-2018. *NORC at the University of Chicago*.
- Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using google search data. *Journal of Public Economics*, *118*, 26–40.

- Taub, A. (2020). A new covid-19 crisis: Domestic abuse rises worldwide. *New York Times*.
<https://www.nytimes.com/2020/04/06/world/coronavirus-domestic-violence.html>
- Towers, S., Afzal, S., Bernal, G., Bliss, N., Brown, S., Espinoza, B., Jackson, J., Judson-Garcia, J., Khan, M., Lin, M., et al. (2015). Mass media and the contagion of fear: The case of ebola in america. *PloS one*, 10(6).
- Unemployment insurance weekly claims data. (2021). <https://oui.doleta.gov/unemploy/claims.asp>.
- Zimmerman, J. (2014). No, we don't literally want to ban men. but 2014 was the year women got even. *The Guardian*.