

NORTHWESTERN UNIVERSITY

Essays in Applied Microeconomics

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Economics

By

Susan Ou

EVANSTON, ILLINOIS

June 2018

© Copyright by Susan Ou 2018
All Rights Reserved

ABSTRACT

Essays in Applied Microeconomics

Susan Ou

The first two chapters of my thesis are related to health economics, and explore how individual decisions affecting health can be impacted by different factors, including by government policy. The third chapter of my thesis (coauthored with Heyu Xiong) focuses on public economics in the Chinese context.

In the first chapter, I study the causal impact of neighborhoods on health. Through exploiting variation in the number of years individuals have lived in their neighborhood, I examine the causal effects of exposure to high and low body mass index (BMI) neighborhoods on one's own BMI. The identifying assumption is that there are no unobserved individual level characteristics correlated with both BMI and moving, after controlling for observables. I find that neighborhoods do not have a causal impact on health.

In the second chapter, I study the effect of information provision on consumer welfare in the context of the 2006 trans fat labeling legislation. I develop and estimate a structural discrete choice model, featuring heterogeneity in the valuation of information from the label and taste for trans fats. Revealed preference estimates indicate that consumers would be better off in a labeling regime than a ban regime, though the gains in consumer surplus are small. However, a normative approach, informed by the health costs of trans fats found in the medical literature, suggests that a ban would lead to higher welfare gains than a label.

In the final chapter, Heyu and I study the role of media in the transmission of ideology in the context of the Cultural Revolution. We develop a novel identification strategy by interacting the strength of radio signals and linguistic compatibility of local dialects to the broadcast language, Mandarin. A stronger signal is found to increase revolutionary intensity in counties where Mandarin was better understood. Through investigation of participation in the Send Down Movement, we provide evidence that one mechanism underlying our findings is the direct effect of exposure on individuals rather than differences in local policies induced by media. The effects of propaganda are persistent, as evidenced by Communist Party membership in later life.

Acknowledgements

This dissertation would not have been possible without the help of my advisors, faculty, peers, and friends. I am lucky to have had the support of so many during my time at Northwestern. First, I would like to thank Matt Notowidigdo, who introduced me to the world of applied microeconomics in his labor economics class during my second year. I thank Matt for his guidance throughout my graduate studies and for helping me make the transition from student to researcher. I would also like to thank Lori Beaman and Mar Reguant. Their concrete advice and insights were invaluable to helping me sharpen my work. I also thank Nancy Qian for her patience and generosity of time, and for teaching me how to become a better writer, presenter and researcher.

I am grateful for the encouragement and help from my friends and classmates. They have been endlessly supportive from the first day of graduate school, all the way through to the end of writing my thesis. In particular, I would like to thank Max Tabord-Meehan for his perspective, encouragement and unwavering support.

Finally, I would like to thank my parents for always believing in me.

Table of Contents

ABSTRACT	3
Acknowledgements	5
Table of Contents	6
List of Tables	8
List of Figures	10
Chapter 1. Does Where You Live Affect Your Health? An Analysis of Neighborhood Exposure Effects	12
1.1. Introduction	12
1.2. Data	17
1.3. Empirical Model and Identification Strategy	18
1.4. Results	25
1.5. Robustness	27
1.6. Conclusion	29
Chapter 2. Labels and Consumer Welfare: Evidence from Trans Fat Regulation	31
2.1. Introduction	31
2.2. Institutional Background	38
2.3. Data and Descriptive Statistics	41
2.4. Reduced Form	45
2.5. Demand Model and Identification	48
2.6. Robustness	59

	7
2.7. Welfare	61
2.8. Conclusion	72
Chapter 3. Mass Persuasion and the Ideological Origins of the Chinese Cultural Revolution	74
3.1. Introduction	74
3.2. Background	82
3.3. Data	91
3.4. Empirical Strategy & Results	96
3.5. Mechanisms	105
3.6. Conclusion	116
Tables and Figures	118
Tables and Figures for Chapter 1	118
Tables and Figures for Chapter 2	129
Tables and Figures for Chapter 3	140
References	151
Appendix	164

List of Tables

1.1	Demographics of Movers and Non-Movers	118
1.2	Individual Level Regressions, Selected Regressors	118
1.3	Median BMI in High and Low Residualized BMI Zip Codes	119
1.4	Demographics of Movers to High Residualized BMI Zip Codes	119
1.5	Median BMI in High and Low BMI Zip Codes	119
1.6	Treatment Effects	120
1.7	Explaining Zip Code Residuals	121
1.8	Quantile Regressions	121
2.1	Market Share by Revenue	129
2.2	Popcorn Categories	129
2.3	Summary Statistics	129
2.4	Demographics: Popcorn	130
2.5	Logit Parameter Estimates	131
2.6	Heterogeneity	132
2.7	Robustness	133
2.8	Welfare Gains from Different Regimes	133
3.1	Summary Statistics	140
3.2	Correlates of Signal Strength	140
3.3	Media and Intensity of Conflict: Mandarin Indicator	141

3.4	Alternative Dependent and Independent Variable Definitions	141
3.5	Restricted Samples	142
3.6	Clustering at Different Geographic Units	142
3.7	Additional Controls	143
3.8	Heterogeneity	143
3.9	Summary Statistics	144
3.10	Individual Behavior	145
3.11	Send Down Movement: Mandarin Speakers in Mandarin-Speaking Counties	146
1A.1	Treatment Effects, 5th Percentile	165
2B.1	Prices Per Ounce	169
2B.2	First Stage	169
2C.1	Summary Statistics	171
2C.2	Demographics: Cookies	171
3A.1	Mandarin Comprehension	174
3A.2	Robustness: Alternative Samples & Signal Cutoffs	174
3A.3	Send Down 5-year Estimates	175
3A.4	Falsification Test Individual Behavior	176

List of Figures

1.1	Histogram of Years Living in Current Neighborhood	122
1.2	Individual BMI Distribution	122
1.3	Zip Code BMI Distribution	123
1.4	Zip Code BMI Distribution for non-movers	123
1.5	California BMI Distribution	124
1.6	Zip Code Residual BMI Histogram	125
1.7	California Residualized BMI Distribution	126
1.8	Covariate Stability	127
1.9	Effect of Location on BMI, 95th Percentile	128
2.1	Nutrition Fact Label	134
2.2	Trans Fat Content in Different Product Groups	134
2.3	Demand for Trans Fat Popcorn	135
2.4	Prices of Trans Fat Popcorn	136
2.5	Simple Example: 2 Periods	137
2.6	Simple Example: Many Periods	138
2.7	Demand for Trans Fat Popcorn	139
3.1	Radio Diffusion Exchange	147
3.2	Extent of the Violence during the Cultural Revolution	147
3.3	Predicted Radio Signal	148

		11
3.4	Linguistic Distance to Mandarin	149
3.5	Participation in the Send Down Movement: Coefficients of Mandarin x Age Cohort	150
1A.1	Treatment Effect, 5th Percentile of Residualized BMI	164
2B.1	Popcorn Prices	168
3A.1	Communist Party Membership: Coefficients of Mandarin x Age Cohort	177

CHAPTER 1

Does Where You Live Affect Your Health? An Analysis of Neighborhood Exposure Effects

1.1. Introduction

There exists great regional variation in health across the United States. Residents of some areas are healthier than those of others, by measures of diet, exercise habits, rates of chronic diseases, mortality, and life expectancy. This geographic variation can be due to the causal effects of neighborhoods, or due to the sorting of individuals choosing to live among those who are similar to themselves. In this paper, I explore the role of neighborhood causal effects.

Neighborhoods may affect health through both the social and physical environment (Diez Roux and Mair, 2010). The physical environment refers to the built structural environment in which individuals live. Health and epidemiology literature have found that factors such as being located near supermarkets and grocery stores are beneficial to health (Macintyre et al., 2002). Walkability of a neighborhood (Xu and Wang, 2015) and proximity to parks, recreational areas and green spaces encourage physical activity as well (McCormack et al., 2010). Children with better access to parks have been found to be less likely to experience increases in BMI (Durand et al., 2011). Proximity of health services is also an important determinant of health (Diez Roux and Mair, 2010). The social environment refers to social norms, safety and violence, social stressors, and social

I would like to thank Matt Notowidigdo for his advice and guidance. I would also like to thank Lori Beaman, Lee Lockwood, Cynthia Kinnan, Mar Reguant and participants of the Northwestern Economics 501 seminar for helpful conversations. I am grateful to the Data Access Center at the UCLA Center for Health Policy Institute for data assistance.

cohesion. High crime and social disorder have been associated with poor mental health, and social connectivity has been associated with lower BMI (Diez Roux and Mair, 2010).

In the economics literature, the relationship between the physical environment and health has been found to be less conclusive. Currie et al. (2010) find that living near fast food restaurants may be detrimental to health for schoolchildren and pregnant women. On the other hand, Anderson and Matsa (2011) find that living near restaurants have no impact on obesity. Handbury et al. (2015) find that disparities in food choices among socioeconomic groups persist even when controlling for food access, indicating that policies to improve healthy food access would be ineffective for the very groups these policies would target. Allcott et al. (2017) find that neighborhood environments do not have meaningful effects on healthy eating.

In addition to these environmental factors, peers can also affect health. Christakis and Fowler (2007) study obesity among networks of people. They find that obesity is contagious, and that having friends who are overweight leads oneself to gain weight. Carrell et al. (2011) study peer effects of physical fitness, using randomly assigned friends in the military. They find that peer effects are statistically significant, with the effects concentrated among those with poor fitness.

In this paper, I identify if there exist causal effects of neighborhoods on health, and in particular, on BMI. I do not distinguish between the role of the social or physical environment on health. Many of the previous epidemiological studies linking physical environments and social environments to health are not sufficient to prove causation between health and location. In addition to the causal effects of neighborhoods, there are other reasons why there are differences in health behavior across regions. It is unclear whether neighborhoods, both the social and physical environment, make people unhealthy, or if health is more ingrained and is unchangeable with environment. Choices may reflect our own characteristics, rather than the causal channels of the physical and social environment. Selection effects may be at play, and simply observing correlations between health and healthier environments and healthier friends may not be enough to identify causality.

To explore neighborhood causal effects, I exploit variation in neighborhood exposure, defined as the number of years an individual has lived in his neighborhood. I implement this strategy using the California Health Interview Survey, a dataset from the state of California which documents the number of years individuals have lived in their current neighborhoods. I use median neighborhood BMI as a proxy for the influence of both the social and physical environment. I select zip codes with the largest potential treatment effects, which are those that have a large potential location effect on BMI. This is calculated by examining residuals aggregated at the zip code level of a regression of individual BMI on individual covariates. An area with a large potential location effect on BMI would have a large residualized BMI, meaning the BMI of non-movers is higher than can be explained by observed individual characteristics alone. This indicates that these areas have certain attributes which make them heavier above and beyond what can be predicted by individual covariates such as income, age and education. These attributes can include unobserved neighborhood characteristics, as well as unobserved individual characteristics. I seek to disentangle these effects in my analysis. I then use the zip code level non-mover BMI in these areas as a reference point to analyze how the BMI of movers into these zip codes change over time. As with non-movers, I assume that the BMI of movers can be explained by both individual covariates, as well as a residual term that includes potential location effects. In my study, the difference between initial movers and non-movers in high treatment effect zip codes is around 1.5 points in BMI, or about 10 pounds.

By examining how the parameter on the location effect changes over time as a person lives in his neighborhood for longer, I find that locations barely affect BMI for movers. The BMI of movers does not converge to the BMI of non-movers in areas of high residualized BMI, indicating that location does not affect BMI in high residualized BMI areas. The identifying assumption I make is that conditional on observed characteristics, there are no unobserved characteristics correlated with both moving and BMI. This is a strong assumption, but if the assumption were violated, estimates of the location effect would be

biased away from zero. I find minimal effects of neighborhoods on health, which alleviates concerns of endogeneity in my setting.

One potential explanation for these results is that the geographic variation in BMI is primarily driven by sorting. Additionally, my results do not rule out habit formation in early life. It is possible that neighborhoods may have a causal impact on individuals, but only during one's youth.¹ Neighborhood effects are potentially no longer salient for adults. In the population that I study, I do not find convergence of the BMI of movers to the BMI of the non-movers, indicating that for adults, the role of habit formation may be limited. Adults may not respond to converge to their surroundings because their habits have been potentially ingrained during childhood.

My paper closely complements work studying health outcomes in the Moving to Opportunity (MTO) program. MTO was a random assignment program in which low income individuals were given vouchers to move into higher income neighborhoods. Kling et al. (2007) examine health consequences five years after the implementation of MTO, finding that moving to a lower poverty neighborhood has no effect on physical well-being or economic self-sufficiency, but has substantial effects on mental health for adults and female youth. They do find a substantial effect on obesity but caution that this statistically significant result was one of many analyzed from a family of outcomes from the MTO study, and thus may be a false positive. Ludwig et al. (2013) also examine the effects of MTO on health, but over a longer time horizon of 10 to 15 years. They find that over the longer time horizon, MTO improved physical and mental health but had no effect on economic outcomes. There were no effects on obesity (BMI greater than 30), but cases of extreme obesity (BMI greater than 40) were reduced. I contribute to the literature by studying a broader population, rather than focusing on low income movers who volunteered to take part in the experiment. I also use a treatment which is more specific to the outcome of interest, BMI, rather than a treatment of moving to a lower poverty neighborhood.

¹Due to data limitations, I only study the adult population.

Other papers using a movers strategy to disentangle person effects from place effects include work by Bronnenberg et al. (2012), who study the evolution of brand preferences. By using consumer migration history, and comparing preferences in a state pre-move and post-move, they find that brand preferences are highly persistent. They also use a habit formation model as a lens to interpret the results, and show that over time, a consumers' brand preferences do converge to the brand preferences of the neighborhood they move to. Finkelstein et al. (2014) determine the percentage of variation in health care expenditure attributable to location characteristics versus person characteristics, using patients who move from low utilization areas to high utilization areas. They find little evidence to support the hypothesis of habit formation, since the utilization rate of movers does not converge over time to the utilization rate of the non-movers (there is a discrete jump in rate upon move but no convergence). Chetty and Hendren (2015) study the causal channels of neighborhoods on children's earnings and intergenerational mobility using those who move across counties. Aaronson (1998) estimates the impact of neighborhoods on children's educational outcomes using sibling data to eliminate family-specific heterogeneity associated with neighborhood selection. Handbury et al. (2015) examines changes in access to grocery stores, and Allcott et al. (2017) exploits both supermarket entry and households' moves to healthier neighborhoods to study why wealthier households purchase healthier foods.

This work also builds upon Chetty et al. (2016), who study the association between income and life expectancy in the United States. They find that although life expectancy generally increases with income, the degree of association varies substantially across areas. This indicates that even after controlling for income, there is a high degree of variation in life expectancy. The authors correlate this variation with several location specific characteristics. In my paper, I seek to determine if these differences in neighborhood health are causal and due to location specific characteristics, or if similar people choose to live together. I do not determine which particular location specific characteristics

might be most salient in driving differences in places; rather, I group all of these potential characteristics as a general “location effect.”

The paper proceeds as follows. Section 2 documents the data. Section 3 discusses the empirical framework and identification strategy. Section 4 presents the results. Section 5 presents robustness checks. Section 6 concludes

1.2. Data

I use data from the California Health Interview Survey (CHIS) from 2003 through 2014. The data consists of five rounds of surveys, totaling more than 250,000 observations in the cross section. Respondents answer questions relating to their individual health status and lifestyle habits, including their BMI. Demographic information, such as age, gender, race, zip code of residence, length of time living in neighborhood, country of birth, household income, and education level, is also observed.

My geographic dimension of analysis is on a zip code level. The data includes 1678 unique zip codes, out of a total of 2590 in the state of California.

For further analysis, I also use data documenting the physical characteristics, such as walk score, bike score, and transit score of each zip code.² Walk score is determined based on distance to amenities and pedestrian friendliness, measured by block length and intersection density. Transit score is based on distance to the nearest public transit route, and bike score is measured based on availability of bike infrastructure and number of bike commuters.

1.2.1. Variable Definitions and Summary Statistics

I define individuals to be “non-movers” if they have lived in their neighborhood for greater than 15 years, and “movers” if they have lived in their neighborhoods for less than 15 years. I interpret “neighborhood” to mean zip code, and assume that moving is inter-zip code. Figure 1.1 shows the distribution of length of time that individuals have lived in

²The scores are obtained from the website www.walkscore.com.

their current neighborhood. In the dataset, about 50% of individuals have moved in the last 10 years. Table 1.1 displays summary statistics for movers categorized by years living in current neighborhood, as well as for non-movers. Overall, recent movers are younger, earn less, and are less educated than non-movers and more established movers.

BMI, the health outcome I use, is calculated using the formula, $weight(kg)/height(m)^2$. A BMI less than 18.5 is considered underweight, between 18.5 and 25 is considered normal, between 25 and 30 is overweight, and over 30 is obese. About three quarters of Americans are overweight, and one third are obese. Figure 1.2 shows the distribution of individual BMI and Figure 1.3 shows the distribution of median zip code BMI in the CHIS dataset. Figure 1.5 shows the geographic distribution of median zip code BMI for non-movers, with bluer shades indicating lower BMI areas and redder shades indicating higher BMI areas.

1.3. Empirical Model and Identification Strategy

1.3.1. Regression on Non-Movers: Establishing a Measure of Health

After having established the geographic heterogeneity in BMI, I determine if this heterogeneity is due to the causal effects of place or due to the differences among the people who live there. I do not study what specific factors might cause certain neighborhoods to have higher BMIs than others; rather, I categorize everything into an all-encompassing neighborhood effect. The ideal experiment would be to randomly assign individuals to new neighborhoods, and trace their BMI trajectory over time. Since the ideal experiment is not feasible, I create a research design and make assumptions in order to identify the causal effects of neighborhoods. I assume that in the CHIS dataset, there are no unobserved individual level characteristics correlated with both BMI and moving, after controlling for observables. The CHIS survey also does not track individuals over time. Therefore, the analysis focuses on moves to extreme areas, with a very high degree of “health” and very low degree of “health,” where BMI is a proxy for health. Extreme BMI

neighborhoods are necessary to establish that on average, individuals have a BMI below the neighborhood median BMI when they move to very unhealthy neighborhoods, and vice versa when they move to healthy neighborhoods. If instead, I study a neighborhood of average health, individuals who move to these neighborhoods can come from both higher and lower BMI neighborhoods, which means that only the net effect can be observed. In this case, it would be impossible to distinguish if the neighborhood should have a positive treatment effect or negative treatment effect, since I do not observe the mover's initial BMI. Therefore, I study only neighborhoods with extreme measures of health.

To establish an index for the degree of “healthiness” of a neighborhood, I calculate each zip code's potential location effect on BMI. This location effect is calculated as the aggregated zip code-level residuals from a regression of individual BMI on individual observables for non-movers. I thus attribute the location effect to the unobservables on the zip code level that determine BMI once individual characteristics are controlled for. More precisely, the residuals are from the following regression on non-movers:

$$(1.1) \quad y_{ij} = X_i' \beta + \epsilon_{ij}.$$

Individuals are indexed by i and zip codes by j . y_{ij} is individual BMI, X_i are individual level covariates, including age, education, household income, race, gender, marital status, employment status, size of family, country of birth and year of survey fixed effects. The error term includes unobservables that are idiosyncratic for person i in zip code j . I take the residuals, $\widehat{\epsilon}_{ij}$ from this regression, and calculate the median residuals by zip code for non-movers, $\widehat{\epsilon}_j$. This median zip code residual is a potential location effect that serves as an index for how unhealthy or healthy a neighborhood is, above and beyond individual observable characteristics. The zip code level residuals are ranked, and zip codes at the top of this ranking with the highest residuals are areas which have high BMI, above and beyond what can be explained by individual observables. These areas have high potential treatment effects that are especially likely to raise or lower a person's BMI. An interpretation of this would be that there are environmental and societal factors which

make these places high BMI places, which cannot be explained by observed individual characteristics alone. High residualized BMI places are most likely to have a causal effect on health, since high BMI is not only attributed to the observable characteristics of the people who live there, but also due to unobservables, potentially resulting from location effects. It is also possible that areas with a high residualized BMI have the same individual level unobservables for recent movers. Through my analysis, I determine whether the unobservables are individual-level or neighborhood-level.

The distribution of these zip code level residuals are shown in Figure 1.6. Compared to the distribution of BMI for non-movers, as shown in Figure 1.4, the residuals still exhibit a high degree of variance compared to the histogram of the BMI levels, meaning that even after observables are controlled for, variation across zip codes persist. In addition, the R-squared for Regression (1.1) is 6.5%, which also indicates that observed individual heterogeneity explains a small percent of the variation in BMI. Figure 1.7 shows the geographic variation in residuals. This figure is similar to Figure 1.5, which shows the geographic variation in BMI level across California, although it is interesting to note that some zip codes with high residualized BMI do not have high BMI levels. Table 1.2 presents the estimated coefficients of the control variables from Equation (1.1).

Table 1.3 shows the summary statistics of movers into places of high and low residualized BMI. Individuals just moving (within a year) to high residualized BMI areas have a median BMI of 26.23 (60th percentile of median zip code BMI), and move to zip codes with a median BMI of 27.9 (>95th percentile of median zip code BMI). The difference in BMI is more than 1.5 points on the BMI scale, or more than 10 pounds in weight. This separation between movers and non-movers in high residualized BMI areas enables me to study how the BMI of movers converges to the BMI of non-movers over time in these high residualized BMI areas. These areas are the most likely to have the greatest potential causal impact on health.

Table 1.4 shows the demographics of movers and non-movers to high residualized BMI areas. Compared to Table 1.1, which shows the demographics of all movers, the patterns

are similar. More recent movers are younger, less educated, and have a lower income than those who moved less recently and non-movers. However, movers to high residualized BMI areas are younger, and have a lower level of education and income in each category of time since move in comparison to all movers.

On the other hand, movers to low residualized BMI areas have almost the same BMI as the non-movers of the zip code they are moving to. Movers into the lowest 5% of BMI areas have a median BMI of 24.09, while the median non-mover BMI in these areas is 24.23, essentially indistinguishable from the mover median BMI. The median mover BMI is actually lower than the median non-mover BMI. These levels of BMI are around the 5th to 10th percentile of zip code median BMI. Even after accounting for individual level observables, movers and non-movers into low residual zip codes have very similar BMIs. Thus, I cannot distinguish which zip codes have large treatment effects in low BMI areas.

1.3.1.1. Why Focus on Residuals as a Measure of Health? One might think that a more natural analysis would involve studying movers to zip codes with the highest and lowest BMI levels, instead of residuals. However, using residualized BMI allows me to isolate neighborhoods with a higher BMI due to unobservables, which I attribute to the potential effect of the neighborhood. In addition, using BMI levels instead of residuals does not provide an adequate separation between mover and non-mover BMI. When levels are used, movers to high level BMI neighborhoods have the same BMI as non-movers. Table 1.5 shows the summary statistics of movers into high and low BMI level zip codes. These movers have about the same BMI as non-movers in the destination, in both high and low BMI areas. This limitation makes it impossible to separate out the impact of location on movers, since they have the same BMI as non-movers. Instead, by using residualized BMI, I am able to identify zip codes which have high potential location effects, where movers and non-movers differ in their BMI.

Using residuals to rank healthy and unhealthy places enables me to isolate areas with high potential treatment effects, as well as obtain an adequate separation between mover

and non-mover BMI. I then am able to study how the BMI of movers converges to the BMI of non-movers in high residualized BMI areas.

1.3.2. Regression on Movers: Analysis of Neighborhood Exposure Effects

I restrict my sample to movers who move to the zip codes with the five percent highest zip code residuals since only among this sample is there a separation between median mover BMI and non-mover BMI. I assume that these individuals are moving from areas of lower BMI. As with non-movers, individual mover BMI can be attributed to personal characteristics and a residual which includes the effect of location on BMI. For movers, this location effect describes how BMI responds to exposure to high BMI neighborhoods, which may grow depending on how long an individual has been living in the neighborhood. I examine if and how this location effect changes over time for movers to high residualized BMI areas.

Restricting to movers to high residualized BMI areas, I estimate the following regression specification:

$$(1.2) \quad y_{i\tau} = X_i' \beta + \delta_{2-5} \mathbb{1}(2 < \tau \leq 5) + \delta_{5-10} \mathbb{1}(5 < \tau \leq 10) + \delta_{10-15} \mathbb{1}(10 < \tau \leq 15) + \epsilon_{i\tau}$$

Here, i indexes individuals and τ is length of time living in current neighborhood. X_i contains the same vector of individual observables as in regression (1.1), which are age, education, household income, race, gender, marital status, employment status, size of family, country of birth and year of survey fixed effects. The coefficients of interest are the δ 's, which indicate how each year of living in a high BMI place for movers affects their BMI. These coefficients track how the location effect grows over time. In the base specification, the δ variables are jointly estimated for $\tau \in (2, 5]$, $(5, 10]$, and $(10, 15]$.

The identifying assumption I make is that there are no unobserved individual level characteristics correlated with both BMI and moving, after controlling for observables.

In other words,

$$E(\epsilon_{i\tau}|X_i, M_i) = 0.$$

M is migration status, and incorporates both time in neighborhood, τ , and the choice to move to the current neighborhood. The identifying assumption may be violated if individuals move to neighborhoods where residents are like themselves. Though there may be an initial divide between movers and non-movers to high residualized places, movers may be choosing these destinations because they anticipate changing their habits to match the habits of their destination neighborhood. If this were the case, the estimate of the δ parameters would be upward biased for movers to high residualized BMI areas. Therefore, I assume that the estimates are an upper bound for the effect of neighborhoods on health. If the identifying assumption were violated, estimates would be biased away from zero. However, I still find minimal effects of neighborhoods on health, indicating that endogeneity is not a large concern. I discuss this more in the results section.

I note that the BMI of movers and non-movers can be differentially affected by individual level covariates. Thus, I do not restrict the coefficients of the covariates to be the same between movers and non-movers. I explore robustness to this in the robustness section, where I include non-movers in the regression and do restrict the coefficients on the covariates to be the same.

Another specification I use is a parametric logarithmic regression, where instead of using categorical variables for time, I estimate the coefficient of the continuous $\log(\tau)$ variable, which is the logarithm of time spent in neighborhood. This parametric regression models the process of habit formation. According to psychology literature, habit formation follows an asymptotic curve (Lally et al., 2010). In addition, weight loss has been found to follow the “plateau principle,” where the initial few pounds are easier to lose than subsequent pounds.³

³Mayo Clinic

For completeness, even though there is no separation between mover and non-mover BMI for low residualized BMI areas, I estimate a similar regression for movers to low residualized BMI areas.

1.3.3. Tests of Research Design

I also examine covariate stability of the observables. If observable individual level covariates are not balanced over time spent in current neighborhood, this could be evidence that individual level unobservables are not balanced either, which would give more reason to believe that the identifying assumption is violated. If the unobserved individual level characteristics of movers are changing over time in neighborhood — for example, if individuals’ covariates are changing across time, and these covariates are correlated with unobservables that affect BMI, then this could be a reason why the BMI of movers would change (reasons that are unrelated to neighborhood exposure effects). In order to test for covariate stability, I use a method from Card et al. (2015). I regress the BMI of movers to high residualized BMI areas on time invariant control covariates for all movers (all covariates excluding age and income, which increase mechanically with time), to see if pre-determined covariates evolve over time in a way that would affect BMI. I then plot the predicted \widehat{BMI}_{mover} against years since move. The results are shown in Figure 1.8. The predicted BMI remains stable over the years with a range of about 0.1 BMI points. This is very small compared to the difference in BMI between those who just moved and non-movers, which is 1.67, indicating that there is covariate stability and *changes* in covariates would not be likely to affect BMI.⁴ Therefore any changes in BMI should not be attributed to changes in individual level unobservables over time. Instead, any potential changes in BMI should be attributed to the effect of the neighborhood.

⁴I note that the differences in BMI between movers and non-movers can still be attributed to differences in *levels* of unobservables, as opposed to changes in unobservables.

1.4. Results

1.4.1. Model Estimates

The graphical results of estimating Equation (1.2) for the highest 5th percentile of residualized BMI zip codes are shown in Figures 1.10a and 1.10b. The table versions of these figures are in Table 1.6. Figure 1.10a shows the results from the baseline specification, the non parametric regression where the δ coefficients are jointly estimated for $\tau = 2-5$, $5-10$, and $10-15$, in comparison to the reference group $\tau = 0-2$. Movers who moved to a high residualized BMI zip code within the last two years start out about 1.5 BMI points lower than non-movers, whose median BMI is represented by the continuous solid line at 27.9. This difference is equivalent to about 10 pounds in weight. The solid line shows the point estimate of the change in BMI after 2-5 years, 5-10 years, and 10-15 years, compared to a baseline BMI calculated from an average of 0-2 years of neighborhood exposure. The dotted lines are the 95% confidence intervals. The upper end of this confidence interval is still economically small. The graph shows very little impact of neighborhood median BMI (a measure of the neighborhood's social and physical environment and its influence on health) on the BMI of the mover, implying minimal neighborhood exposure effects. In addition, since the estimates are potentially upward biased, the true point estimate of the effect of location on BMI could be even smaller. Figure 1.10b shows the results of the parametric logarithmic specification, where $\log(\tau)$ is the treatment, instead of a categorical variable for time in neighborhood. The dotted lines indicate the 95 % confidence interval. The coefficient on the logarithm of time is 0.029, a small number, which reinforces the findings from the baseline specification. These results show that the BMI of movers does not converge to the BMI of non-movers. This means that high residualized BMI areas have high residuals most likely due to individual level unobservables rather than due to the causal effects of place.

Graphical results for estimating Equation (1.2) for the lowest 5th percentile of residualized BMI zip codes are shown in Figure 1A.1 for completeness (even though movers have

nearly the same BMI as non-movers). The solid continuous line is the median non-mover BMI, at 24.2. The initial blue line at 0-2 years is the average of BMI among those who moved to a low BMI place within the last 2 years. The subsequent solid blue lines are the treatment effects, or the effect on BMI after 2-5 years, 5-10 years, and 10-15 years, which again are jointly estimated. The table version of this figure is in Table 1A.1.

1.4.2. Explaining the Residuals

I attempt to explain some of the zip code level variation in residuals from Equation (1.1), in order to determine if there are systematic characteristics of zip codes which are driving the differences in residuals which were not accounted for in the individual level regressions. I regress the median zip code level residuals on zip code level controls of physical neighborhood characteristics. The equation I estimate is:

$$(1.3) \quad r_j = Z_j' \beta + \epsilon_j.$$

where r_j is the zip code level residual and Z_j are neighborhood physical characteristics are walk score, bike score, and transit score. These scores are only available for a subset of the high residual zip codes. Results are shown in Table 1.7. A higher walk score is associated with a smaller residual, consistent with the intuition that “healthier” neighborhoods should have smaller residuals. The R-squared of this regression is 14%, indicating that these zip code level controls somewhat explain the variation in residuals.

In addition, I attempt to explain the variation in the residuals that can be attributed to the fixed effects of place. In order to do this, I regress individual residuals from Equation (1.1) on a place fixed effect:

$$(1.4) \quad r_{ij} = \gamma_j + \epsilon_{ij}.$$

where γ_j is the zip code fixed effect. I find that the R-squared of this regression is very small, around 1%, indicating that place effects do not explain individual variation in the

zip code level residuals. This further points to evidence that high residualized BMI areas have unobserved individual-level characteristics, rather than unobserved location-level characteristics.

1.5. Robustness

1.5.1. Changing the Definition of Movers

Instead of defining movers to be those who have lived in their neighborhoods for less than 15 years, I change the definition to be those who have lived in their neighborhoods for less than 10 years. I then recalculate the zip codes with high residualized BMI using the new definition of movers. I find similar results, that the BMI of movers does not converge to the BMI of non-movers, as shown in Figure 1.10c and Table 1.6.

1.5.2. Quantile Regression

Up to this point, I have focused on examining average treatment effects. However, the average may not be reflective of the impact of moving on different quantiles in the distribution of BMI. In order to investigate heterogeneous treatment effects at different points in the distribution of BMI, I estimate quantile regressions for the 10th, 20th, 50th, 80th and 90th quantiles instead of OLS regressions of Equation (1.2). Results are shown in Table 1.8. Even though the OLS results show very small neighborhood effects, the top 90th percentile of high BMI movers actually experience a much larger effect on BMI than the rest of the movers. The largest effect on BMI for those who have lived in their zip codes for between 2 and 5 years is seen in the upper decile, with a point estimate of 0.511, while the estimates for the other deciles are between -0.1 and 0.2. Even though on average, there is a minimal effect of neighborhood on BMI, those at the upper decile are affected the most and do experience a higher increase in BMI than other movers. However, the BMI of movers still does not converge to the BMI of non-movers in these

high residualized BMI neighborhoods. I also note that the standard errors are large, so the parameters are imprecisely estimated.

1.5.3. Restricting Movers and Non-Movers To Have The Same Coefficients

The baseline specification is restricted to movers, which allows movers to have different coefficients on the individual level observables from non-movers. The baseline allows the BMI of movers and non-movers to be affected differently by observables. I explore robustness to this specification by restricting the coefficients for non-movers and movers to be the same.

In order to do this, I run the following regression on all residents, both movers and non-movers, of high residualized BMI neighborhoods:

$$(1.5) \quad y_{i\tau} = X'_i\beta + \delta_{2-5}\mathbb{1}(2 < \tau \leq 5) + \delta_{5-10}\mathbb{1}(5 < \tau \leq 10) + \delta_{10-15}\mathbb{1}(10 < \tau \leq 15) + \delta_{15+}\mathbb{1}(\tau > 15) + \epsilon_{i\tau}$$

The parameters are defined similarly as before. τ indexes time in neighborhood. The δ parameters show how additional time in neighborhood affects the BMI of movers. I note that here, there is also a δ_{15+} parameter for non-movers ($\tau > 15$), who are included in this regression specification. As before, X_i are individual level covariates, including age, education, household income, race, gender, marital status, employment status, size of family, country of birth and year of survey fixed effects. The results are shown in Column (4) of Table 1.6. The results are consistent with the baseline specification: additional time in high residualized neighborhoods does not affect movers' BMI, and the BMI of movers to these neighborhoods do not converge to the BMI of non-movers. The coefficients on the categorical variables for time in neighborhood for movers ($\tau \leq 15$) are small and insignificant. The coefficient for non-movers, δ_{15+} , is 0.656, which indicates that of the 1.5 point difference in BMI between movers and non-movers, 0.656 points remain unexplained by observable covariates. This difference can be attributed to a potential location effect;

however, because the BMI of movers do not converge to the BMI of non-movers, it is more likely that this difference reflects unobservable individual level covariates.

1.6. Conclusion

In this paper, I show that the geographic variation in health does not appear to be causal, for those at the upper end of the BMI spectrum. The BMI of movers to zip codes with a high potential location effect, measured through having a high residualized BMI, does not converge to the BMI of non-movers over time. Some plausible explanations for the variation in BMI across neighborhoods include sorting, as well as habit formation in childhood. The zip code level variance in the residuals of the regression of BMI on observables is also high (shown in Table 1.3), implying a high degree of sorting on unobservables across neighborhoods. However, through quantile regression, I do find small neighborhood effects among those at the upper decile of BMI distribution.

The result that neighborhoods generally do not have a causal impact on health are in accordance with Anderson and Matsa (2011), Allcott et al. (2017), and Handbury et al. (2015), who find little impact of neighborhoods on obesity and healthy eating. The results are also in line with the long term outcomes of the MTO study by Ludwig et al. (2013), who find no effects on obesity from this program.

On the other hand, the result that neighborhoods are not important for health behaviors are in contrast to several other studies examining habit formation and the effects of neighborhoods on non-health related outcomes. Chetty and Hendren (2015), Bronnenberg et al. (2012), and Finkelstein et al. (2014) find large location-based effects on earnings, brand preference, and health care utilization. This suggests that health habits are different from earnings potential or the demand for certain brands. Health habits may be particularly difficult to change. People are able to change their preferences for goods over time, but when it comes to choices in their own lifestyle, individuals are slow to change. Future work will also examine the effects of neighborhoods on the health behaviors of

children. This paper focuses on adults, whose behaviors may have already been cemented over time.

In light of these results, the differences in longevity across regions after controlling for income that Chetty et al. (2016) find can be also be attributed to sorting. They correlate the differences with health behaviors and local area characteristics, but it is likely that these local area characteristics do not have a causal impact on health. Rather, “healthier” individuals choose to live in “healthier” areas. In addition, these results suggest that the correlations between neighborhoods and health found in the health and epidemiology literature are not causal. From a policy perspective, these results would suggest that implementing neighborhood based initiatives to increase physical activity or to increase access to recreational areas may not affect the health of residents. Future work involves further investigation of the root factors that cause certain people to be healthier than others.

CHAPTER 2

Labels and Consumer Welfare: Evidence from Trans Fat Regulation

2.1. Introduction

There has been an ongoing debate in food policy about if, and how, to regulate foods that are harmful for health. If consumers are not completely informed about the options among which they are choosing, their decisions may not maximize welfare. A tool that could make information more apparent is mandatory labeling, and recent examples include labeling of genetically modified foods, labeling of added sugar content, and labeling of trans fat content. In addition to food-related policies (Abaluck, 2011; Jin and Leslie, 2003; Mathios, 2000), labeling and disclosure are also important in the context of financial decisions (Agarwal et al., 2014), financial securities (Greenstone et al., 2006), take-up of social programs (Bhargava and Manoli, 2015; Guyton et al., 2016; Manoli and Turner, 2014), health insurance plans (Jin and Sorensen, 2006; Sanbonmatsu et al., 2012; Scanlon et al., 2002), hospital choice (Pope, 2009), energy efficiency (Allcott et al., 2017), and school choice (Figlio and Lucas, 2004; Hastings et al., 2015).

I am grateful to Matt Notowidigdo, Lori Beaman, and Mar Reguant for their guidance and support. I would also like to thank Seema Jayachandran, Cynthia Kinnan, Aviv Nevo, Robert Porter, Max Tabord-Meehan, and participants of the Northwestern Applied Microeconomics seminar for helpful comments. I also thank Michael Jacobson of the Center for Science in the Public Interest for providing me with trans fat data. This research was supported in part through the computational resources and staff contributions provided for the Social Sciences Computing Cluster (SSCC) at Northwestern University. Calculated based on data from The Nielsen Company (US), LLC and marketing databases provided by the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. Nielsen Data Copyright: Copyright 2018 The Nielsen Company (US), LLC. All Rights Reserved

Alternatively, another policy option is a ban on unhealthy ingredients or properties. Examples include bans on large soda sizes, or bans on food additives. Consumers who prefer a product without the unhealthy property could be better off in a ban regime because of the availability of new reformulated products. Banning would also be better for welfare if consumers are still unaware of the presence or full health effects of the harmful ingredient after labeling, and if they would make different decisions if they were aware. However, those who still derive utility from consuming a product despite the negative health costs would be made better off in a labeling regime than a ban. Bans would eliminate potentially favored properties from products and ignore heterogeneity in preferences. Because of this, labeling and disclosure are favored over more draconian measures in the policy realm. Sunstein and Thaler (2003) advocate for disclosure instead of hard regulation, and “nudging” individuals to make welfare-maximizing decisions.

In this paper, I study the effect of information provision in the form of trans fat labeling on consumer behavior. I compare the welfare gains from labeling to the welfare gains from a counterfactual trans fat ban, finding that if preferences are taken as given, a ban would be worse for consumers than labeling, but still better than a no-regulation regime on average. However, I find that if consumers’ true preferences are consistent with a normative benchmark, calculated from the health and medical literature, where I assume that their valuation of the label is consistent with their valuation of health in other contexts, a ban would lead to higher welfare gains than a labeling regime. In particular, I study the legislation enacted by the Food and Drug Administration (FDA) in 2006 mandating the labeling of trans fat content in packaged goods. This legislation is a natural experiment to study the effects of labeling on consumers. I focus on the microwave popcorn market, which was one of the product groups with the highest presence of trans fat (both in content per serving as well as prevalence across brands). During this time period, awareness of the health costs of trans fats were high, but awareness of the trans fat content of particular foods was low before the labeling legislation. In response to the legislation, several brands chose to reformulate their products to eliminate trans

fat. Consumers responded to the legislation as well: I show reduced form evidence that consumers reduced their demand for brands with trans fat.

To conduct a welfare analysis of the labeling legislation, I develop and estimate a random coefficients logit discrete choice model which features heterogeneity in valuation of trans fat information from the label as well as in taste for trans fat. Relative to a reduced form analysis, where it is difficult to account for spillovers to the control group, a demand system captures substitution patterns to all goods, including those without trans fat. The identification strategy uses the trans fat labeling legislation as a natural experiment by exploiting both the change in information from the label, as well as the resulting product reformulation decisions, in a difference-in-differences-style framework. One identifying assumption is that there are no time-varying unobservables correlated with the decision of whether to reformulate. I provide evidence that the decision to reformulate is plausibly exogenous controlling for brand fixed effects, and is unrelated to a brand's consumer base. To identify consumer valuation of the label, I compare the changes in demand between products that always had trans fat to products that never had trans fat (the control group). Products that always had trans fat were unlabeled in the pre-period but were subject to the label in the post period, while the group that never had trans fat were not directly affected by the labeling legislation in either period. Thus, comparing the group that always had trans fat to the group that never had trans fat would recover the consumer valuation of the label.

Another parameter of interest is taste for trans fat. To identify consumer taste for trans fat, I compare changes in trends between products that reformulate to eliminate trans fat, to products that never had trans fat. Changes in demand to the products that reformulated would reflect preferences for taste because these products contained trans fat in the pre-period, but changed to having no trans fat in the post period. These products were also never directly affected by a label. To identify both taste and valuation of the label, an additional necessary assumption is that in the absence of the legislation, the counterfactual trends of product demand would have trended the same way as before the

legislation, conditional on price and brand fixed effects. I provide evidence of this in the paper.

I estimate the average valuation of the label to be $-\$0.013$ per ounce, which means that consumers must be compensated $\$0.013$ more per ounce when an item is labeled as having trans fat. As a reference, the average price of popcorn is $\$0.14$ per ounce. The magnitude of the average valuation of taste is found to be much smaller than the valuation of the label.

Next, I find the welfare gains of the labeling legislation taking preferences as given. I distinguish between decision utility and experienced utility.¹ While a consumer purchases a product according to decision utility, the utility when the item is consumed is experienced utility. In the pre-period, choices are made without the explicit knowledge of trans fat content. Decision utility may not be the same as experienced utility if the consumer would have chosen differently with more information. However, in the post period, I assume that decision utility is the same as experienced utility and that these choices represent the individual's true preferences. To evaluate the welfare gains of the legislation taking preferences as given, I compare the experienced utilities of the choices made before the legislation with the utilities of the choices made after the legislation, treating the post legislation decision utility as experienced utility. If consumers prefer products with trans fat less after they become informed about them, they may not have chosen optimally before the legislation, because they lacked information about trans fat content. The average welfare gain from the label is calculated to be $\$0.002$ per ounce choice situation. Per year, the welfare gain is $\$0.60$, or about 4% of the average spent on popcorn per year. I also find that without the product reformulation induced by the policy, the welfare gain would be much smaller.

I then investigate the welfare changes from a counterfactual ban on trans fat, taking preferences as given by again assuming that the choices made post-2006 represent an

¹This is in the spirit of Kahneman et al. (1997).

individual's true preferences. I assume that no products exit the market in this counterfactual, and that all products reformulate to eliminate trans fat. To simulate this ban regime, I first estimate new prices in a Nash-Bertrand equilibrium. With these new prices, I compare the welfare gain in a ban to a no-regulation regime, using the estimated utility specification which takes preferences as given. I find a welfare gain of \$0.001 per ounce choice situation, when moving from a no-regulation regime to a ban. This indicates that consumers would be better off on average compared to a regulation-free regime, but worse off compared to a labeling regime.

Consumers who prefer products without trans fat can be made better off in a ban regime, where trans fat-free versions of products replace the versions with trans fat. This holds if the reformulated products without trans fat are not so much more expensive than the products with trans fat, so that the utility gain from not having trans fat is offset by the increase in price.² However, consumers who highly value the taste of trans fat are made worse off in a ban regime, since products they prefer with trans fat may no longer contain them. In this particular situation, even though on average consumers place little value on taste, there are some consumers for whom taste matters greatly. These consumers drive the result that a ban would be worse than a labeling regime for consumers on average. Labeling would be best because it enables both those who prefer the taste of trans fat as well as those who value their health highly to choose products with and without trans fat, respectively, even if their favorite product is not available in the version they prefer.

Finally, I construct a normative benchmark to evaluate how consumers would value the trans fat label if their true preferences were consistent with their valuation of health in other contexts. I use estimates from the medical and value of statistical life literatures to inform my calculations of the health costs of trans fat. I assume and provide evidence that consumers were not aware of the trans fat content in their food before the labeling

²I note that in a world where brands are homogeneous and are perfect substitutes, any regime where both reformulated and non reformulated products exist would be best for welfare. In my setting, brands are imperfect substitutes, and their reformulation decisions contribute to the welfare changes.

legislation. The normative benchmark valuation of the trans fat label is calculated to be $-\$0.41$ per ounce, compared to revealed preference estimates of $-\$0.013$ per ounce, indicating that consumers vastly undervalue the label. This discrepancy between estimated consumer valuation of the label and the benchmark valuation can exist due to several reasons. Consumers may not be informed of the label, or even if they are informed of the label, may not be fully informed of the health consequences. Alternatively, consumers may be informed of the health consequences of trans fat, but may exhibit time inconsistent behavior.

Welfare estimates using this normative benchmark suggest that a ban on trans fat would lead to a much higher consumer surplus than a label or no-regulation regime. In fact, if consumers behaved according to the benchmark valuation of trans fat, the demand for products with trans fat would fall to zero, forcing producers to reformulate, and resulting in the outcome in a ban. I find that the average welfare gain in a ban using the normative benchmark is $\$0.07$ per ounce choice situation. This is more than twice as large as the welfare gain using the benchmark utility from the labeling legislation (which is calculated to be $\$0.03$).

My paper relates to several strands of literature. Broadly, it relates to papers studying the effects of information provision or information disclosure, in many different contexts, including restaurants, education and healthcare, as outlined by Dranove and Jin (2011). They find several examples of when consumers respond to quality disclosure, but less evidence of responses by sellers to boost quality. More specifically, my paper relates to the literature studying the responses to food-related labeling. Ippolito and Mathios (1991) find that when a regulatory ban against advertising of health claims about fiber was lifted, more high fiber cereals entered the market. Mathios (2000) studies the effect on demand of labeling fat content in salad dressings, finding that demand for salad dressings high in fat declines. Kiesel and Villas-Boas (2013) investigate consumer reaction to the labeling of organic milk, and find that the probability of organic milk purchases increases with a USDA organic seal. Bollinger et al. (2011) find the effects of calorie posting on demand for

Starbucks menu items. They discover that the calorie per transaction amount declines, but this is driven primarily by food purchases, not beverage purchases. Jin and Leslie (2003) examine the effects on restaurant demand after restaurant health grade cards are implemented, finding that the demand for lower rated restaurants declines. Oster (2017) finds that in response to a diabetes diagnosis, customers show significant but relatively small calorie reductions. These results are in accordance with my paper, where I show reduced form evidence of both consumer and producer response to the trans fat labeling legislation.

My paper is also related to work studying welfare gains of food-related policies. Dubois et al. (2017) examine the potato chip market, and simulate a counterfactual in which advertising is banned. They find that the potential health benefits of a ban are offset by lower prices and substitution to other types of junk food. Hut and Oster (2018) find that the average household's food purchases do not respond to disease diagnosis, changes in government diet recommendations, or major research findings. My paper is most similar to that of Abaluck (2011), who studies the effect of the National Labeling Education Act (NLEA), which created the nutrition facts panel. He conducts a welfare analysis of the labeling, finding that consumers are insufficiently responsive to the information about the nutrient content of foods. This result is similar to my paper, where I find that consumers are insufficiently responsive to trans fat information compared to a normative benchmark based on health costs. However, in my paper, I also take into account the supply response when calculating welfare gains to consumers.

More specifically, there have also been papers in the health literature studying the trans fat labeling legislation. Niederdeppe and Frosch (2009) study the effects of news coverage on trans fat demand, finding that consumers reduce their consumption after exposure to news in the short run, but the effects dissipate in the long run. Restrepo and Rieger (2016a) and Restrepo and Rieger (2016b) find the effects of trans fat bans on incidences of coronary heart disease in Denmark and New York City restaurants,

respectively. Mohapatra et al. (2017) find the welfare effects of a counterfactual trans fat ban, but do not take into account changes in taste.

I build on the previous literature by studying the introduction of labeling to an unambiguously unhealthy but potentially tasty ingredient and the associated regulatory counterfactuals. It is difficult to think about alternatives to labeling when it comes to nutrition in general, as some properties of food, such as calories, are not inherently bad. By leveraging supply changes, I estimate consumer taste for the unhealthy ingredient, and take this into account in my counterfactual analyses for alternative regulatory policies. Without understanding how much consumers value the taste of trans fat, it would not be possible to understand the effects of a ban on trans fat. The framework developed in my paper can be used to evaluate the gains from information provision in broader contexts as well.

The paper proceeds as follows. Section 2 provides institutional background. Section 3 describes the data and descriptive statistics. Section 4 shows reduced form evidence of the effects of the trans fat legislation. Section 5 describes the structural model, Section 6 describes robustness checks, and Section 7 describes the welfare analyses. Section 8 concludes.

2.2. Institutional Background

In this section, I provide background information on the details of the trans fat legislation. I summarize what trans fats are, as well as why they are appealing to producers and might be unappealing to consumers.

In 2003, the FDA mandated that by 2006, conventional packaged foods and dietary supplements with more than 0.5 grams of trans fat per serving must declare the amount of trans fat in the nutrition facts label, under total fat.³ Foods with less than 0.5 grams of trans fat per serving could state 0 grams on the label. As a result of this legislation,

³We may worry that the firms' response would be to strategically change their serving sizes, but I do not find evidence of this.

many companies reformulated their products using different types of oils to substitute for trans fats.⁴ This was the first major overhaul of the nutrition panel since its inception in 1990. Figure 2.1 shows an example of how the added label appears on the nutrition facts panel.

Trans fat, or partially hydrogenated oil (PHO), are artificially produced from vegetable oils through a process of hydrogenation, or adding hydrogen atoms to the molecule of fat.⁵ Hydrogenation causes the fat to become solid at room temperature.

2.2.1. Trans Fat and Producers

Trans fats are appealing to producers because they enhance the taste and texture of goods, and prolong shelf life. This makes trans fat a cost-effective ingredient. Trans fats were widely used in shortenings and margarine products, as well as cookies, baked goods, frosting, and salty snacks, before the FDA legislation and widespread public awareness of the health consequences of trans fat. Naturally, the largest opponents of the labeling legislation were packaged food lobbies. Producers worried that labels could turn consumers against products, and claimed that the form of information provided was misleading.⁶ Even after the legislation, some producers were reluctant to reformulate because they faced large costs to do so, and the risk that consumers would not like taste of the new product. Figure 2.2 shows the average amount of trans fat per serving across different product groups in the years preceding the legislation. In this paper, I focus on the microwave popcorn market, which was one of the product groups with the highest amount

⁴More recently, in 2015, the FDA legislated that it would no longer be legal to have any amount of trans fat in packaged food products by 2018.

⁵Trans fats are naturally occurring in small quantities in meat and dairy products. Studies have shown that these naturally occurring fats do not have the same health consequences as the artificially produced ones.

⁶For instance, opponents of trans fat labeling argued that an extra claim informing consumers to avoid products with trans fats would be misleading because it might lead consumers to ignore saturated fat content and focus on trans fat.

of trans fat. The following is a quote from a newspaper article about trans fats in popcorns: “‘We’ve mastered it, and I’m not going to tell you how we did it,’ laughed Pamela Newell, a senior director of product development at ConAgra. It took ‘a lot of money,’ she added, since many replacement oil blends limited or reduced the flavor of the popcorn.”⁷ ConAgra is the parent company of the popcorn brand Orville Redenbacher’s.

Different types of food products also faced differences in costs. In another newspaper article, a spokeswoman for ConAgra was quoted as saying “‘Some foods are more challenging than others in regard to the removal of some kinds of fats, due to certain taste or texture expectations...We are using a variety of ingredients and preparation techniques to achieve taste and food quality consumers expect.’”⁸ Within a product category, reformulation costs should be comparable, as the substituted fats are usually the same (Unnevehr and Jagmanaitis, 2008).

In addition to the costs of reformulation, producers were also reluctant to reformulate because of the risk of alienating consumers who were loyal to the taste of a product. In two informal taste tests, Pop Secret, a brand with trans fats at the time, won against many brands without trans fat.⁹ According to a Popcorn Pulse survey by Jolly Time, “nine of 10 respondents watch TV or movies while munching, and 24 percent crunch while on the computer. Some 86 percent consider popcorn a healthy snack, 27 percent look for healthiest/low fat popcorn attributes, yet 84 percent concede that taste is more important than health benefits. More than half the respondents (58 percent) seek intense butter flavor.”¹⁰

The FDA estimated the costs of reformulation to be \$400,000 per product (Department of Health and Human Services, 2003). Part of the argument in passing the labeling

⁷www.csmonitor.com/Business/Latest-News-Wires/2013/1109/Microwave-popcorn-Can-it-survive-a-trans-fat-ban

⁸<http://www.chicagotribune.com/business/ct-trans-fat-food-transition-0628-biz-20150626-story.html>

⁹<https://www.thrillist.com/eat/nation/the-best-microwave-popcorn-money-can-buy>,
<http://www.serious-eats.com/2010/08/whats-the-best-microwavable-popcorn-brand-taste-test-pop-secret-orville-redenbacher-slideshow.html>

¹⁰<http://www.foodprocessing.com/articles/2006/093/>

legislation was that the benefits of labeling in terms of health far outweighed the costs of labeling and reformulation.

2.2.2. Trans Fat and Consumers

Trans fats are unappealing to consumers because of their health consequences. The medical literature indicates that the consumption of trans fats has been linked to many adverse health conditions, with the strongest evidence pointing to elevated risk of coronary heart disease (CHD). Trans fats raise bad cholesterol and triglycerides, and lower good cholesterol. They increase the plasma activity of the cholesterol ester transfer protein, which transfers cholesterol esters from high density lipoprotein (HDL) to low density lipoprotein (LDL) (Mozaffarian et al., 2006). Non-profit groups such as BanTransFats.com and Center for Science in the Public Interest (CSPI) were the largest proponents of the trans fat legislation. A lawsuit was launched by BanTransFats.com against Kraft in 2002, claiming that Oreos were too dangerous for children to eat because of the trans fat content.

Before the legislation, public awareness of the health concerns surrounding trans fat was high. According to a survey administered by the FDA between October 2004 and January 2005 (one year before the legislation), 67% of respondents had heard of trans fat (Lin and Yen, 2010). However, packaged foods were not labeled with the quantity of trans fat per serving. Thus, consumers were most likely unaware of the trans fat content particular food products prior to the legislation. The labeling legislation can therefore be treated as an information shock to consumers.

2.3. Data and Descriptive Statistics

2.3.1. Purchase Data

Purchase data comes from the Nielsen Homescan Panel, accessed through the James M. Kilts Center for Marketing at the University of Chicago Booth School of Business. The panel data starts in 2004 and continues to present day. For the purposes of this project,

I only use data from 2004 to 2007, which spans two years before the legislation and two years after. Later data is not used because there may be concerns that shocks to demand for products more than two years after the trans fat labeling legislation are no longer related to the legislation. Demographic information for each panelist is recorded, including education, household income, number of children, education, and region of residency.

Panelist households are required to scan the bar code of each item they buy, from grocery stores, supermarkets, convenience stores, and other retail outlets for packaged goods. For each transaction, the UPC code, brand description, UPC description, product category, date of purchase, item size, and price paid are observed.

The Nielsen Homescan Consumer Panel is a stratified survey which is balanced on household size, income, head age, head education, presence of children, race, and type of occupation to reflect the distribution of demographics in the United States. Households are randomly recruited, and are incentivized with monthly prize drawings, sweepstakes, and gift points awarded for transmission of data. See Einav et al. (2010) for a validation study of the Homescan Panel.

2.3.2. Popcorn

In this paper, I focus on the microwave popcorn market, one of the snack foods with the highest prevalence of trans fats immediately following the labeling legislation (see Figure 2.2). Even after the legislation, popcorn manufacturers were slow to reformulate their products.¹¹ Popcorn is also a relatively homogeneous product, making substitution across different brands more plausible.

I conduct my analysis at the brand level, because trans fat content varies at the brand level. I categorize the top six brands of popcorn in revenue from 2004-2007 as separate

¹¹For instance, Pop Secret did not reformulate until pressured to do so by the 2018 trans fat ban.

brands, and aggregate up the rest as a basket of “other” brands as a seventh brand.¹² The top six brands of popcorn comprise about 85% of the total revenue of popcorn sold. Generic store brands contribute another 10% share of revenue, and the rest are small brands which contribute less than 1% to the overall share of revenue. I also categorize all low fat varieties from all brands into an eighth separate brand. Low fat varieties are different from the full fat versions, and may target different consumers. Low fat varieties also have less than 0.5 grams of trans fat per serving, which I categorize as not having trans fats. My definition of having or not having trans fats is consistent with the definition used in the labeling legislation. If a product has under 0.5 grams of trans fat per serving, and is thus labeled on the nutrition facts panel as 0 grams, I also treat the product as having 0 grams of trans fat in my data. Table 2.1 shows the market share by revenue of the top brands, along with the average price. Further details about data construction are in Appendix A.

The popcorn brands include those that never had over 0.5 grams of trans fats per serving (Smart Balance and low fat varieties), those that reformulated to replace trans fats in January 2006, when the labeling legislation was enacted (ACT II, Orville Redenbacher’s), and those that always had trans fats in the time period studied (Pop Weaver, Pop Secret, Jolly Time, other). In the rest of the paper, I refer to these groups respectively as *never TF*, *reformulators*, and *non reformulators* (see Table 2.2). Summary statistics about purchases in the popcorn market are shown in Table 3.9.

¹²The total revenue shares were calculated within the Consumer Panel Dataset. I verify that the top brands are similar in the Nielsen Retail Scanner Dataset from 2006-2007, which provides data from point of sale systems. I do not use the Retail Scanner Dataset in my analysis because this dataset starts in 2006, not 2004. In the Retail Scanner Dataset, the top six brands are the same as in the Consumer Panel, except for one brand, Pop Weaver, which has a lower revenue share in the Scanner Dataset. However, the difference is minor, as in both datasets, each additional brand after the top six adds 1% or less of additional revenue share.

2.3.3. Trans Fat Nutrition Data

I construct a novel dataset of trans fat nutritional information. This data was collected by hand through internet searches for historical news articles and images, as well as from the Center for Science through the Public Interest. See Appendix A for details. I do not record the quantity of trans fats. Instead, I record only an indicator for having trans fats or not, at a certain time t . This is because it is very difficult to ascertain the level of trans fat in a product before the labeling legislation. Popcorn with trans fat usually contain 5 grams of trans fat per serving, and a serving is usually around 30 grams. Reformulation of trans fat products occurs generally at a brand level (Pop Secret did not reformulate and Orville Redenbacher's did, for example), and affects all subbrands. Trans fat content can only be observed on nutrition facts panels after 2006, but often, news releases can be found when companies decided to reformulate. Thus, I am able to ascertain whether or not a product contained trans fat before 2006. After 2006, brands that did not have trans fat labeled their products with a "zero trans fat" label. This does not affect the interpretation of the results.

2.3.4. Who Reformulates?

In this section, I discuss the kinds of brands that reformulate. In general, the reformulators were firms with higher revenue (see Tables 2.1 and 2.2). Since reformulation was expensive, it is possible that only the larger firms could afford to do so in time for the enactment of the labeling legislation. Alternatively, larger firms may have also been at risk of lawsuits, and had more incentive to retain a better image.¹³ The brands studied are well known brands, and it is not obvious from public consensus that some brands are superior to others. The typical consumer is most likely unaware about the difference in size of the top popcorn producers, or at least is unlikely to consider size of the manufacturer as a factor in their purchase decisions.¹⁴

¹³Oreo's, the largest cookie brand in revenue, was the target of a lawsuit concerning trans fats in 2003.

¹⁴In various Google searches of "best popcorn brands," the brands I study appear as contenders.

The consumer base of the reformulators and non reformulators were not very different. Table 2.4 shows the average demographics of consumers of reformulators versus non-reformulators in the pre and post periods. The income, education, and presence of children are very similar between the two groups in the pre-period. Furthermore, as evidenced in Table 2.1, there is no systematic relationship between the price of reformulators and non reformulators, though the never TF group does have higher prices than both.

2.4. Reduced Form

In this section, I show evidence that the legislation to label trans fats led to a decline in demand for trans fat products relative to products without trans fat. To do this, I conduct an event study examining the change in demand for trans fat products by running the following regression:

$$(2.1) \quad demand_{pcjt} = \alpha p_{jt} + \gamma TF_{jt} + \sum_t \beta_t quarter_t + \sum_t \delta_t quarter_t * TF_{jt} + \epsilon_{jt}$$

Each observation is at the quarter-brand level (even though household level data is available, for the purposes of this exercise it is more useful visualize aggregate trends, since individual households make few or no purchases in a given brand-quarter). The left hand side, $demand_{pcjt}$, is per capita demand, measured in number of ounces. This is calculated by taking the total weight per quarter-brand in the Nielsen dataset, and dividing by the number of households in the dataset who purchased popcorn at least once over the four year period. p_{jt} is the average price paid for product j in quarter t , TF_{jt} is an indicator for product j having trans fats in quarter t , and $quarter_t$ is a quarter fixed effect. The coefficients of interest are δ_t , the coefficients on the interaction between TF_{jt} and a quarter dummy. These coefficients plot the progression of demand for trans fat products over time in relation to products without trans fat.

I note that the framework here is slightly different from a traditional difference-in-differences, because the treatment and control groups are changing over time, due to the

“reformulators,” who switch from the treatment to control group. These products reformulated in 2006, thus changing from having trans fats to having no trans fats. Therefore, a drop in demand in 2006 can be attributed to product reformulations, in addition to substitution away from trans fat products. To fix ideas, suppose that there was no substitution due to the labeling, and consumers continued to purchase the same quantities of each product as they had purchased before the labeling. This still means that they would be purchasing fewer goods with trans fat after the legislation, due to product reformulation.

Results plotting the δ_t coefficients are shown in Figure 2.3(a) where the dotted lines are the 90% confidence intervals. After the labeling legislation, demand for trans fat products drop. The figure highlights the drop in demand after the legislation. The horizontal red line in the post period represents the difference-in-differences estimate, which is -2.51 with standard error 0.558. For reference, the average demand is 2.52 ounces per quarter-brand per capita.

Next, I find how much of the decline in demand for trans fat products can be attributed to substitution away from trans fat products, as opposed to product reformulation. To do this, I estimate the same regression as Equation 2.1, but add brand fixed effects. This demans the outcome of demand by brand. Product reformulations occur at the brand level. Therefore with brand fixed effects, the demand change is within brand, and due to substitution from or to products with trans fat, as opposed to product reformulation. Referring back to the earlier example, if the demand for all brands remained the same before and after reformulation, there would be no drop in demand for brands with trans fat at the legislation, if brand fixed effects are included. Within brand changes in the treatment group are compared to within brand changes in the control group. The results are plotted in Figure 2.3(b). Essentially, this plots the changes in demand for *brands* with trans fat over time. The horizontal line represents the difference-in-differences estimate, -0.615 with standard error 0.176. This is smaller than the effect without brand fixed

effects, and indicates that the drop in demand for trans fat products is driven more by reformulation rather than by substitution.

In addition to plotting the regression-adjusted demand for trans fat products, I also present an analogous graph for price per ounce in Figure 2.4. Brand fixed effects are included. Prices for trans fat products do drop after the legislation, but the change is small compared to the an average price per ounce of \$0.14.

There are several reasons why the results in this section should only be interpreted as descriptive. Although the event study is a transparent way to examine trends in trans fat products over time, it may not be the most accurate way to evaluate the quantitative effects of the legislation, and the difference-in-differences estimate may not represent the average treatment effect. In particular, there may be spillovers to the control group when consumers substitute from products with trans fat to products without trans fat. The control group without trans fat thus may not represent the counterfactual demand in the absence of legislation, so the difference-in-differences estimate may be an overestimate. Therefore, I turn to a demand system which can help us to better understand the impacts of the legislation on demand for all the different groups. This demand system will allow substitution to be built into the model, as well as provide a framework to estimate the welfare effects of the legislation and conduct policy counterfactuals. Another reason why the difference-in-differences estimate should not be interpreted as the average treatment effect is because the reduced form trends shown in Figure 2.3(b) show that demand for products with trans fat slopes downward even before the legislation. This could be due to increasing substitution to lower fat popcorn (which did not have trans fat) over time unrelated to the legislation. I take into account this trend in the full demand analysis. This emphasis in this section is the descriptive evidence that a drop in demand for trans fat products occurred when the legislation was enacted in 2006.

2.5. Demand Model and Identification

The reduced form section provides qualitative evidence that consumers (as well as producers) do respond to the trans fat legislation. I complement the reduced form with a structural logit model to evaluate the welfare effects of the labeling legislation and to conduct policy counterfactuals. I construct a discrete choice logit model which features heterogeneity in taste for trans fats as well as valuation of the label, by allowing for random coefficients as well as explicit individual demographic interaction terms. In each time period, household i purchases a product with brand j . The household can also choose to purchase no popcorn at all, which is denoted choice 0. To avoid conditioning on purchase, I aggregate shopping trips into a monthly basis. Each time a particular brand is purchased, that counts as a choice for that brand. The average number of times popcorn is purchased is about 4 times yearly. If in a particular month a household does not make a purchase, I treat this as the “no purchase” or outside option. If households make more than one purchase a month (in less than 5% of observations), I model these as independent observations. The utility of no purchase is normalized to zero. The model is as follows:

$$(2.2) \quad u_{ijtr} = \begin{cases} \alpha_i p_{jtr} + \gamma_i TF_{jt} + \delta_i (TF_{jt} \cdot label_t) + brand_j + year_t + \epsilon_{ijtr} & j \in \{Ref., NRef.\} \\ \alpha_i p_{jtr} + brand_j + year_t + \rho_1 t + \epsilon_{ijtr} & j \in \{Never TF\} \\ \epsilon_{ijtr} & j \in \{outside\} \end{cases}$$

The left hand side is utility of individual i purchasing one ounce of product j in time t in region r . Price of product j in time t and region r is p_{jtr} . Since I only observe the price of the purchased good, and not alternatives, I impute the price of the alternatives by taking the average at the region - quarter - brand level. The country is divided into 9 regions by Nielsen. The variables TF and $label$ are defined as before, where TF_{jt} is an indicator

representing if product j has trans fats in time t , regardless of the label, and $label_t$ is an indicator representing being in the post-2006 labeling regime. The utility specification has three cases, the first for a brand belonging to the reformulators or non reformulators, the second for a brand belonging to the never TF group, and the third for the outside no purchase option. Only the first case includes terms with TF , since the reformulators and non reformulators are the only products that ever contain trans fat. I assume that the brand fixed effects capture all unobserved heterogeneity between brands. I include a time trend ρ_1 on the never TF group to capture a differential change in trends between this group and the treated groups. I also include a year fixed effect on the inside options. I discuss this and further assumptions in the next section.

The parameters α_i , γ_i and δ_i are parameters representing heterogeneity. In particular, $\alpha_i = \alpha_{1i} + X_i\beta_1$, and $\delta_i = \delta_{1i} + X_i\beta_2$ where X_i is a vector of demographics including income, education, and presence of kids. α_i and δ_i include random coefficient parameters, α_{1i} and δ_{1i} , as well as terms governing explicit interactions with observed demographics. α_i captures household i 's price elasticity, γ_i captures household i 's taste for trans fats, and δ_i captures response to the trans fat label. γ_i , the coefficient of TF , captures taste because the TF indicator is turned on whenever trans fat is present in a product, regardless of whether or not consumers are explicitly aware of its presence. δ_i is the additional effect on demand from the explicit knowledge that there is trans fat in a product. I specify the negative of α_{1i} to be distributed lognormal with mean μ_a and variance σ_a . γ_i and δ_{1i} are distributed normal with mean μ_g and μ_d , and variances σ_g and σ_d , respectively. ϵ_{ijtr} is distributed independent Extreme Value Type I, with variance $\frac{\pi}{6}$.

2.5.1. Identification

The parameters to be identified are α , γ , and δ , the coefficients on price, trans fat, and trans fat interacted with labeling, respectively. A necessary assumption is that the error term is independent of the explanatory variables. The identification of α , the effect of

price on demand, exploits variation in price within brands over time. Brand fixed effects control for unobserved product characteristics, and a year fixed effect controls for time-varying unobservables that affect all products the same way. The time trend on the never TF group allows for never TF products to become more desirable over time.

In addition, to control for time-varying brand and region-specific demand shocks, I use a Hausman instrument, which is the average price of a good in other markets in that time period. The price of goods in other markets is correlated with own price, but should be exogenous to demand in the same period. The idea of the control function is to control for the correlation of the error term with the endogenous regressor by explicitly controlling for the correlated part of the error. The first stage is to regress price on the instrument and controls, then predict the residuals and run the logit flexibly controlling for the predicted residuals (Train, 2009). In this case, I control for the predicted residual and residual squared. The first stage results all have high F-statistics, and results are shown in Appendix B. The part of the error term that is correlated with price is a function of the predicted residuals from the first stage, so by controlling for this explicitly, the remaining error term should no longer correlated with price.

The identifying assumption is that there are no national-level unobserved taste shocks that are differential across brands over time and correlated with price. However, since prices stay relatively constant throughout the time period studied, as evidenced in Figure 2.4, time-varying shocks are less of a concern.

To identify the label (δ) and taste (γ) parameters, I use ideas from a difference-in-differences framework. We can think of two separate treatments: one, labeling trans fats, and two, reformulating. To identify consumer valuation of the label, I compare the changes in demand between products that always had trans fat to products that never had trans fat (the control group). Products that always had trans fat were unlabeled in the pre-period but were subject to the label in the post period, while the group that never had trans fat were not directly affected by the labeling legislation in either period. Thus, loosely speaking, comparing the group that always had trans fat to the group that never

had trans fat would recover the consumer valuation of the label. To identify consumer taste for trans fat, I compare changes in demand between products that reformulate to eliminate trans fat, to products that never had trans fat. Changes in demand to the products that reformulated would reflect preferences for taste because these products contained trans fats in the pre-period, but changed to having no trans fats in the post period. These products were also never directly affected by a label.

To more concretely discuss identification, I begin with a stylized two-period model without time trends. In Figure 2.5(a), I plot the level shares of the reformulators, non reformulators, outside no purchase group, and never TF groups in the simple two-period model. In this stylized example, the labeling policy enactment occurs in Period 2. Due to this policy, the outside no purchase share remains the same, the share of reformulators and non reformulators decrease, and the share of the never TF group increases. In 2.5(b), I plot the associated log of the market shares for the non reformulators and the never TF group (left), and the reformulators and the never TF group (right). The reason that the log of the shares is the plotted outcome of interest is because in a logit, the log of the probability odds is linear in the parameters of the model. In particular, $\log \frac{P(\text{treated})}{P(\text{never TF})} = v_{\text{treated}} - v_{\text{never TF}}$, where v_{treated} is the mean utility of consuming a brand in the treated category, and $v_{\text{never TF}}$ is the mean utility of consuming a brand in the never TF category. $v_{\text{treated}} = \alpha_i p_{jtr} + \gamma_i TF_{jt} + \delta_i (TF_{jt} \cdot \text{label}_t) + \text{brand}_j + \text{year}_t$, and $v_{\text{never TF}} = \alpha_i p_{jtr} + \text{brand}_j + \text{year}_t + \rho_1 t$. In the simple model, there are no prices, so the only parameters are TF , $TF \cdot \text{label}$, a time fixed effect, and a fixed effect for the category of product. The average probability that a particular product is purchased over the population is the market share.

To identify δ , the label parameter, we can think of a standard difference-in-differences setup, ignoring the reformulators for now. The idea is to compare the trends between the non reformulators who always had trans fat with the never TF group. Since the non reformulators always had trans fat, and changed from having no label to having a label, but the never TF group was never labeled with trans fats, the difference-in-differences style comparison between the log shares of the non reformulators and the never TF group

will recover the estimate for δ . This is shown graphically in the left panel of Figure 2.5(b). The counterfactual trend for the non reformulators is plotted with dotted lines. This counterfactual is based on the trend of the never TF group. The estimate for δ is the difference between the observed trend and the counterfactual trend.¹⁵ The parameters are estimated as in a traditional difference-in-differences.

For γ , the taste parameter, the identification is more subtle. Conceptually, to identify the taste parameter, we could compare brands with trans fat to ones without trans fat cross-sectionally in the pre-period. However, since there may be inherent differences in taste across brands, I include brand fixed effects, and instead exploit product reformulation decisions across time to identify the taste parameter. Reformulators change from having unlabeled trans fats to no trans fats, so their change in demand in comparison to the never TF control group would recover the taste parameter. The right side panel of Figure 2.5(b) shows the identification of γ graphically. The difference-in-differences estimate for γ would be the difference between the counterfactual log share (dotted line) and actual log share for the reformulators in period two.

A time fixed effect fixes the share of outside purchases, so the estimates come from variation in the inside shares.¹⁶ In this two period model without time trends, a concern for interpreting γ and δ as causal effects on demand is that there are unobservables in the error term that are correlated with the reformulation decision of brands, or the trans fat content of different brands, that affect demand. To address this concern, I include brand

¹⁵The counterfactual trend is not necessarily the trend in the absence of legislation, if substitution occurred from the treated group to the control group. If the substitution were entirely to the outside option, then the magnitude of the change would be halved. This is internally consistent with the assumption that there is an additional welfare gain from consuming a product without trans fats, rather than not purchasing at all. One may also be worried about “spillovers” when substitution occurs from treatment to control, but the parameters here are not the average treatment effect, but rather parameters in a utility model.

¹⁶I explore robustness to not using a time fixed effect because there may be concerns that the fixed effect will absorb variation from when people substitute to the outside option to avoid trans fats. However, by focusing on variation in inside shares, substitution to the outside shares will still be taken into account, since the inside shares will necessarily be smaller if the outside shares are larger. I explore robustness to using a time trend in both the treated groups and a time trend in the never TF group in the robustness section. It is not possible to use a fixed effect in both control groups, as this will absorb all variation.

fixed effects, and assume that brand fixed effects completely capture unobserved product characteristics that affect demand and are correlated with the reformulation decision. In other words, I assume that there are no time-varying brand-level unobservables that both affect a brand's decision to reformulate and are correlated with demand. To provide evidence that this assumption is plausible, I will show in the next section that the pre-trends in log market share between the reformulators and non reformulators are comparable.

In addition, I assume that the treatment effects of reformulation and labeling are homogeneous across products (i.e, that valuation of taste and label, γ_i and δ_i , are the same across products, though may be different across individuals). If, for example, products that reformulated were targeting consumers who were healthier and the non reformulators were targeting consumers who were more unhealthy, the δ_i and γ_i we measure will be an underestimate than if the reformulation affected all products with trans fats randomly. I discussed in Section 3.4 and demonstrated in Table 2.4 that reformulators and non reformulators are very similar in terms of demographic levels. The differences between reformulators and non reformulators do not vary in a way that would affect how consumers respond to reformulation. The difference between the firms seems to be related to firm size instead. Therefore, the assumption of homogeneous treatment effects is plausible.

Next, I extend the model to a many-period model, where time trends would be a concern. From the reduced form graph in Figure 2.3(b), we see that demand for brands with trans fat slopes downward in the pre-period, compared to brands without trans fat. This trend is unrelated to the labeling legislation. To account for this trend, I assume that the *change* in trends between the treated groups (reformulators and non reformulators) and never TF group should be constant over time in the absence of the legislation. The assumption is that in the absence of the labeling legislation, the difference in the log of the shares between control and treatment groups would have trended the same linear way in the post-period as in the pre-period. In other words, the timing of the legislation is uncorrelated with the linear trend in the difference of log shares, conditional on controls. This would imply that there is no deviation from a linear trend in difference of log shares

in the time periods leading up to the legislation. The δ and γ will be estimated from the deviation from linear trends in the difference of log shares between the treated and never TF groups, in the post period. I show evidence from the data in the next section that motivates this linear pre-trend.

To implement the revised assumption, I include a linear time trend in the utility specification for the never TF group. Since utility is relative in a logit model, this means that the time trend is in reference to the treated groups (reformulators and non reformulators). Again, time fixed effects are included to fix the outside share, and focus on variation in the inside shares. The log of the difference in probability odds between the treated group and the never TF group is $v_{treated} - v_{never\ TF} = \alpha_i p_{j_t tr} + \gamma_i TF_{j_t t} + \delta_i (TF_{j_t t} \cdot label_t) + brand_{j_t} - (\alpha_i p_{j_c tr} + brand_{j_c} + year_t + \rho_1 t)$, and thus should be linear in time, controlling for the explanatory variables.¹⁷

To illustrate these time trends, I show another simple example in Figure 2.6, extending the examples shown in Figure 2.5. In these graphs, the hypothetical legislation occurs at Period 3. The two graphs on the top panel show why a time trend is necessary. The left side graph shows a hypothetical evolution of trends between reformulators, non reformulators and the never TF group. The right side graph shows the same hypothetical evolution, but differenced. The difference is in the log shares between the treated groups and the never TF group. In this first set of graphs, there is no deviation from trend in the log difference of market share between Periods 2 and 3, when the legislation occurs. Thus, we should expect that the the label parameter δ and taste parameter γ should be zero, since the legislation did not affect how the log shares of these groups were trending relative to the never TF group. However, without including a time trend, the estimates for these parameters would be nonzero, because the downward slope in demand would be instead attributed to the parameters of interest instead of to time. With a time trend, the parameters δ and γ would be estimated to be zero, since there is no deviation from trend for which to estimate the parameters.

¹⁷I use a subscript j_t for treated brand and j_c for control brand in the never TF group.

In the center, the two analogous graphs show the log shares and difference in log shares when $\gamma = 0$ but $\delta \neq 0$. The reformulators do not trend differently from the control group, but the non reformulators do. The log shares of the non reformulators declines. Substitution occurs proportionally to both the reformulators and never TF group from the non reformulators. Thus, we would expect that the γ (taste) parameter should be zero, but δ (label) should be nonzero. γ should be zero since there is no differential substitution to this group compared to the never TF group. δ is nonzero because there is substitution away from this group compared to the never TF group. The dotted lines on the right side graph show the counterfactual trends in the absence of legislation. The slope of the dotted lines represent the estimated ρ from the time trend. As is shown, the δ parameter is estimated from the deviation from the counterfactual trend. γ is zero because there is no deviation from trends.

On the two graphs in the bottom panel, I plot the analogous graphs in the case where the legislation leads to a change in trends in both reformulators and non reformulators, leading to nonzero γ and δ . In this case, the log shares of the non reformulators declines, and substitution occurs from non reformulators to reformulators and the never TF group, but disproportionately. The log of market share of the reformulators declines faster relative to the never TF group. This indicates that the reformulators are now less preferred, which is attributed to taste in the model. In the right side graph, the dotted lines represent the counterfactual trends, and the estimates of γ and δ are the deviation of the actual trend from the counterfactual trend.

2.5.1.1. Plotting Trends. In this section, I show visual evidence that the assumptions to identify γ and δ are likely to be satisfied, by plotting the trends of the difference in log shares, between the reformulators and never TF group, and non reformulators and never TF group, after adjusting for price and brand fixed effects. The first assumption is that there are no time varying unobservables correlated with the reformulation decision. I provide evidence that the pre-trends of the reformulators are similar to the non reformulators. In addition, I assume and show that these pre-trends are linear in the log of

the shares. To do this, I conduct an event study by running the following regression:

$$(2.3) \quad \ln marketshare_{jt} = \alpha price_{jt} + \sum_t \phi_t Non\ Reformulator_j * quarter_t + \sum_t \omega_t Reformulator_j * quarter_t + quarter_t + brand_j + \epsilon_{jt}$$

On the left hand side is the log of the market share of brand j in quarter t . Market share is defined as the total instances a particular brand is purchased, divided by the size of the market. The total market size assumes one purchase per month per household. These are the total opportunities for purchase.

I plot the difference in adjusted shares between the treated non reformulators and reformulators, and the never TF group. This is the ϕ_t and ω_t coefficients, since the reference group in the regression specification is the never TF group. The trends are plotted in Figure 2.7, and the assumption of linearity in the pre-trends do seem plausible. In addition, the difference in log shares for reformulators and non reformulators also trend similarly in the pre-period, providing further evidence that there are no time-varying unobservables correlated with the reformulation decision of different brands. See Appendix B for further evidence of similar pre-trends in prices. The log difference in shares is normalized to 0 at one quarter before the legislation. The dotted line shows the estimated time trend using the pre-2006 data.¹⁸ In the pre-period, the estimated trend approximates the trend in the data, and in the post period, it shows the extrapolated counterfactual trends in the absence of legislation. The γ and δ estimates are the deviation from this time trend in the post period. This deviation from the linear trend is where the identification for γ and δ comes from.

¹⁸This time trend is estimated from the aggregated data for visual purposes only. Since the data is aggregated, individual heterogeneity is masked. Therefore, this is not the same time trend as in the full logit specification, which takes into account individual heterogeneity.

2.5.2. Estimation

Estimation is conducted using maximum simulated likelihood (MSL).¹⁹ The likelihood is similar to the likelihood in a regular logit, except the probabilities are simulated for a given value of the parameters of the distribution of the random variables (call these parameters θ). The average simulated probabilities are the average over the likelihood of purchase for particular draws of the random variables. The simulated probabilities are then plugged into a log-likelihood equation. The maximum simulated likelihood estimator is the value of θ that maximizes the simulated log likelihood.

2.5.3. Parameter Estimates

The results are shown in Table 2.5. In the first column, I show a specification using a plain logit specification, with no explicit individual heterogeneity. In the second column, I include explicit heterogeneity by interacting variables with demographics, and then in the third column, I estimate a random coefficients logit model with individual heterogeneity terms. The IV specification is in the fourth column. This is the preferred specification because it addresses the potential endogeneity of prices. It can also be seen that this specification provides the most realistic elasticities.²⁰

Since in a logit, scale and location are unidentified, I normalize the coefficients by dividing by the price coefficient, α_i . This is the marginal rate of substitution between the coefficient and price, or a willingness to pay measure. For each regression specification, I calculate the willingness to pay for trans fats for a consumer with average demographics. This is shown in the bottom panel. In particular, the parameter of interest is $\frac{\delta_i}{\alpha_i}$, which is the willingness to pay for the trans fat label, or the willingness to pay to avoid the health consequences of trans fat, controlling for taste. This is the also the marginal rate of substitution between $TF \cdot label$ and p . By dividing by the price coefficient, this provides a comparison of the reduction in utility from $TF \cdot label$ to the reduction of

¹⁹I use Stata's mixlogit implementation.

²⁰Elasticities are simulated for the third and fourth specifications.

utility from an increase in price. The willingness to pay for the label is the change in price that has the same impact on consumption, or keeps utility constant, given the presence of the trans fat label in a product. It reflects the willingness to pay to avoid trans fat once consumers explicitly know the trans fat content of a product, presumably because of their health consequences. Individuals must be compensated this amount more for a product with labeled trans fat. The willingness to pay estimate from the preferred specification is -\$0.013 per ounce for a consumer with average demographic characteristics. The convention used in this paper is that a negative represents an unwanted quality of the product that an individual must be compensated for.

I also calculate the willingness to pay for the taste of trans fats, $\frac{\gamma_i}{\alpha_i}$. This is found to be much smaller than the willingness to pay for the label. Individuals are willing to pay \$0.002 more per ounce for the taste of trans fat on average.

The comparison between the willingness to pay for trans fat taste and the willingness to pay for the label represents the tradeoffs between taste versus health. If the ratio $\frac{WTP\ label}{WTP\ taste} = \frac{\delta_i}{\gamma_i}$ is large, this means that the utility gain from the taste of trans fats must be large to compensate for a loss in utility from health. If the ratio is small, on the other hand, this means that health costs do not matter as much, as a loss in a unit of utility from health can be compensated for by a small gain in utility from taste.

2.5.4. Heterogeneity

To explore the differences in willingness to pay for the label among different demographic groups, I restrict the sample along dimensions of income, education, frequency of purchase, and smokers.²¹ For each cut of the data, I calculate the willingness to pay for the label and taste of trans fats, for the average consumer within the restricted sample. The results are shown in Table 2.6. Willingness to pay for the label is noticeably higher for those of higher income, higher for those with higher education, and higher for more frequent buyers. The willingness to pay for taste is an order of magnitude smaller than the willingness to pay for

²¹I designate a household as a smoking household if the panelist purchases cigarettes.

the label for most groups. Notably, the willingness to pay for taste is highest for frequent purchasers, indicating that this group values taste highly (although they also value the label more than less frequent purchasers). When the sample is restricted to smokers, the willingness to pay for the label is actually higher than in the unrestricted sample.

2.6. Robustness

In this section, I conduct several exercises to assess the robustness of the results to different specifications. I omit 1 month pre and post legislation, estimate a logit in shares, and use different time trends. Results showing the WTP for the label and WTP for taste for a consumer with average demographics are in Table 2.7. Overall, WTP estimates for the label from the robustness exercises are comparable to the estimates from the base specification. In addition, the WTP for taste are in most cases smaller in magnitude than the label WTP, as in the base specification.

2.6.1. Omitting 1 Month Pre and Post

There may be concerns that some reformulated brands were still selling backstock in January 2006, or if products rolled out their new reformulated brands early. Thus, I test for robustness by omitting December 2005 and January 2006 sales from the data. I find that the average WTP estimate for the label is -0.011 per ounce, which is similar to the base specification. The average WTP for taste is negative, but is very small (I note that this is the average, which means that some consumers still have a positive WTP).

2.6.2. Alternative Time Trends

In this section, I explore robustness to using a linear time trend on the never TF group as well as time fixed effects. First, I replace the time fixed effect with a time trend in the treated groups (reformulators and non reformulators), in addition to the existing linear time trend on the never TF group. This is to address concerns that the time fixed effect

absorbs all variation in substitution to the outside option. I show that the results are very similar when a time trend is included instead of a fixed effect in the treated groups.

I also explore robustness to using just a time fixed effect without any time trend. As expected, the WTP for the label estimate is larger because it does not take into account that the difference between the never TF group and the treated groups was trending downward over time to begin with. Without taking this trend into account, the effect of the legislation may be overestimated.

2.6.3. Shares Logit

In the individual level logit model, only the choice between brands is modeled and taken into account, not size or quantity. There may be concerns that this is obscuring changes to total weight purchased that are independent of the frequency of purchase. Thus, I aggregate the data to estimate a logit using shares, which accounts for total weight purchased. The downside of this approach is that the granularity of the data is reduced, and individual heterogeneity becomes aggregated.

Individuals are assumed to choose the option which gives them the highest utility. Market shares are empirically calculated as the total weight per quarter-region-brand purchased divided by the potential market. I assume that the potential market is 1 ounce of good per household every day, where the total sample of households is those who made at least one popcorn purchase in the four years. These market shares can be thought of as the average of individual choice probabilities. Since in the shares logit, individual level heterogeneity is masked and the level of granularity is coarser than at the individual level, I make the level of granularity finer in the time dimension by using quarter time fixed effects on all inside goods and a time trend at the quarter level on the never TF group. The probability or share of those purchasing popcorn j in time t in region r is:

$$s_{jtr} = \frac{\exp v_{jtr}}{1 + v_{0tr} + \sum_k^J \exp v_{ktr}}$$

After inversion using market shares, we have:

$$\begin{aligned} \ln(s_{jtdr}) - \ln(s_{0tdr}) &= v_{jtdr} - v_{0tdr} \\ &= \alpha p_{jtr} + \gamma TF_{jt} + \delta TF_{jt} \cdot label_t + brand_j + \rho(\mathbf{1}(\text{never TF}) \cdot quarter_t) + \omega_t + \xi_{ijtr} \end{aligned}$$

where 0 is the outside option of no purchase and v_{0tdr} , the mean utility of no purchase, is normalized to 0. I also instrument for prices using the Hausman instrument of prices in other markets described earlier. The parameter ρ is the time trend for products in the never TF group, and ω_t is a quarter fixed effect. I find that the magnitude of willingness to pay for the label, $\frac{\delta}{\alpha}$, is $-\$0.0005$ with standard error 0.005. The WTP for the label estimate is smaller and not statistically significant, possibly because the granularity of the data has been aggregated.

2.7. Welfare

I conduct several different welfare calculations in this section using the baseline estimates from the demand model. First, I find the welfare gain of the labeling legislation taking preferences as given. Then, I find the welfare gain in a counterfactual trans fat ban, taking preferences as given. In this ban, I assume that all producers reformulate and none exit the market. I also solve for new prices in this regime. Finally, I find the welfare gain if individuals behaved according to a benchmark utility, which is calculated from the health costs of trans fat found in the medical and value of statistical life literatures.

2.7.1. Welfare Gain in Labeling Regime Taking Preferences as Given

In this section, I compute the welfare gains from the legislation, taking preferences as given through the lens of the model. To do this, I assume that the decisions made after the labeling legislation reflect the true experienced utility that an individual internalizes.²² More concretely, I distinguish between decision utility and experienced utility. Consumers

²²In the spirit of Kahneman et al. (1997).

use decision utility to decide which good to purchase, but this may be different from their experienced utility. If an individual purchases a good with trans fat before the legislation but stops purchasing the product post-legislation, this suggests that he was not optimizing before the legislation because he changed his decision when he became more informed. Before the legislation, the utility that he experiences is not the same as the utility with which he makes decisions. After the labeling legislation, I assume that there is no longer a discrepancy between experienced utility and decision utility.

To find the welfare gains from the legislation, the idea is to compare the experienced utility of choices pre and post legislation. Utility changes may arise from new information from labeling, which induces different choices, as well as the new availability of reformulated products. Price changes can also induce changes in utility, although I find price changes to be very small. Decisions in the pre-legislation world are made where no reformulation has taken place, using pre-legislation prices. There is a discrepancy between decision utility and experienced utility for products with trans fats. This discrepancy is δ_i . There would be no difference in utility for products that never had trans fat. In the post period, there is no discrepancy between decision utility and experienced utility for any product.

Formally, the welfare gain is:

$$\frac{u^L(\text{decisions under } u^L) - u^L(\text{decisions under } u^{NL})}{MU(\text{income})}$$

where u^L is the estimated label utility specification from the logit model, and u^{NL} is the utility specification pre-legislation. The observed decisions made with a label are assumed to reflect experienced utility. The equation is the difference between the experienced utility of the decision made under the label regime, and the experienced utility of the decisions made without a label. I divide the difference in experienced utility between post and pre-legislation, by the marginal utility of income, to convert the change in utility into a dollar amount.

To calculate the experienced utility of choices pre-legislation without a label, I simulate choices using utility in a regime with no labeling, but calculate the utility of these choices using experienced utility. Choice without the legislation is simulated using pre-2006 uninformed utility:

$$u_{ijtr}^{NL} = \alpha_i p_{jtr} + \gamma_i TF_{jt} + brand_j + \epsilon_{ijt}$$

I note that I keep the year fixed effect and time trend on “never TF” constant when simulating the choices to focus on changes in utility coming solely from the legislation, instead of changes that would have occurred anyway over time. This precludes a welfare increase for all purchases simply for being in the labeling regime (the fixed effects for being in years 2006 and 2007 are small to begin with), because it is impossible to disentangle if changes to utility for popcorn products as a whole are due to idiosyncratic year shocks, or due to the labeling regime. The utility of no purchase is again normalized to 0. I make 500000 simulation draws from the households in the panel dataset in the pre-period at random, using prices in the pre-period. For each household’s decision situation, the option j with the highest u^{NL} is chosen. The utility of this choice is calculated using the experienced utility:

$$u_{ijtr}^L = \alpha_i p_{jtr} + (\gamma_i + \delta_i) TF_{jt} + brand_j + \epsilon_{ijt}$$

I note that no *label* term appears in the experienced utility, because experienced utility does not depend on having a label. It is the internalized utility when an item has trans fat, regardless of whether it is explicitly labeled. The experienced utility of the decision made under the labeling regime is also simulated with 500000 draws, but from the post period panel dataset, with post period prices. In this instance, there is no discrepancy between decision and experienced utility. Decisions are made with u^L (the brand with the highest u^L is chosen), and the utility of that choice is also calculated with u^L . I find that the welfare gain for the average popcorn purchaser is \$0.002 per ounce choice situation,

compared to an average \$0.14 price per ounce of popcorn. Per year, the welfare gain is \$0.60, or about 4% of the average spent on popcorn per year.²³

To determine the importance of information relative to reformulation for welfare, I conduct a similar exercise as above, but hold supply fixed. That is, I assume no change in prices, and no reformulation relative to the pre-period. The idea is to find the welfare gain if individuals were able to decide with a label in the pre-period, holding supply fixed. The welfare gain formula is the same as above, except that the *decisions under u^L* are the decisions in the pre-period instead of post period, with consumers knowing about the trans fat content of products. I find the welfare gain from information alone to be $\$3.6 * 10^{-6}$ per ounce choice situation, indicating that product reformulation, instead of information, drives the gains in welfare from the legislation. This is consistent with the findings in the reduced form analysis, where the decline in consumer purchases of trans fat products were driven more by reformulation than by substitution away from brands with trans fat.

2.7.2. Welfare Gain in Ban Regime Taking Preferences as Given

A counterfactual policy of interest is a ban on trans fats. I simulate the welfare gain in this counterfactual regime, again assuming that consumer behavior in the labeling regime reflects their true preferences and experienced utility. I also assume that in this equilibrium all producers reformulate instead of exiting the market. This is a plausible assumption to make in the popcorn market, as the products with trans fat have a substantial portion of market share. For instance, we would not expect Pop Secret popcorn to exit the market. In addition, Tom Brenna, a professor of human nutrition and chemistry at Cornell University, was quoted as saying in 2003, “‘When the ban takes effect in three years, companies aren’t likely to discontinue a product because they can’t figure out reformulation...Trans fat does not have magical properties,’ he said. ‘There are other

²³In my model, households are simulated to make a purchase decision once per month, and on average each purchase of popcorn is 25 ounces.

ways to do almost everything. ... I guarantee these guys are thinking hard about it.”²⁴ Further evidence is provided by observed events leading up to the 2018 trans fat ban (we do not see brands exiting the market).

The welfare comparison is analogous to the previous section. I compute the utility gain with a ban on trans fats, compared to a no-regulation regime. Ex ante, the welfare gain from a ban is unclear, even compared to labeling. A ban could be better than labeling if no consumers favor products with trans fat, and thus prefer a particular product without trans fat to one with trans fat. The following equation gives the difference in welfare between a ban and no-regulation regime:

$$\frac{u^B(\text{decisions under } u^B) - u^L(\text{decisions under } u^{NL})}{MU(\text{income})}$$

The second term in the numerator is calculated the same way as before, using the distinction between decision and true experienced utility. Again, the assumption is that choices in the post labeling world reflect consumers’ true preferences. The first term in the numerator is calculated using the following utility expression:

$$u_{ijtr}^B = \alpha_i p_{jtr} + \text{brand}_j + \epsilon_{ijt}$$

No trans fat terms appear because trans fats have been eliminated from all products and can no longer enter utility. Choices are made using u^B (the product with the highest u^B is chosen), and the utility of these choices is also calculated using u^B .

However, the prices in a ban are different from observed prices because in this scenario, all firms reformulate. Prices are solved for in a Nash-Bertrand equilibrium, as in Nevo (2001). Using his notation, suppose there are J brands, each owned by a different firm. The firms maximize profits: $\Pi_{jtr} = (p_{jtr} - mc_{jtr})Ms_{jtr}(\mathbf{p}) - C$, where $s_{jtr}(\mathbf{p})$ is the market share of brand j in time t and region r , equal to $\frac{1}{n} \sum_{i=1}^n \frac{\exp(v_{ijtr})}{\exp(1 + \sum_{k=1}^J v_{iktr})}$, where n is the

²⁴<http://www.chicagotribune.com/business/ct-trans-fat-food-transition-0628-biz-20150626-story.html>

number of individuals. v_{ijtr} is the mean utility for person i choosing brand j . C is a fixed cost and M is size of the market.

The first order condition is: $s_{jtr}(\mathbf{p}) + (p_{jtr} - mc_{jtr}) \frac{\partial s_{jtr}(\mathbf{p})}{\partial p_{jtr}} = 0$. The specific formulation of the first order condition changes depending on the regime, through the $s_{jtr}(\mathbf{p})$ shares term that is a function of mean utility. For instance, in a labeling regime, the mean utility would include terms with trans fat, but not in a ban regime. I solve for mc_{jtr} in the pre-labeling period, using the pre-period utility to simulate shares. I then plug the median marginal cost by brand, mc_j , into the first order conditions of a ban regime to solve for prices. The assumption is that marginal costs are the same regardless of trans fat content.²⁵

In this scenario, I find a welfare gain of \$0.001 per ounce choice occasion, or \$0.30 per year, when moving from a no regulation regime to a ban. This means that consumers are better off in a ban than a no-regulation regime, but worse off than in a labeling regime on average. Some consumers prefer products without trans fats to ones that do. Firms do not need to respond to this in the no-regulation regime because consumers are not explicitly aware of trans fat and thus cannot take trans fat information into account when making decisions. The result that consumers are worse off in a ban compared to a labeling regime is driven by consumers who strongly prefer the taste of trans fat. In a labeling regime, some products still contain trans fats, while some reformulate. Labeling can accommodate both those who prefer the taste of trans fat and those who prefer trans fat free products. However, because the same product is never available in both versions, there can always be consumers who will gain from one regime more than another.

2.7.3. Benchmark Willingness to Pay

In this section, I calculate a normative benchmark valuation of the trans fat label to determine whether or not the estimated willingness to pay magnitude of \$0.013 is large

²⁵The marginal costs are very similar when calculated using the pre-period data versus post period data. This is because the prices in the pre-period and post period are very similar. See Appendix B for details.

or small. Then, I use this benchmark willingness to pay to construct a benchmark utility function, with which I reevaluate welfare gains from a label and ban regime. This benchmark reflects consumers' willingness to pay for the information contained in the trans fat label, if their true valuation of the label were consistent with their valuation of health in other contexts. To calculate this normative benchmark, I use the medical literature to calculate the risks of trans fat consumption, which uses data from prospective cohort studies as well as randomized controlled trials. I then map this risk to willingness to pay using VSL numbers from the literature.²⁶ I also assume that consumers were unaware of the presence of trans fat in microwave popcorn before the labeling legislation, and thus were not already internalizing the costs of trans fat pre-2006. As mentioned in Section 2.2, this is a reasonable assumption to make.

The harmful effects of trans fats have been demonstrated in both controlled trials as well as population cohort analyses. See Appendix D for a more thorough discussion of the medical literature.

The baseline figure I use is from the meta-analysis conducted by Mozaffarian et al. (2009), who find that replacing 2% of total daily caloric intake from saturated fat with trans fat leads to a 20% increase in adjusted relative risk of developing CHD over the course of 10-20 years.²⁷ This particular statistic is relevant because saturated fat is the most common type of fat used to replace trans fat when products are reformulated. A 20% increased relative risk of CHD over 10 to 20 years is 1.11% - 2.21% increased relative annual risk. I argue that these figures do not suffer from endogeneity bias since trans fats were actually not known to be harmful, or at least worse than saturated fat, before 1990, when the data were collected. It would not seem reasonable that an individual would

²⁶The value of statistical life is the dollar amount that an individual is willing to pay to avoid probability of death, divided by that probability. It is typically calculated from hedonic wage regressions and compensating differentials. For instance, if individuals are willing to be compensated by \$10,000 to reduce the probability of death by 10%, the VSL is $\$10,000/0.1 = \$100,000$. Different VSLs have been calculated from different contexts, ranging from occupational hazards to helmet and seatbelt wearing, to smoking. These values range from 0.8 million for smokers, to 6.4 million calculated from occupational hazards.

²⁷The 95% confidence interval is 7% to 34%.

choose butter over margarine (which is abundant with trans fat), for example, on the basis of health. In fact, margarine was once marketed as being healthier than butter. Thus, the increased rates of CHD through trans fat consumption over saturated fat consumption can be seen as a reasonable estimate that does not suffer from endogeneity. It may even be an underestimate if health conscious individuals were trying to be healthier.

The next step is to find the percent increase in calories from trans fat from each ounce of product purchased. Since the studies surveyed by Mozaffarian et al. (2009) were conducted over a period of 10 to 20 years, I assume that the 2% increase in calories from trans fat is sustained over the course of the study, and that each incremental calorie from trans fat linearly affects the probability of developing CHD.²⁸ A lower bound on the amount of trans fat in microwave popcorn that contains trans fat is 2 grams of trans fat per serving (most of the products which did not reformulate contained 5 grams per serving), and each serving is around 1 ounce. 2 grams of fat is 18 calories, or 0.9% of a 2000 calorie diet. Each additional ounce of a trans fat product consumed would increase the percentage of calories from trans fat by $0.9\% * \frac{1}{365 * years}$. I calculated above that a 2% increase in daily caloric intake from trans fats leads to a 1.11% - 2.21% annual increase in risk of CHD. For each ounce of product, this results in a $6.94 * 10^{-5}\%$ to $2.76 * 10^{-4}\%$ increase in annual risk of CHD.

In my study, I only focus on the health costs from loss of life from CHD due to trans fat consumption. I also assume that the only path to death from trans fat consumption is through a CHD event. These assumptions provide a lower bound estimate on the risk of death from trans fat consumption. In the Mozaffarian et al. (2009) study, CHD was characterized as death from CHD, or nonfatal myocardial infarction. The risk of death in the year following a first episode of CHD attack is 74% (T and WB, 1971). The

²⁸As discussed in Appendix D, Mozaffarian et al. (2009) cite a study finding a linear relationship between trans fatty acid consumption and LDL cholesterol levels. Stamler et al. (1986) finds that the effect of blood cholesterol levels on CHD are continuous, and Verschuren et al. (1995) find a linear effect of cholesterol levels on CHD.

probability of death from CHD through one ounce of trans fat consumption in one year is thus $6.94 * 10^{-5}\% * 74\%$.²⁹

I map the risk of death from CHD to a VSL of 0.8 million (converted into 2006 dollars). This is the lower bound from a range of studies surveyed by Viscusi and Aldy (2003). The number comes from Ippolito and Ippolito (1984), who observe how demand changes when cigarette smokers became more informed about the health effects of smoking, in the 1970s when knowledge of the risks of smoking became widespread. The VSL of cigarette smokers should be seen as a lower bound benchmark, since cigarette smokers have a higher risk tolerance than the rest of the population.³⁰ In addition, Ippolito and Ippolito (1984) assume that individuals are fully informed of the risk when they decide whether to smoke, when they calculate the VSL of 0.8 million for cigarette smokers. If consumers are less than fully informed, the VSL would be larger. From this VSL, the benchmark consistent willingness to pay for a consumer with average demographics, $\frac{\delta^*}{\alpha_i}$, is calculated to be -\$0.41 per ounce of product purchased to reduce the relative risk of death per ounce purchase of a trans fat good. Since -\$0.41 is higher than the price per ounce of popcorn, this implies that individuals would need to be compensated for consuming popcorn with trans fat. As another benchmark for comparison, The implied VSL from the estimated willingness to pay for the label taking preferences as given is less than 0.2 million.

The magnitude of the benchmark willingness to pay is much larger than the magnitude of the estimated willingness to pay from the data, $\frac{\delta_i}{\alpha_i}$, of \$0.013. Even if we scale this WTP by the percentage who check nutrition fact panels for trans fat before buying food (Eckel et al., 2009), 25%, the estimated WTP magnitude is still too low, at \$0.04. Reasons why consumers may be undervaluing the label include lack of information about the health

²⁹We can also take the view that one ounce purchased does not mean one ounce consumed. The average household size is 2, so if consumption is split evenly, then the risk of death from CHD would be halved from each one ounce purchase.

³⁰Hersch and Viscusi (1990) find that smokers receive a lower compensating differential for risk than non smokers.

costs of trans fat, lack of awareness of the label, time inconsistent preferences, or bounded rationality.

2.7.4. Welfare Gain in Label Regime Using Benchmark Utility

The welfare gains taking preferences as given understates the effects of the legislation, if the post-2006 label utility is still not the true experienced utility. In this section, I reevaluate the welfare gain from the legislation, using the benchmark utility as true experienced utility. The benchmark utility is the utility that uses the benchmark WTP for the label. From the benchmark willingness to pay measure, I calculate the benchmark δ_i^* , the parameter governing response to the trans fat label. To do this, I multiply the benchmark WTP by α_i . The welfare gain from the legislation can then be revisited. Instead of treating the estimated u^L as the experienced utility, the benchmark experienced utility is:

$$u_{ijtr}^* = \alpha_i p_{jtr} + (\gamma_i + \delta_i^*) TF_{jt} + brand_j + \epsilon_{ijtr}$$

Again, I note that no *label* term appears in the experienced utility, because the experienced utility does not depend on the explicit presence of a label. This is similar to the experienced utility taking preferences as given, except δ_i^* replaces δ_i . The welfare gain from the legislation is then:

$$\frac{u^*(decisions\ under\ u^L) - u^*(decisions\ under\ u^{NL})}{MU(income)}$$

where the post and pre-legislation choices are predicted through the model as before, using post and pre-legislation prices respectively. The only difference is that the utility function used to calculate the utility of these decisions is no longer the estimated label utility, but the benchmark utility. More concretely, *decisions under u^L* would be simulated with u^L , where the alternative with the highest u^L is chosen. The utility of this choice is calculated

using u^* . Similarly, *decisions under u^{NL}* would be simulated with u^{NL} , but the utility of these choices would again be calculated using u^* .

I find the welfare gain using the benchmark utility to be \$0.03 per ounce choice occasion, or \$9 per year, which is much higher than the revealed preference welfare gain.

2.7.5. Welfare Gain in Ban Regime Using Benchmark Utility

The potential welfare gain is even higher if individuals behaved according to the benchmark utility function, with the benchmark willingness to pay for trans fat. In this counterfactual, where consumers are fully responsive to the health consequences of trans fats, I find that demand for trans fat products would fall to zero. Even when I simulate a world where prices of trans fat products are zero, keeping prices of non trans fat products the same, demand for trans fat products would still be zero.³¹ Therefore, it is reasonable to expect that no producers would still manufacture products with trans fat if consumers behaved according to the benchmark willingness to pay for the label. This scenario would be the same as the outcome resulting from a ban on trans fats.

The welfare gain from the counterfactual ban is as follows:

$$\frac{u^B(\text{decisions under } u^B) - u^*(\text{decisions under } u^{NL})}{MU(\text{income})}$$

The first term in the numerator is the ban utility of the choices made in a ban, calculated as in Section 7.2. Once gain, $u_{ijtr}^B = \alpha_i p_{jtr} + \text{brand}_j + \epsilon_{ijtr}$. Prices in the ban regime are solved for, as in Section 7.2. The second term is the benchmark utility of the decisions made under a no-regulation regime, calculated as in 7.4.

The welfare gain using the benchmark δ_i^* in the counterfactual ban is \$0.07 per ounce occasion, or \$21 per year. This reflects the welfare gains both if consumers behaved

³¹The prices of the non-trans fat products would likely increase in this scenario, but not likely to a degree where the demand for trans fat products would rise above zero. Strategically, the products with trans fat become irrelevant, and non trans fat products compete among themselves. I do not model this explicitly because it is unreasonable to believe that producers would continue manufacturing products with trans fats in this scenario.

according to their valuation of health, or simply if their true valuation of the label were consistent with their valuation of health, even if they do not behave this way.

2.8. Conclusion

The regulation of ingredients that may be harmful for health has continued to be a heated topic of debate in food policy. In this paper, I find that the consumer surplus resulting from the trans fat labeling legislation is \$0.002 per ounce choice situation. Per year, the welfare gain is \$0.60, or around 4% of the average spent on popcorn for the average buyer. Most of this welfare gain comes from product reformulation rather than from substitution away from brands with trans fat induced by the label. A counterfactual ban regime would make consumers worse off, but still better off than in a no-regulation regime. However, a normative benchmark valuation of the trans fat label calculated from the health costs of trans fat found in the medical literature suggests that the revealed preference valuation of the label is several magnitudes too small. This could be due to several reasons, including lack of information about the label and about the health consequences of trans fat, time inconsistent preferences, or bounded rationality. Using this benchmark valuation, a ban on trans fat is found to be better for welfare than labeling trans fat content. Thus, if we believe that consumers' valuation of the trans fat label should be consistent with their valuation of health in other contexts, a ban on trans fats would be the best way to regulate trans fat content in foods. This recommendation falls in line with the FDA's decision to ban trans fat in 2018. However, if consumers' preferences are taken as given, they will be worse off in a ban than a labeling regime.

Another counterfactual policy of possible interest is a tax, similar to sin taxes on cigarettes or sugar. It would be easy to incorporate a tax into this framework, but given how harmful I find trans fats to be, a tax would need to be so high that demand for trans fat products would fall to zero. The same outcome as a ban would result. One might wonder why a tax would not need to be prohibitively high in other contexts, such as with cigarettes. Cigarette taxes do not map into this framework because the health effects of

cigarettes are more widely known and accurately predicted.³² Many papers assume that consumers rationally consume cigarettes, fully internalizing their health consequences. The trans fat labeling legislation is an exogenous shock that enables me to to examine how consumers respond to information. Without this natural experiment, it would not be possible to understand if consumers have internalized the risks of trans fats.

The framework developed in this paper can also be used to shed light on future nutrition fact labeling decisions. Labeling of added sugar and GMOs have been recent topics of contention. I show that labeling can reduce consumer demand for labeled products, induce product reformulation, and lead to modest welfare gains. However, potential welfare gains could be much higher, given that the risks of trans fat have been well established. Future research would involve investigating the channels for exactly why consumer response to labels may inconsistent with their valuation of health.

³²In fact, Viscusi (1990) finds that individuals overestimate the risks of cigarette smoking.

CHAPTER 3

Mass Persuasion and the Ideological Origins of the Chinese Cultural Revolution

3.1. Introduction

A striking aspect of several regimes during the 20th century is their ability to orchestrate mass campaigns over large geographic areas. Examples include the Red Terror under the Soviets, the organized killings of Jews in Nazi Germany, the “Killing Fields” of Khmer Rouge, and the Great Leap Forward in China. These violent movements frequently relied on the mobilization of civilians who were otherwise only tangentially connected to the political process. Ruling bodies wielded a level of “soft” power and administrative capacity previously unseen in history, which was made possible through the use of new communication technologies that allowed states to project influence across space more easily. Recent literature has emerged discussing how such technology increased civilian participation in many different settings such as the Rwandan Genocide (Yanagizawa-Drott, 2014) and Nazi Germany (Adena et al., 2015; Voigtländer and Voth, 2014).¹ These papers typically focus on short or medium-run outcomes and allow for the possibility that radio exposure matters because it affects both the local environment and because it persuades individuals to act (even absent changes in the local policy or enforcement). Our paper makes

This chapter is coauthored with Heyu Xiong. We are grateful to Georgy Egorov, Joel Mokyr, Matt Notowidigdo, and Nancy Qian for generous advice and feedback on this work. We also thank Lori Beaman, Matthias Doepke, Joe Ferrie, Seema Jayachandran, Ruixue Jia, Cynthia Kinnan, Lee Lockwood, Mara Squicciarini, and participants of the Northwestern Applied Microeconomics Seminar, History Seminar, NBER Conference on the Chinese Economy, and DEVPEEC Conference at UC Berkeley for helpful advice and comments. Heyu Xiong is grateful for financial support from the Balzan Foundation and Northwestern’s Center for Economic History.

¹We will discuss these papers in more detail later in the introduction.

progress on this agenda by examining the short and long-run (up to 45 years after exposure) effect of media in a new context — the Chinese Cultural Revolution — and by providing evidence that individual persuasion is an important mechanism through which radio exposure affects outcomes.

The empirical setting we study is the Chinese Cultural Revolution (1966-1976), a period characterized by collective violence and political persecution. Violence against political opponents was state sanctioned, but perpetrated by ordinary citizens. Estimates of the number of fatalities range from 250,000 to 1.5 million, while the total number of victims, including those persecuted, is greater than 35 million.² In the prelude to the Cultural Revolution, the Communist party developed a sophisticated wired radio infrastructure from which politicized media was regularly broadcasted.

First, we show that state-sponsored media led to more killings during the Cultural Revolution at the county level. Next, using retrospective micro-data, we examine the specific mechanisms through which media could have induced violence. We isolate the direct effect of persuasive communication on individuals by holding the contextual or place-based effects in the area exposed constant. In particular, we provide evidence on the bottom-up dynamics of mobilization by showing that individual responses vary with exposure to media in an environment where top-down policy is uniform. Finally, we investigate the long run consequences of exposure to state-sponsored media on life trajectories.

The main empirical challenge in estimating the effect of media is identifying exogenous variation in exposure. The method for establishing identification that has become standard in the media literature is to use the spatial variation in the predicted quality of broadcast signal. This was first developed in Olken (2009) and later refined in (DellaVigna et al., 2014; Enikolopov et al., 2011; Yanagizawa-Drott, 2014). While compelling, a concern is that the geographic variation in reception may be correlated with unobserved

²The death toll exceeds some of the modern era’s worst incidents of politically-induced mortality, such as the Soviet “Great Terror” of 1937-38, the Rwandan genocide of 1994, and the Indonesian coup and massacres of suspected communists in 1965-66. The death figure also likely understates the true extent of the violence due to the presumably larger number of those who were imprisoned or otherwise persecuted.

factors that can affect the outcomes of interest. We address this concern by introducing a second source of variation: linguistic distance to Standard Mandarin, which was the national language mandated for use in all central broadcasts. Thereby in our setting, the extent of exposure was jointly determined by availability of radio broadcast signal as well as the proximity of the local dialect to the broadcast language.

Our county-level regression specification uses a difference-in-differences strategy to exploit these two sources of variation. This dramatically relaxes the identification assumptions. We examine the contemporaneous effects of radio on local revolutionary intensity. Specifically, we regress the number of killings directly attributed to the Cultural Revolution on the interaction between the county's linguistic distance from Mandarin (based on the primary language in the county), and the strength of radio signal locally. For our estimates to be causal, we require that the unobserved differences between Mandarin and non-Mandarin counties with high broadcast signal to be comparable to the unobserved differences between Mandarin and non-Mandarin counties with low broadcast signal, in a counterfactual world absent of radios.

We argue that this assumption is credible. In China, linguistic differences are plausibly exogenous to determinants of revolutionary behavior. China is composed of hundreds of mutually unintelligible spoken dialects united by a common written script. Relative to other modern nation states with the same level of linguistic diversity, linguistic variation in China stems less from ethnic differences. Linguistic differences reflect historical migration patterns and diffusion of groups that are often no longer salient in modern times.³ The unique intersection of linguistic heterogeneity and ethnic homogeneity lends credibility to the research design.

³For instance, there is no sense that an individual from Shanghai or Guangzhou who does not speak Mandarin is any less Chinese than an individual from Beijing who does (although there are some minority non-Mandarin speaking groups who are not of Han ethnicity — we will address this later in the paper).

The empirical analysis reveals that radio signals induced conflict more in the areas where Mandarin was better understood. We find that a standard deviation shift in exposure leads to more than a quarter standard deviation shift in percent killed in a county. The results are robust to inclusion of varying controls and alternative definitions of treatment.

An important question which follows from our first result is the mechanism through which increased violence occurred. Media can cause violence through two channels: by agitating ordinary citizens directly (bottom-up dynamics), or by motivating local bureaucrats to promote more violent tactics (top-down organization). Media can legitimize individual behavior directly by leading citizens to commit violent acts of their own volition upon listening to the ideological messages (Yanagizawa-Drott, 2014). Local bureaucrats or political agents might also escalate violence, in response to media provocation, by coercing the people to act or by increasing recruitment (Rogall, 2014). In other words, the differences in outcomes can reflect either the differential response on the part of the listening public or differential enforcement of policies across locations arising from media. Follow-up work to Yanagizawa-Drott (2014), such as Rogall (2014), has emphasized the latter channel in the Rwandan context.

To study the specific mechanisms, we examine within-county variation in a complementary outcome pertinent to the Cultural Revolution: participation in the Send Down Movement. The Send Down Movement entailed a program of rustication in which youths were sent down to the countryside to be reeducated alongside farmers. This movement was partly compulsory and enforced through the local government, but partly voluntary as well. By utilizing individual variation *within* small localities — that is to say, by controlling for county fixed effects — we control for the regional differences in enforcement, thus isolating the individual component of participation. We consider participation in the Send Down Movement as a proxy for revolutionary behavior, as it belongs to the bundle of actions endorsed by the state during the course of the Cultural Revolution (Zhou, 2004). The data we use comes from the China Family Panel Studies (CFPS) survey,

where we observe individual decisions to join the Send Down Movement, as well as the dialect spoken at home.

More specifically, to identify the individual component of participation, we exploit a natural experiment generated by the differential receptiveness to media by Mandarin speakers of different birth cohorts during this time period.⁴ The identification framework involves a *difference-in-differences* strategy in which we focus on cohorts who lived through the Cultural Revolution during their youth. We compare the difference in participation between Mandarin speakers and non-Mandarin speakers of that cohort with other cohorts in small geographic cells, where the dimension of media access due to radio signal would be similar. We find that individuals aged 10 to 21 at the start of the Cultural Revolution who understood Mandarin were more likely to participate in the Send Down Movement than their non-Mandarin speaking peers and Mandarin speakers from other cohorts. We interpret this as capturing the differential exposure to propaganda by the Mandarin speakers of this particular cohort. By isolating *within* county variation, we provide evidence that one channel through which media operates is through direct persuasion of individuals.

Another feature of our setting is that we can distinguish direct exposure through personal media or propaganda consumption (proxied by individual Mandarin comprehension) from indirect exposure through interaction with peers (proxied by residing in a predominantly Mandarin speaking county). We find a positive and significant interaction between the two channels, which is consistent with the existence of social interactions or local complementarities. This sheds light on how individual behavior translates into collective action.

Lastly, we move beyond contemporaneous outcomes to examine the long-term consequences of exposure to propaganda on behavior. The previous literature has found that living through Communist regimes has a lasting effect on individual beliefs (Alesina and

⁴The inclusion of location fixed effects removes the variation in radio exposure due to differences in radio signal.

Fuchs-Schündeln, 2007).⁵ We attempt to provide one explanation for why this occurs. In general, studies have found that political communication rarely has long term effects.⁶ However, specific evidence on the persistent effect of propaganda itself is elusive. The political continuity in the Chinese context presents a unique opportunity to study this question. China has not experienced a regime change since the founding of the People's Republic of China, and the party responsible for the Cultural Revolution-era propaganda remains in control.⁷ Hence, we are able to examine whether propaganda is effective at cultivating a more permanent support of the government.

We consider Communist party membership in later life as an outcome. Joining the Communist Party is a competitive process where self initiated applicants are screened based on ideological rigor.⁸ Utilizing the same econometric framework as before, we find that Mandarin speakers who were of the impressionable age cohort during the Cultural Revolution were more likely to join the Communist Party later in life. The evidence suggests that media facilitated the recruitment and supply of party members.

This paper relates to several distinct strands of literature. It complements recent work exploring the political effects of media. DellaVigna and Kaplan (2007), Gerber et al. (2009), and Chiang and Knight (2011) investigate media influence on voting behavior in developed democracies. Enikolopov et al. (2011) and DellaVigna et al. (2014) show effects of media on voting behavior in transitional democracies, namely Russia and Croatia. Adena et al. (2015) attributes the rise of Nazi support partially to the influence of radio propaganda. Much less is known regarding the impact of media in non-democracies. Notably, Yanagizawa-Drott (2014) finds that radio broadcasts encouraging violence during the Rwandan genocide increased militia violence.

⁵Similarly, Voigtländer and Voth (2015) show Nazi indoctrination had persistent effect on fostering anti-Semitism in Germany.

⁶The half life of political advertisement in a US political campaigns is merely one week (Hill et al., 2013).

⁷The post-Mao shift in party policy and leadership did not result in the outright repudiation of Mao and his policies. With the exception of the prosecution of the Gang of Four, public admission of the failure of Mao-era policies were understated. It stands to reason that 1950s propaganda would still be consequential for the current regime.

⁸See background section.

Our paper adds to this existing literature by exploring the effect of propaganda in a novel setting — the Cultural Revolution, which is substantially distinct from the existing studies. Prior papers have shown the capacity of media to exploit pre-existing ethnic cleavages and instigate violence exclusively along the ethnic dimension. Our paper shows the ability of the state to carry out mass violence through media that is not ethnically motivated and unrelated to ethnic predispositions. This setting also allows us to devise a novel identification strategy based on isolating variation in the spoken language in a largely ethnically homogenous environment. The success of indoctrination depended not only on the availability of broadcast infrastructure, but also on the linguistic compatibility of the listening public.

In addition, previous studies have not attempted to study if media affects behavior by influencing local policies or administration, or by influencing individuals directly. We provide evidence of the bottom-up dynamics of mass mobilization by showing that individual responses vary with exposure to media in an environment where bureaucratic enforcement is uniform. We also investigate the long run consequences of exposure to state-sponsored media, which is relatively unexplored in the prior literature. Our setting is particularly well suited to answer this question, since the political environment is relatively stable over the time period we study.

This paper also contributes to the literature on how linguistic diversity shapes economic and political outcomes. Linguistic fractionalization is a barrier to state capacity, hindering the government's ability to implement policy. Outcomes due to this limitation on state capacity are found to have an unfavorable impact on the country. Numerous studies have attributed ethnolinguistic fragmentation to political instability, poor political and economic institutions and low economic growth. For instance, Easterly and Levine (1997) find that GDP growth is inversely related to fractionalization across a large sample of countries. La Porta et al. (1999) show fractionalization is important in determining the quality of government. Alesina and La Ferrara (2000) document that participation in social activities is lower in more ethnically or racially fragmented localities in the United

States. More recently, Michalopoulos (2012) and Bazzi et al. (2017) explore the causes and consequences of ethnolinguistic diversity in the setting of developing countries.

We provide a clear mechanism to the limits of centralization when there is linguistic fragmentation using intra-country evidence. During the Cultural Revolution in China, linguistic differences precluded the state from carrying out its goals by constraining the audience of state sponsored media, limiting the scope of persuasion. This provides new evidence that promoting linguistic homogeneity through standardization of language augments state capacity, broadly conceived.

Finally, our paper adds to the growing empirical literature on the Chinese Cultural Revolution. These studies have typically focused on the outcomes for areas and the cohorts who experienced it (Bai, 2014; Gong et al., 2014, 2015; Kinnan et al., 2015; Meng and Gregory, 2007; Zhou, 2013).⁹ There is relatively little work studying the local determinants of the Cultural Revolution. To our knowledge, we are the first to provide rigorous empirical evidence on the cause of the violence, as well as the motivating factors compelling individuals to be voluntarily rusticated. By showing that media influences ideology and individual behavior during the Cultural Revolution, our study also complements two recent studies about the determinants of ideology (Cantoni et al., 2017) and foreign media uptake (Chen and Yang, 2017) in China today.

The rest of the paper proceeds as follows: Section 2 provides a brief history of radio broadcasting in China and the Cultural Revolution, Section 3 describes our data, and Section 4 explains the contemporaneous effects of media on violence during the Cultural

⁹Meng and Gregory (2007) find that those whose education was disrupted during this time period faced a decrease in lifetime earnings. Within this literature, special attention is paid to the impact of the Send Down Movement specifically. Zhou (2013) explores the long term effects of the Send Down Movement on individuals, finding that those who were sent down in fact have better economic outcomes. Gong et al. (2015) explores the persistent effect of China's the Send Down Movement on beliefs, finding that those individuals who were sent down to work in the countryside are less likely to believe that external circumstances such as luck, control their lives. Gong et al. (2014) also find that these sent-down youth are more likely to experience mental health problems and chronic diseases. Bai (2014) investigates the economic legacies of violence during the Cultural Revolution, finding that more revolutionary regions were slower to industrialize and have a lower GDP. Kinnan et al. (2015) finds the effects of lasting inter-province links created by migration due to the Send Down Movement.

Revolution. An analysis of mechanisms and the persistence of ideology is presented in Section 5. Concluding remarks are offered in Section 6.

3.2. Background

It is beyond the scope of our paper to provide an extensive background of the events of the Cultural Revolution. Instead, in this section we will focus on the particular details that are relevant for our analysis of the relationship between radio propaganda and individual behavior. We document that the Communist government undertook a campaign to mobilize the public into action during this time period. The media, and especially radio, were particularly salient ways to communicate with the masses and facilitating the transmission of state ideology.

The Cultural Revolution was a large-scale political campaign launched by Mao Zedong in 1966 with a purported intent to “cleanse the class ranks of bourgeois elements”.¹⁰ The violence of the Cultural Revolution was pervasive and widespread, especially during the first two years, from 1966 to 1968. Even though much of the writing concerning violence during the Cultural Revolution focuses on urban violence, more recent work has documented the extent of violence in the rural areas (Walder and Su, 2003). Individuals deemed incompatible with the socialist system were persecuted, including intellectuals, senior party officials, rich peasants, teachers and elites. These “class enemies” were subject to public denunciations, forced self criticisms, and beatings if not outright death.¹¹ The violence and political purges during the Cultural Revolution were typically perpetrated by ordinary individuals within a radicalized community rather than agents of the central government. Oftentimes, the perpetrators and victims knew each other.

¹⁰The Cultural Revolution was borne out of factional competition between Mao and other senior party leaders within the Chinese Communist Party (CCP) leadership. Following the failure of the Great Leap Forward, Mao found himself increasingly marginalized and isolated from the political process in the central committee. As a result, Mao launched the campaign in order to purge his political enemies and regain effective decision making.

¹¹The land-owning class and the educated were targeted in this punitive campaign. Even though land had already been redistributed in the earlier movements, individuals who formerly owned land were nevertheless targeted.

The nature of the violence was primitive. Instead of guns and armed weaponry, victims were often beaten to death with blunt objects, or forced to jump off cliffs. The perpetrators of violence were not deemed as criminals, but rather, were accepted as someone who was acting on behalf of the community. These communities were “willing” agents of the central government (Su, 2011).

State apparatuses such as the police and the military were largely paralyzed. Legitimacy during the Cultural Revolution grew from association with Mao. Hitherto, Red Guards, a revolutionary youth organization composed of ordinary civilians, became the paramilitary force of the movement. According to varying sources, between 200,000 and 1.5 million were killed during the Cultural Revolution, and many more were victims of persecution (Walder and Su, 2003).

The Cultural Revolution required mass involvement and compliance at many levels of the public. Because the movement was principally initiated by Mao to re-assert control over the party, it did not have consensus support within the government itself, particularly provincially. Thereby, in implementing the movement, Mao bypassed traditional party structures and appealed directly to the masses. Popular consent and participation constituted the instrument of *de facto* political power. Radio became the means through which the overwhelming mass response was realized.

Throughout the Cultural Revolution, radio allowed for the direct communication between Beijing and the local communities. Radio broadcasts served the dual purpose of communicating Mao’s directions and goals as well as inciting the masses to action and to carry out said goals. State rhetoric heightened the revolutionary fervor among ordinary civilians who were otherwise tangentially connected to the political process. In the next few sections, we provide an overview of the buildup of radio infrastructure in the years leading up to the Cultural Revolution, the content of the broadcasts, and constraints on its effectiveness.

3.2.1. Broadcasting Infrastructure in Communist China

The official use of radio by the Communist Party of China (CPC) dates from September 1945, when it established a radio station within the CPC controlled territory in Ya'an. Following political consolidation in 1949, the regime nationalized existing private stations and developed an extensive network of mass communication that was centrally operated. Detailed instructions regarding its administration and establishment were given in April 1950, when "Decisions Regarding the Establishment of Radio-Receiving Networks" was announced (Pye, 2015).

Similar to the organization of the broadcasting network in the Soviet Union at the time, the Chinese government organized its broadcasting network at three distinct operational levels: central, regional, and local, each of which corresponded to their respective geographic units and political authority. The central tier referred to the Central People's Broadcasting Station, or Radio Beijing. It created programming, in particular national and international news, and dictated policy. Radio transmission from the central station was relayed wirelessly, first to municipal or provincial radio stations, before being re-directed to radio-receiving stations located at each county seat. The local radio stations had to rebroadcast content created by the Central People's Broadcasting Station, and could only add local news content (Liu, 1971).

From the county radio station, wires that carry the broadcasts are extended to the rural villages in its jurisdiction and connected to strategically-situated loudspeakers. Typical locations included floors of manufacturing plants, poles in market places, roofs of government buildings, communes, and dormitories. These public places were chosen for their visibility to facilitate collective listening. The overall engineering schematic is illustrated in Figure 3.1. Strictly speaking, this was not a broadcasting system, but a system of point-to-point radio communication, with dissemination of selected programming at points of reception through means of wired loudspeakers. This system and its predecessors were known as "radio diffusion exchanges" in the Soviet Union (Houn, 1957).

The construction of these exchanges circumvented the lack of ordinary radio equipment. By 1956, there were reportedly only 1,500,000 radio receiver sets capable of receiving programs from medium wave stations nationally, including those that were controlled and operated by the government (Jan, 1967). Because personal receiver sets were costly to manufacture and private ownership was scarce, collective listening via public loudspeakers constituted the bulk of radio reception through the 1950s and 1960s. Wired loudspeakers were more economical to build en masse, and also allowed the Communist government to completely regulate listening habits. Per governmental figures, the total cost for building a new radio diffusion exchange, together with 150 wired speakers, was about 7,000 *yuan* and monthly expenses for operating such exchanges did not exceed 90 *yuan*. In comparison, the cost of 150 regular radios was more than 20,000 *yuan* and their monthly maintenance was estimated at 1,500 to 2,000 *yuan* (Jan, 1967).

In the decade immediately prior to the Cultural Revolution, there was an extensive buildup in the stock of this broadcasting infrastructure. The Third National Radio Broadcasting Conference, held in December 1955, announced a schedule to build more than 1,169 wired radio broadcasting stations in 1956 with 781,942 loudspeakers attached, 80% of which would be installed in villages, starting from a baseline of 107 rural receiving stations and 56 provincial or municipal stations. This conference projected that by the end of 1957 there would be more than 1,800 wired radio broadcasting stations with more than 1,360,000 loudspeakers in villages. The collectivization campaigns and introduction of communes during the Great Leap Forward facilitated the extension of broadcasting into rural areas. This initiative had dramatic local consequences. Loudspeakers were installed in peasants' homes and commune offices. According to government sources at the time, by 1963, 95% of all counties had access to loudspeaker facilities, although this number may have been inflated (Jan, 1967).

Sustained radio operation required a consistent source of electrical power. The expansion in equipment was accompanied by an increase in the electrical power dedicated to broadcasting. The combined strength of the stations in 1952 rose to 475.2 kilowatts;

in 1954 the figure was more than nine times that of 1952. In 1957, total kilowatts had increased by 470% from 1954. However, by the mid 1960s, widespread electrification still eluded significant portions of China. Large power plants were concentrated in coastal locations and regions of Manchuria which were formerly occupied by the Japanese. This was despite the call for the installation of large hydro and thermal plants in Beijing's first Five Year Plan, initiated in 1956. The actuality of the electrification campaign consisted of crude generating plants of every conceivable method: small hydro motors, hand generators, gas motors, wind motors, etc. (Liu, 1964).

Consequently, radio development was confined to areas where the power supply was relatively sufficient. This pattern of provincial and municipal radio stations is confirmed in the data, as radio station presence is most prevalent along the two main railroad routes at the time: from Beijing to Hangzhou, and from Guandong to Hangzhou, as well as near tributaries of major waterways. In some parts of the country, radio infrastructure was bootstrapped from older telephone lines.¹²

Nevertheless, on the eve of the Cultural Revolution in 1964, there was a robust system of radio communication in China, consisting of 141 provincial and municipal stations (including the central station in Beijing), 1,975 rural receiving stations, and approximately 6 million loudspeakers across the country, amounting to 1 loudspeaker per every 160 persons (Latham, 2007; Liu, 1964). Radio was an invaluable tool for state transmissions and projected national authority directly to its intended recipients.

3.2.2. Radio Content and Propaganda

From the outset, broadcasting infrastructure in China was designed expressly for the purpose of mass persuasion. The political leaders were keenly aware of mass media as an instrument for state indoctrination and for the transmission of ideology. Virtually all senior members of the party were actively involved in media and propaganda activities at

¹²However, the invention of the transistor allowed the development of battery-powered loudspeaker systems, obviating the need for electricity for loudspeaker systems (Cook, 2014).

some point in their careers (Volland, 2003). The alacrity with which Chinese government developed its network of state sponsored media following its establishment in 1949 was proclaimed by contemporary Western sources as the “most extensive propaganda effort” in history (Howse, 1960). Both scholars writing in the midst of the Cultural Revolution and researchers writing retrospectively have noted the pervasiveness of politicized rhetoric and have contemplated its role in facilitating the events (Howse, 1960; Jan, 1967; Markham and Liu, 1969).

Initially, radio broadcasting was used in adjunct to the press. Due to persistent illiteracy, the growth in the number and circulation of newspapers from 1952 to 1959 was modest. Despite a sustained plan to eradicate illiteracy, the literacy rate remained at only around 30% among those in the age group of 14 to 40, and substantially lower for older cohorts. Peasants who had become literate often lapsed back into illiteracy after the conclusion of compulsory education. Additionally, due to the varied and oftentimes physically inaccessible terrain, the medium of radio proved a particularly effective channel of persuasion and state communication.

The content of radio programming was highly integrated to national policy and focused public attention to immediate goals of the state. During the Cultural Revolution, political propaganda actively promoted the demagoguery of Mao and emphasized *Maoist* thought over any semblance of orthodox Marxism. The dominant style of broadcasting focused on mass agitation and the coverage of mass campaigns with the goal of arousing support and increasing mobilization. This was especially pronounced from the onset of the Cultural Revolution.

Propaganda constituted a significant portion of the content over the airwaves. On average, radio broadcasts in 1964 would last 435 minutes per day, consisting of program announcements (5%), educational programs (16%), newscasts (29%), weather (2%), agricultural programs (7%), and entertainment (41%).

Newscasts consisted of broadcasts from the Central People’s Broadcasting Station, the official party station, including programs such as “Quotations from Chairman Mao,”

“Selected Reading of Chairman Mao’s Works,” and leaders’ speeches (Jan, 1967). Some examples of quotes from Mao’s Little Red Book include: “A revolution is not a dinner party, or writing an essay, or painting a picture, or doing embroidery; it cannot be so refined, so leisurely and gentle, so temperate, kind, courteous, restrained and magnanimous. A revolution is an insurrection, an act of violence by which one class overthrows another.” This serves as an example of the types of behavior sanctioned and even promoted during the Cultural Revolution. The chief function of news broadcasts was to discredit Mao’s enemies, and they were almost exclusively devoted to sensational exposés of those who were purged.

Entertainment included revolutionary songs and dances, aimed at political agitation. All traditional or foreign cultural influences, including music and opera, were purged, leaving only party propaganda. The media propagandized the “literature of workers and peasants” over high culture. The Communist Party demanded programs be “nationalistic and populist” not “intellectual and foreign” (Liu, 1964).

Agricultural programming consisted of half technical advice and half propaganda, which included speeches from model farmers. In total, propaganda and indoctrination constituted 85% of broadcast time. These radio broadcasts encouraged listeners to participate in the revolutionary cause to build a greater China together (Latham, 2007).

To enforce compulsory listening, a system of collective listening was adopted in villages and communes. This was implemented in two ways: broadcasting assemblies and institutional listening. In the former, heterogeneous audiences of peasants were gathered together in “radio auditoriums” and listened to designated programs in groups, commonly monitored by party cadres. In the latter, loudspeakers broadcast for a set number of hours each day in public places such as governmental offices, factories, and schools, where employees, workers, and students were captive audience members. In the lead up to the Cultural Revolution, the practices of institutional listening became greatly promoted. A report from Shanghai dated August 9th, 1966, the day after Radio Beijing had broadcast the Central Committee’s decision on the Cultural Revolution, stated:

The broad revolutionary people enthusiastically listened to the broadcast of the Central Committee's decision last night. Early this morning, parade columns appeared in major streets of Shanghai... The commune members in the suburbs, who were busy in reaping and planting, listened to the broadcast and were greatly excited (Liu, 1971).

Throughout the Cultural Revolution, radio would remain the direct communication link between the central government and the people, bypassing the interference of local political power and bureaucrats. Local stations were forced to rebroadcast programs from the Central People's Broadcasting Station, and only had power to create local news content. In fact, whenever a purge within a given regional party occurred, the radio stations in that region stopped broadcasting regional content entirely and only broadcast news from Radio Beijing.

The radical politicization of media content in this time period created a favorable climate of opinion for the Cultural Revolution: it heightened the morale of Mao's followers and identified the ideological enemies, whether perceived or real. Exposure to radio increased the impression of universality of the political struggle, induced otherwise apathetic peasants to become state agents, and emboldened them to act. Following the conclusion of the Cultural Revolution, radio broadcast became much more subdued and sanitized.

3.2.3. Language Standardization and Policy

Despite the pervasive infrastructure of the radio network across China, the comprehension of the messages were constrained. The uniformity of content in the broadcasts were also reflected in the uniformity of broadcast language.

In an attempt to combat localism, it was mandated that all official news broadcasts through the wired rural broadcast network be conducted in Standard Mandarin. As a result, in provincial counties where the native dialect was not Mandarin, only an estimated

15% of audience members could actually comprehend the centrally-relayed broadcasts from Beijing (Liu, 1971). Zhou Enlai himself remarked in 1958:

Radio and the cinema are powerful publicity instruments. But as our common speech has not yet been made universal, their effectiveness in the districts where only local dialects are spoken is inevitably limited (Liu, 1971).

China is a linguistically diverse nation, within which the predominant language is Chinese, or *Hanyu*.¹³ Chinese itself refers to a collection of related but often mutually-unintelligible dialects. The varieties of Chinese resemble distinct spoken languages united by a single written script and shared cognates.¹⁴ The varieties of Chinese differ mainly in their phonology and to a lesser degree, syntax and vocabulary. Linguists have categorized the varieties in several different ways, but most agree that there are between seven and ten groups, only one of which is the Mandarin family.¹⁵ There is a general consensus of a North and South division with more pronounced variation in the rugged South.¹⁶

The choice of Mandarin-language broadcast in counties where it was not understood should not be considered an oblivious oversight by a non-optimizing bureaucrat. Rather, it was a calculated decision reflecting the careful tradeoff between the static efficiency of persuasion and the dynamic efficiency of linguistic standardization. A unified language was thought to be instrumental in unified policy for a unified country. Before the ascendance of the Communist party in 1949, there was no national standard language.¹⁷ In the

¹³There are also at least nine groups of non-Chinese languages spoken by ethnic minorities within the present PRC borders.

¹⁴To fix ideas, a helpful analogy can be drawn to variations within the Western Romance languages. But unlike the Romance languages, the differences in Chinese dialects reflect only the spoken form. There is only one written form of Chinese, which would be used for anyone writing or reading Chinese. Henceforth we will refer to the varieties of Chinese as dialects for simplicity.

¹⁵The common agreed upon groups include: Mandarin, Hakka, Cantonese, Wu, Gan, Min, and Hui. This level of categorization subsumes a great deal of underlying differences as there is local variation even within these broad groups.

¹⁶Northern China is composed of flat central plains whereas the South is riddled with mountains and rivers.

¹⁷A common governmental language existed in the form of *Guanhua* during the dynastic periods but it was used only by the upper echelons of bureaucrats and magistrates (Ramsey, 1987).

incipient years of Communist rule, the regime was quite tolerant of minority dialects as the Communist revolution had drawn its support from the largely dialect-speaking rural peasant population. However, by 1955, the government had become highly cognizant of linguistic barriers to national construction. In a little publicized Conference on Standardization of Chinese Language in 1956, a simplified version of Standard Mandarin was devised and promoted as the common tongue.¹⁸

This language policy was implemented in practice through different mechanisms including laws, regulations, education, exams, and restriction of minority language use in public spaces (Barnes, 1982). In 1956, primary school teachers were trained in the standard language and the use of minority dialects in schools or over the airwaves was rebuked. Even the primacy of propaganda was subordinate to the directory of language standardization.

3.3. Data

Our aggregate analysis makes use of a number of datasets, including: (i) revolutionary intensity at the county level, proxied by number of deaths, from Walder and Su (2003), (ii) the predicted strength of radio signal received at each county seat, (iii) the extent of Mandarin intelligibility locally, and (iv) information on county characteristics from historical censuses and gazetteers. The rest of this section describes each of these in detail.

3.3.1. Cultural Revolution Intensity

The main outcome of interest is the intensity of the Cultural Revolution. On the county level, this is proxied by the number of killings due to revolutionary violence. Our analysis uses a county-level dataset on revolution-related fatalities and victims, digitized from

¹⁸The National Language Unification Commission established the Beijing dialect of Mandarin as the standard language of the country in 1932. For expediency, the People's Republic retained this standard when they took power in 1949. However, active enforcement and promotion only began from 1956 onward.

regional gazetteers (Walder and Su, 2003). To the best of our knowledge, this is the most comprehensive dataset of casualties available.

Gazetteers are book-length “encyclopedias” detailing local histories, demographics, economics, etc. Gazetteers covering the Cultural Revolution were published in the late-1980s as a consequence of a central policy directive issued in 1978. Each county was instructed to conduct official investigations into the period in order to rehabilitate “wrongful” victims and compensate remaining family members (Su, 2011).¹⁹ Although the resulting annals contained varying degrees of details, they all included specifics on the number of abnormal deaths attributable to the revolution.

For instance, these deaths include suicides of individuals under persecution, deaths in clashes with military or factions, deaths in struggle sessions or as a result of imprisonment, and executions during political campaigns. Walder and Su (2003) collected this information along with the number of victims in each county, more loosely defined. The authors discuss the issue of data quality extensively. They conclude that the degree of under-reporting in the data should not be correlated with the severity of the Cultural Revolution locally.

Figure 3.2 illustrates the spatial variation in revolutionary intensity for the entire sample. The average percent killed in a county is 0.048%, with a standard deviation of 0.164. As noted in previous work, considerable variation in revolutionary intensity exists (Su, 2011; Walder and Su, 2003).

3.3.2. Radio Signal

We construct a measure of radio reception using information on the location of provincial and municipal radio stations. This data is obtained from Liu (1971), who identified the location of 141 known provincial and municipal stations in 1964, just prior to the

¹⁹National standards were established with much room for interpretation. Some of the annals were very lengthy and detailed, including details on the method of killing, while others were brief and conservative.

Cultural Revolution. Based on this cross sectional data, we apply the Irregular Terrain Model (Hufford, 2002) to calculate the predicted radio signal strength in all localities.

The Irregular Terrain Model (ITM) was originally developed by the US government for frequency-planning purposes and allows one to accurately predict signal strength across narrow geographical cells (Phillips et al., 2011). The model computes the signal loss between transmitting and receiving locations accounting for the physical distance and topography that lies in between. It has also been employed by Olken (2009), Enikolopov et al. (2011), and DellaVigna et al. (2014). To implement the ITM algorithm, we utilize information on radio locations along with a high resolution geo-topographical map of China. For each county, we predict the radio signal strength at the county seat, where historically the county receiving radio stations would have been located.

Because the previous studies have demonstrated that the quality of the broadcasts varies non-linearly as a function of signal strength, we define our explanatory variable, $Signal_c$, as a binary indicator for if the strength of signal in county c is above the median.²⁰ We also explore robustness to specifications using continuous measures of signal strength.

To account for potential endogeneity in the location of radio stations, we follow Olken (2009) and simulate the hypothetical signal quality in free space (i.e., assuming terrain is flat and absent of any geomorphological obstacles). Conditional on the “free space” signal, which captures variation in signal strength driven by proximity to transmitters, the coefficient of $Signal_c$ is identified only from the idiosyncratic variation in propagation patterns caused by topography, which is plausibly exogenous.

Although this assumption is fundamentally untestable, we provide indirect evidence of conditional independence by examining the correlation of radio coverage with local characteristics that can determine participation in violence. The correlates we consider are: total population, population density, industrialization, gender ratio, township administrative classification, linguistic fragmentation, and historical development (as measured by number of Buddhist temples constructed prior to 1920). Table 3.2 shows the relationship

²⁰Similar threshold-based designs are used in Bursztyrn and Cantoni (2016) and Durante et al. (2015).

between signal strength and these county characteristics. The first column presents the coefficients from univariate regressions in which a dummy for having above median signal in a county is regressed on each of the correlates. In the second column, we add the “free space” signal variable along with a vector of geographic controls. The geographic features are geographic coordinates, an indicator for containing rivers, an indicator for being coastal, the ruggedness of terrain, and having railroad access.

As should be expected, the unconditional distribution of radio signals is not random, as the raw correlations are statistically significant. However, with controls, $Signal_c$ is not significantly related to most observable conditions prior to the Cultural Revolution (with the notable exceptions of agrarian population and the township dummy for whether the county is classified as a *shixiaqiu*). We find that controlling for “free space” signal along with geographic characteristics alleviates much of the concern over selection. Therefore, we explicitly control for these covariates in our main specification. Figure 3.3 displays the geographic variation in actual and hypothetical signal strength.

3.3.3. Linguistic Data

Another source of identifying variation is generated by differences in regional vernacular dialects and heterogeneity in their compatibility with Standard Mandarin. We assemble this data in two steps. First, we identify the spatial variation in dialects across China from *The Language Atlas of China*. The Language Atlas is a compilation of local linguistic studies documenting Chinese dialects and their genealogical relationships. The digitized data is organized at the county level. It records the primary dialect spoken in each county, other minor dialects if present, and the dialect families they belong to.²¹

Second, to measure each dialect’s linguistic distance to Mandarin, we appeal to experimental data collected by linguists in the field. Tang and Van Heuven (2009) study

²¹Although the majority of the languages are spoken by those of Han ethnicity, in our robustness section we restrict our attention to linguistic diversity within the ethnically Han population and exclude observations, primarily in autonomous regions, where non Chinese languages, such as Turkic, Altaic, or Mongolian languages are observed. This mitigates possible confounding bias of ethnicity.

the strength of pairwise mutual intelligibility between Chinese dialects. They relate functional intelligibility between dialects to proximity in lexical structural and phonological regularity.

The authors conducted an extensive experiment in order to find the mutual intelligibility between pairs of Chinese dialects. 150 native speakers of each of 15 different Chinese dialects were subjected to a listening exam where they were asked to identify words and sentences read by speakers of another dialect, including their own. The listening exam was administered via a recording of 288 standard Chinese core words read by a native speaker of each dialect. The participants resided in rural areas and were around the age of 50 in 2009, and thus, were youths at the time of the Cultural Revolution.²² These participants were also selected because they had never traveled outside of their home province. From this experiment a measure of bilateral intelligibility between dialects was compiled. A reproduction of their findings are shown in Figure 3A.1. From this chart, we focus on the row “Beijing,” which represents the ability of listeners of each dialect to correctly identify words from the Beijing Mandarin dialect.²³ Since only 15 dialects were studied instead of the entirety of Chinese dialects, we use each of these dialects as the representative of the family of dialects that they originate from. We standardize these comprehension measures by subtracting the percent understood of one’s native dialect.

We construct our analytical dataset by combining the spatial data on the dialects spoken by county and the proximity of each of these dialects to the Mandarin family and Standard Mandarin. Figure 3.4 maps the geographical variation of the underlying data.

3.3.4. Control Data

We complement the above dependent and explanatory variables with additional sources of control data, which we briefly outline here. The socio-economic and demographics

²²Therefore, their intelligibility scores would closely reflect intelligibility of youths during the time period we are interested in.

²³The pronunciation of Standard Mandarin is based on Beijing Mandarin.

information come from the 1964 Census, which was the last census enumeration prior to the Cultural Revolution. This data is obtained from the University of Michigan’s China Data Center. We also use the China Historical GIS (CHGIS) digital map collection of Harvard University. Using the maps, we compute the proximity to nearest navigable river and distance to the coast from the centroid of each county as well as the number of historical Buddhist temples contained in each county. Data on county-level railroad access as of 1961 is created from rail network files provided by Baum-Snow et al. (2017). Ruggedness and terrain feature data are constructed as per instructions from Nunn and Puga (2012). Table 3.1 shows the summary statistics of select explanatory and dependent variables.

3.4. Empirical Strategy & Results

In this section, we study how radio propaganda affected the intensity of the Chinese Cultural Revolution contemporaneously. The outcome we study is the number of killings directly attributable to the Cultural Revolution. Our empirical strategy is motivated by the institutional features of the historical episode, where both the quality of broadcast signal and the local people’s comprehension of Mandarin affected the strength of exposure. We define treatment as the interaction of high signal strength with linguistic distance from Mandarin. Equation 3.1 describes the baseline specification:

$$(3.1) \quad y_c = \beta \text{Mandarin}_c \cdot \text{Signal}_c + \alpha \text{Signal}_c + \sigma \text{Mandarin}_c + \gamma \mathbf{X}_c + \lambda_p + \epsilon_c,$$

where Signal_c is an indicator variable for having above median radio signal in county c , Mandarin_c is the mutual intelligibility of dialect in county c with Mandarin, \mathbf{X}_c is a vector of county characteristics and λ_p is a province fixed effect. The outcome, y_c , is the casualty rate directly attributed to the events of the Cultural Revolution in county c . The coefficient of interest is β , which measures the effect of the interaction of language with radio reception.

In practice, the $Mandarin_c$ variable is constructed in two ways. First, we define an indicator variable for if the main local dialect in county c belongs to the Mandarin language group. Second, we use a continuous measure of linguistic distance, *Experimental Intelligibility_c*. This is the percentage of the 288 core Mandarin words correctly identified in a listening exam, by a sample of 150 speakers of the dialect in county c .²⁴

The continuous intelligibility measure, which utilizes more variation than a binary Mandarin indicator, may also be more exogenous than the binary indicator for Mandarin, dissuading concerns that our results are driven by fundamental differences between Mandarin and non-Mandarin counties.²⁵ Because treatment varies at the dialect level and unobservables differ along this dimension, we cluster our standard errors at the level of the local dialect.

The baseline controls, \mathbf{X}_c , include contemporaneous information such as a county's population in 1964, 1964 share of agricultural population, 1964 gender ratio, 1964 number of households, and railroad access in 1962; historical controls such as the number of Buddhist temples within the county, an index of ethnolinguistic fractionalization; as well as robust set of geographic controls, such as the ruggedness of the terrain, river and waterway access, county area, latitude and longitude, distance to provincial capitals, distance to Beijing. To account for unobservable characteristics that could vary systematically with location of radio stations, we include the free space signal strength and a set of radio station fixed effects.

Formally, the econometric framework relies on the interaction between the two sources of variation, $Signal_c$ and $Mandarin_c$, and only the interaction is to be interpreted as plausibly exogenous. The key identifying assumption is that the interaction term between

²⁴The *Experimental Intelligibility* measure is created from data collected in Tang and Van Heuven (2009), discussed in detail in Section 3.3.

²⁵In the early days of the Cultural Revolution and during the rise of the Communist Party, language standardization had nothing to do with Communism. In addition, in this period, there was no sense in which Mandarin speakers were more educated (although this may be the case today, as Mandarin is the official language taught in schools). Since language standardization was still beginning to take effect at the start of the Cultural Revolution, Mandarin comprehension was a matter of location, not of education.

$Mandarin_c$ and $Signal_c$ is orthogonal to other determinants of violence. The parameter β reflects the relative difference in the level of violence between Mandarin and non-Mandarin counties with the same radio coverage. The interpretation is similar in spirit to a traditional difference-in-differences setup, but exploits comparisons across space rather than time. Analogously, identification requires that the unobserved differences between Mandarin and non-Mandarin counties with high broadcast signal to be comparable to the unobserved differences between Mandarin and non-Mandarin counties with low broadcast signal in a counterfactual world in the absence of radios.

This assumption would be violated if radio stations were installed more densely in places where the Mandarin speaking counties were expected to be more violent. We argue this is implausible, given that this would be a peculiar criteria for selection of radio station placement on the part of Communist regime, and because radio station construction was largely governed by technological constraints.

Even though we require only the interaction between language and radio signal to be exogenous, we account for the potential endogeneity in station location by controlling for the hypothetical signal loss in the absence of geographical obstacles (as a polynomial in the distance to nearest radio station). The residualized variation in signal strength will be driven by variation in exogenous features of the terrain.

3.4.1. Baseline Results

The estimates from our main equation are presented in Table 3.3. The raw relationship without additional controls is reported in column (1). In the tables, we notate $Signal_c$ as $\mathbb{1}\{\text{Signal}\}$ and $Mandarin$ as $\mathbb{1}\{\text{Mandarin}\}$. We find a significant and positive coefficient of 0.043 on the interaction term. In column (2) we control for signal intensity under the assumption of flat terrain (without geomorphological obstacles) with a polynomial in the distance to nearest radio station, a province fixed effect and an indicator for if the given county contains a radio station. We introduce further geographic controls in column (3),

which include a county's latitude and longitude, railroad access, ruggedness, river and coastal access, area size, distance to major cities, and distance to Beijing.

We introduce controls for pre-existing socioeconomic conditions in column (4). This includes: the number of historical Buddhist temples, 1964 county population, 1964 county gender ratio, 1964 fraction of non-agricultural population, the number of households in 1964, and the ethnolinguistic fragmentation. This is our preferred specification. Adding the baseline controls improves the explanatory power of the econometric model. Whereas the coefficients on just $Signal_c$ and $Mandarin_c$ are not significant, the point estimate on the interaction term remains similar in magnitude.

In counties that speak Mandarin, an above median signal quality leads to a 0.039 percentage point increase in the percent killed per county, relative to non-Mandarin speaking counties. This is a meaningful effect, as the average percent killed per county is 0.042 percent. The standardized beta coefficient is 0.239, which implies almost a quarter standard deviation shift in the percent killed per county given a standard deviation shift in the independent variable, the interaction between $Signal_c$ and $Mandarin_c$. The onset of violence is driven by the interaction between radio signal along with the ability to comprehend the propaganda.

In last four columns of Table 3.3, we replicate the analogous results using *Experimental Intelligibility* as a continuous measure of linguistic distance to Mandarin. One might think that Mandarin speakers are somehow systematically different and have different unobservables. This analysis mitigates those concerns by allowing us to utilize linguistic variation within the Mandarin and non-Mandarin groups.

We start our analysis with no controls (column 5), then include province fixed effects and free space signal loss controls (column 6), and finally the baseline controls (columns 7 and 8). The estimates are robust to inclusion of various controls. The effect of radio signal is increasing in the linguistic distance to Mandarin. Counties with strong signal reception where the local dialect was mutually intelligible with Mandarin experienced a

greater degree of violence. Broadcast access without comprehension, or vice versa, was not sufficient in explaining Cultural Revolution conflict.

3.4.2. Robustness

To assess robustness of the results, we conduct the following checks: (i.) alternative definitions of the dependent and independent variables; (ii.) clustering at different geographic units; (iii.) using different subsamples (omitting areas with a high percentage of ethnic minorities); and (iv.) successively including additional controls that correspond to additional determinants of violence.²⁶

Alternative Dependent and Independent Variables

Table 3.4 shows that the effects of radio and language are not dependent on specific constructions of either the dependent or independent variable. First, in columns (1) - (4), we perform our baseline analysis with the number of people killed in each county as the outcome, rather than the fraction of population killed.

Then, in columns (4) - (8) we use percent persecuted instead of percent killed per county as the outcome variable (this is from the same dataset as the violence data).

Finally, in columns (9) - (12) we use the continuous predicted signal itself, rather than a binary indicator for whether the signal is above the median signal strength, as an explanatory variable and in the interaction. All the results using alternative dependent variables are qualitatively similar to the main specification. Additional signal thresholds are explored in the appendix.

Clustering

In our main specification, we clustered our standard errors at the level of dialect group. In Table 3.6, we show robustness to different specifications of clustering. To accommodate for other types of unobserved correlation between observations across other geographic boundaries, we cluster standard errors at the nearest the radio station level (columns 1 and 2), and at the radio-language group level (columns 3 and 4). We also use Conley

²⁶In the appendix, we also censor the sample by omitting observations with extreme values.

standard errors with a 150km cutoff for spatial autocorrelation (columns 5 and 6). The results remain robust.

Subsamples

One potential concern is that our estimates are driven by targeted violence towards ethnic minorities in peripheral or autonomous provinces. This would mean that our results could be confounded by the effect of ethnic based conflict. To address this, we restrict our attention to what is traditionally referred to as China “proper” and areas that are uniformly ethnically Chinese.

Contemporary China consists of 23 provinces and 5 autonomous regions. By restricting samples to only the territories within the Great Wall, or China “proper,” we demonstrate, that our results are stable across sub-populations (see Table 3.5). First, we exclude the Northwestern provinces (Gansu, Qinghai, Ningxia, Xizang (Tibet), and Xinjiang), second the Southwestern provinces (Guangxi and Yunnan), then the Northeastern provinces (Inner Mongolia, Heilongjiang, and Jilin), and finally all of these border provinces, in columns (1), (2), (3), and (4) respectively. This implies that the effect of radio broadcasts on violence is not explained by a story of the core versus periphery parts of China, or by the ethnic-based violence that may have occurred during the Cultural Revolution.

Additional Covariates

We make further progress towards addressing selection and improving the validity of our research design by controlling for omitted variables explicitly. Drawing on the growing literature on determinants of violence, we include additional controls to dissuade potential competing stories. The robustness results are shown in Table 3.7. Column (1) of Panels A and B refers to our baseline specification. In the subsequent columns we successively add more controls. In column (2) for Panel A and B, we interact each of our baseline controls with $Mandarin_c$ and $Experimental\ Intelligence_c$ respectively.

A credible threat to identification is correlation between Mandarin and other controls that affect violence in high radio penetration areas, since identification in our setting draws from the interaction of radio coverage and linguistic distance. The inclusion of

interacted controls allows the effect of any control we include to vary by the language of the county. Thereby the residual variation captured by β is not conflated with the differential impact of observed controls across language groups.

To account for unobservable differences across regions and among the radio towers, we add fixed effects for the nearest radio stations in column (3).

Lastly, our choice of additional controls is guided by the determinants of violence that have been emphasized in the literature. In column (4), we control for crop suitability of the local soil (this includes suitability of grain, wheat, rice, and millet). This is motivated by the literature on effect of crop suitability on long-run economic development, given that initial economic conditions can affect the onset of violence. Crop suitability is an index created by the Global Agro-Ecological Zones (GAEZ) model developed by the Food and Agriculture Organization.²⁷

In column (5) we include controls for historical development variables: the number of civil service entrants and imperial exam qualifiers. This is included to capture the possible effect of economic development on the extent of the violence given the purported goal of punishing former land owners.

We also include covariates on the incidences of historical conflict in column (6), given the persistence in conflict and culture that is documented in the literature. We control for conflict during the Taiping and Boxer Rebellion as well as revolutionaries in the initial Republican revolution.²⁸

The coefficient stays consistent and significant across all specifications. This alleviates concerns regarding our empirical strategy.

²⁷<http://www.fao.org/land-water/en/>

²⁸The Taiping and Boxer Rebellion variables are dummies equal to one in counties affected by the Taiping Rebellion, and where the Boxer operated and killed foreigners, respectively. The revolutionaries in the initial Republican revolution is a count variable. The data on historical development and conflict come from Bai and Jia (2016).

3.4.3. Heterogeneous Effects

We test for heterogeneity in the effect of exposure to media, by splitting the sample across several observable dimensions: the number of schools per province, Communist Party membership provincially, the Great Chinese Famine intensity at the province level, and degree of linguistic fractionalization in the county. Table 3.8 shows the results.

We first test whether radio broadcasting interacts with another source of ideology: schooling. On one hand, education may be complementary to radio broadcasting and amplify its messages. In particular, the classroom was a setting where propaganda was typically bundled with education and ideology was transmitted.²⁹ On the other hand, radio propaganda may be a substitute for curriculum-based persuasion. Schooling can also increase literacy rates and critical thinking, thus leading people to become more skeptical of propaganda (Zaller, 1992). Thus, the effect of radio broadcasts could be potentially more salient in areas with less schooling. We find evidence consistent with the latter story: in provinces where the number of primary schools was below the median, radio was more effective.

Second, we investigate whether significant complementarities exist between local party infrastructure and media propaganda. We proxy for local party organization with the density of Communist Party members in the province.³⁰ We divide the sample into provinces with above and below median density of party members. Columns (3) and (4) display the estimates. We find that the point estimate for our treatment is higher in provinces with an above median density of Communist Party members but the difference is not statistically significant. The magnitude of the effect of radio and language do not vary significantly with the extent of party organization.

Third, we examine the differential impact of radio propaganda across areas with differential intensity of the Great Famine. The Great Chinese Famine was a period in the

²⁹Cantoni et al. (2017) find that school curricula can affect students' political attitudes.

³⁰The data is from Kung and Lin (2003). The province is the smallest geographic level at which we could locate data on party membership.

People’s Republic of China between the years 1959 and 1961 was characterized by widespread famine. It immediately preceded the Chinese Cultural Revolution and the severity of famine has been attributed to Mao-era policies from the Great Leap Forward. Conceptually, we study if experience with poor government policies in the past affects adoption or compliance with government programs in the present. Previous work has found that political repression leads to increased political apathy in subsequent decades (Meng et al., 2015). However, Columns (5) and (6) report that the effect of radio was larger in provinces that experienced greater level of famine mortality (in terms of death rate). This suggests that experience with previous governmental failures did not mitigate the intensity of the Cultural Revolution. Instead, more conflict was perpetrated in areas with higher famine mortality, perhaps due to the buildup of hardship and resentment from the prior episode.

Finally, we examine if there is a differential effect of the impact of radio and speaking Mandarin along the dimension of linguistic fragmentation. For counties with no linguistic fragmentation, the effect of radio propaganda interacted with Mandarin-speaking counties on violence is lower than for counties with linguistic fragmentation. This accords with work by Esteban and Ray (2008) and Esteban et al. (2012), who discuss the effect of ethnic divisions on conflict.

3.4.4. Persuasion Magnitudes

In order to contextualize our estimates, we compute the persuasion coefficient following DellaVigna and Gentzkow (2010). It measures fraction of the population exposed to message but otherwise not predisposed to violence persuaded to kill:

$$f = \frac{\text{percentdead}_t - \text{percentdead}_c}{\text{exposure}_t - \text{exposure}_c} \frac{1}{1 - y_0} \cdot 100$$

y_0 is the counterfactual behavior in the absence of treatment. We calculate the effect of moving from non-Mandarin speaking county (control) to a Mandarin speaking county (treatment), controlling for the existence of radio. Thus, a proxy for y_0 is the average

percent dead in a non-Mandarin speaking county with radio. We define $exposure_t$ as the percent of Mandarin speakers in a Mandarin speaking county, and $exposure_c$ as the percent of Mandarin speakers in a non-Mandarin speaking county.³¹ The *percentdead* values are from our estimated coefficients. The persuasion rate is:

$$f = \frac{0.0003897 - 0.0000781}{0.25 - 0.06} \frac{1}{1 - 0.0003147} \cdot 100 = 0.164\%$$

The persuasion rate implies that radio messages induced 0.164% of individuals, not otherwise predisposed to violence, to commit killings. Though this percentage point estimate is small compared to other measures of persuasion in the literature, which span from 1% to 20%, killing is a much more extreme behavior than, for example, voting. Hence, we believe that our estimate is reasonable. In addition, given the number of people exposed, even a small realized persuasion rate is consequential in explaining the sheer level of violence, especially given that homicide rate in China is typically low.

3.5. Mechanisms

Thus far, we have provided evidence that the violence of the Cultural Revolution was higher in Mandarin-speaking areas with stronger radio signal. This result aligns with the previous literature that demonstrates the effects of media on individual behavior. An important, but relatively unexplored question that follows is the mechanism through which increased violence occurs. The existing literature has emphasized two distinct channels through which media can induce conflict: by inciting ordinary citizens directly (bottom-up dynamics), or by motivating local bureaucrats to promote more violent tactics (top-down organization). Individuals can act of their own volition upon hearing ideological messages, or local bureaucrats might escalate violence by coercing the people to act. Conceptually, the effect of media can be direct or mediated through changes in the local environments

³¹We measure these exposure values using the China Family Panel Studies dataset (described in section 5.1), where we have data on both Mandarin speakers and Mandarin counties.

brought about by media exposure. Using micro-data from the China Family Panel Studies (CFPS) survey, we provide evidence that suggests that media induced differential individual participation, even when local enforcement was presumably constant.

3.5.1. Individual Variation and the Send Down Movement

In this section, we utilize within-county variation to study one potential mechanism through which media affected individual behavior, by examining a complementary outcome pertinent to the period: participation in the Send Down Movement. In contrast to the violence, which had implications for both urban and rural areas, the Send Down Movement, also known as the Rustication Movement, was primarily an urban event. It was a program in which 17 million urban youths were “sent down” to the rural countryside to work alongside peasants in order to learn from them. Although rural rustication was framed as a necessary reeducation, it served a dual purpose to defuse violence and lower unemployment in the cities. It was organized in late 1968 in response to the escalating violence in cities with the intent to de-mobilize the Red Guards in an ideologically expedient fashion.

The movement was partly compulsory and enforced through the local government, but partly voluntary as well. At face value, it would not seem highly desirable to be sent far away from home to perform manual labor. However, Mao depicted this movement as a patriotic event, so former Red Guards and other youth volunteered to be rusticated. Red Guards were disproportionately likely to participate, possibly out of their own volition or selective targeting.³² In total, around 15% of individuals who were sent down volunteered because of a genuine belief in the revolutionary agenda (Pan, 2003).

According to an interview from a former sent-down, “It was like a group outing because we were sent down with our friends—all of our good friends and classmates from the same class at the same school so I thought it would be fun...it was just in popular political

³²Many were also forcibly sent down to the countryside.

fashion for you to leave” (Rene, 2013). We examine if voluntary participation can be attributed to ideological manipulation through state-sponsored media.

A key feature of our data which allows us to determine a mechanism through which propaganda operates is that we are able to utilize within-county variation. This allows us to control for the regional differences in enforcement and isolate the individual component of participation.³³ Our identification strategy involves a difference-in-differences framework in which the two sources of variation are living through the Cultural Revolution during one’s youth, and being a native Mandarin speaker. In this framework, any differences between the Mandarin and non-Mandarin speakers for that cohort in excess of the differences from other cohorts is attributed to the effect of propaganda.³⁴

3.5.1.1. Individual-level Data. Individual-level micro-data comes from the China Family Panel Studies (CFPS), a cross sectional and retrospective survey conducted in 2010 by the Institute of Social Science Survey of Peking University in China. It collected data relating to respondents’ educational outcomes, family dynamics, migration and health. Demographic information include date of birth, province of birth, province of residence at various points throughout one’s childhood, Send Down Movement experience, gender, ethnicity, parents’ occupation, hukou status, and language spoken at home. The CFPS dataset consists of 33,000 individual observations from selected counties across China.³⁵

The specific individual level outcomes that we study in our paper are revolutionary activity. To measure this, we examine the responses to the following question:

- “Have you had any of the following life experiences?” Choices include the “Send-Down” experience, referred to as *send down*.

³³Differences in individual response can arise due to differential sensitivity to pressure from authorities, or differential disposition toward revolutionary ideas. We do not attempt to distinguish between the two.

³⁴We are unable to perform the same specification in the previous section, which uses signal strength as an additional source of variation, because the nature of the data is coded. We cannot observe the actual identities of the counties.

³⁵After dropping observations with relevant missing data we are left with 13,350 observations

In addition to these outcome variables, we also proxy for comprehension of Mandarin during the respondent's youth with a dummy variable indicating if the primary language used in daily communication with the respondent's family is Mandarin Chinese, as opposed to a Chinese local dialect or minority ethnic dialect. This variable is referred to as *Mandarin*. Summary statistics are shown in Table 3.9.

3.5.1.2. Empirical Strategy. We use two key sources of variation to investigate the effect of propaganda on individual revolutionary behavior. The first is being a native Mandarin speaker, since this affects comprehension of the radio broadcasts. The second source of variation is individual age at the start of the Cultural Revolution. Evidence from the psychology literature suggests that adolescents and youths are more susceptible to media than those of other age groups.³⁶ The Cultural Revolution itself also suggests that adolescents might be the most influenced by propaganda during this time period. The Communist Party specifically targeted youth to join the revolutionary cause, and most Red Guards were between 12 and 17 years old (Jing, 1991).³⁷

For these reasons, we consider individuals aged 10 to 21 at the start of the Cultural Revolution, born between 1945 and 1956, as the cohorts differentially exposed to the Cultural Revolution rhetoric. This definition incorporates both the historical details as well as evidence from psychology literature outlining the most impressionable years. We explore the robustness of our empirical results to different definitions of these age categories later in the section.

³⁶The *impressionable years* hypothesis states that the historical environment to which one is exposed to during the transition between adolescence and adulthood has a profound impact on one's attitudes and world views. After this time of plasticity, beliefs become set and permanent. Young adults are especially vulnerable to shifts in attitudes in political beliefs (Alwin and Krosnick, 1991; Flanagan and Sherrod, 1998). There is no consensus for which ages exactly constitute the impressionable ages. Some see the age of 18 as the end point, but others see the process of socialization lasting until the age of 25.

³⁷The Cultural Revolution was not the first instance in history during which youth were the targets of propaganda. During the Nazi regime in Germany, youth were similarly targeted due to their naivete (Hoffmann, 1996). Giuliano and Spilimbergo (2014) examines outcomes on beliefs among those who lived through recessions during their impressionable years. Medical literature has also found strong associations between adolescents and risky behavior (Escobar-Chaves and Anderson, 2008; Klein et al., 1993).

Thus, we explore the effects of understanding Mandarin among the cohorts for whom propaganda was the most salient. The difference-in-differences regression we estimate is as follows:

$$(3.2) \quad y_{ijc} = \alpha \text{Mandarin}_i + \gamma \text{CR cohort}_j + \delta \text{CR cohort}_j \cdot \text{Mandarin}_i + \omega_c + X_i \beta + \epsilon_{ijc}$$

In the equation above, y_{ijc} denotes the participation in the Send Down Movement of individual i of cohort j born in county c . The independent variable, CR cohort_j , is a dummy for being age 10 to 21 at the start of the Cultural Revolution, and Mandarin_i is a dummy for speaking Mandarin at home. ω_c is a county fixed effect which partials out the difference in outcome arising from radio infrastructure or local political conditions. The parameter of interest is δ , the coefficient on the interaction term, which measures the effect of speaking Mandarin for the Cultural Revolution cohort.

In our preferred specification, X_i is a vector of individual level controls, which include gender, education, age, age squared, father's education, mother's education, father's political party, mother's political party, father's occupation, mother's occupation, birth county, urban area of residence dummy, father's birth year, mother's birth year, own birth year, ethnicity, and parents' hukou status. These factors may determine Send Down Movement participation and be correlated with language use at home.

For instance, we control for parental demographics and characteristics because children of parents targeted by the Communist Party for belonging to “bad class backgrounds” may have been more likely to be sent down—either of their own volition to prove their loyalty to the Communist Party—or through coercion. Additional controls include interactions between CR Cohort and education, gender, urban dummy, birth province, and ethnicity.

We note that by utilizing *within* county variation, we allow for radio presence as well as administrative quality to vary flexibly across localities. The purpose of this is to study if individual decisions to participate in Send Down Movement responded to media exposure directly, absent of the mediating effect of local environments. By controlling for

the regional differences in enforcement and radio penetration, we isolate the individual component of participation.

We also estimate a non-parametric event study model by using 5-year birth cohort dummies instead of specifying a treated cohort, because the definition of the impressionable age varies. This lets data “tell the story”. It also allow us to assess the validity of our research design by visualizing the “pre-trend” and examine if differential behavior between Mandarin and non-Mandarin speakers was unique to the Cultural Revolution cohort. We estimate the following fully flexible regression specification:

$$(3.3) \quad y_{ijc} = \alpha \text{Mandarin}_i + \sum_j \gamma_j \text{cohort}_j + \sum_j \delta_j \text{cohort}_j \cdot \text{Mandarin}_i + \omega_c + X_i \beta + \epsilon_i$$

The independent variables, cohort_j , are a set of dummies for the birth cohorts that individuals belonged to. The cohort categories are defined by the age at the start of the Cultural Revolution: 30 and older, 25 to 30, 20 to 25, 15 to 20, 15 to 10, 10 to 5, and 5 to 0 in 1966. The controls are the same as Equation 3.2, except instead of interacting the *CR cohort* with individual controls, cohort_j is interacted. The estimated vectors δ_j , the set of coefficients on the interaction term between cohort_j and Mandarin_i , reveal the effect of speaking Mandarin at home on participation in the Send Down Movement for individuals belonging to each birth cohort. If, for example, the Cultural Revolution propaganda increased participation then we would expect δ_j to be positive and significant only for the birth cohorts 10 to 5 and 15 to 10.

The first two columns of Table 3.10 display the results of estimating Equation 3.2. We find that belonging to the Cultural Revolution Cohort (those aged 10-21 at the start of the Cultural Revolution) and speaking Mandarin at home has a significant effect on participation in the Send Down Movement. The coefficients of interest are the coefficients of *CR Cohort* \times *Mandarin*. The second column includes the full set of interactions between the CR Cohort indicator and education, gender, urban dummy, birth province, and ethnicity. Belonging to the CR Cohort while speaking Mandarin at home leads to a

5.8 percentage point increase in likelihood of joining the Send Down Movement, compared to a baseline value of 1.4 percent probability of participation.³⁸ To interpret this number in the aggregate, we simulate a situation in which the interaction of *CR Cohort* and *Mandarin* is zero, in our data. We compare the aggregate participation rates predicted by our model, between the observed treatment effect and the counterfactual treatment effect. We find that in absence of radio broadcasts, holding all else fixed, a 10.4% decrease in participation in the Send Down Movement would result. In other words, 10.4% of participants were persuaded by media to participate in the Send Down Movement. This number we calculate from the data can be compared to the 15% voluntary participation rate cited in the narrative literature.

The patterns in the data can also be visualized by plotting the coefficients of the interaction terms from the flexible specification. We display the set of coefficients δ_j in relation to γ_j from Equation 3.3 in Figure 3.5. The points plotted here are the level effects for Mandarin speakers and non-Mandarin speakers joining the Send Down Movement. We demonstrate evidence for the parallel trends assumption: conditional on observables, Mandarin speakers and non-Mandarin speakers join the Send Down Movement at approximately the same rate, except for the Cultural Revolution cohort. This indicates that Mandarin speakers were no different from non-Mandarin speakers, except during the time period of the Cultural Revolution, when Mandarin-language propaganda was pervasive.³⁹ This aligns with our hypothesis that due to the historical circumstances, propaganda during the Cultural Revolution was more salient for Mandarin speakers. The point estimates and their associated standard errors are shown in Table 3A.3.

Our results highlight one mechanism for the intensity of the Cultural Revolution. By exploring within-county variation, we show evidence that that Mandarin speakers of

³⁸The effect size is large compared to the mean likelihood because the mean of the treatment variable is also small, at 0.04.

³⁹We also see a differential effect between Mandarin and non-Mandarin speakers for youth aged 6-10 as well, in addition to the designated cohort. It is plausible that younger children are also susceptible to and influenced by media.

the Cultural Revolution Cohort, who thus could better comprehend radio propaganda, complied more with state policy than their otherwise equivalent, non-Mandarin speaking neighbors. This suggests that media influenced individual revolutionary behavior even controlling for the local political atmosphere and bureaucratic enforcement.

3.5.2. Social Multipliers

The analysis above describes the effect of media on atomistic behavior of individuals, but it leaves unexplored the role of contagion or diffusion in the realization of collective action. Decision making among peers is frequently not disjointed but inherently linked. The hypothesis that social influence provides an intermediate channel for political persuasion on behavior dates back to Lazarsfeld et al. (1944) and Katz and Paul (1955).

In this section, we test for the presence of social interactions. In particular, we study whether or not there exists positive spillover effects in propaganda exposure. Evidence which would be consistent with spillover effects include increasing returns to revolutionary behavior if one's peers are also participating. This would mean that there are strategic complementarities to political action, and the effect of media may be amplified in areas where more people are able to comprehend.

One feature of our setting is that we can distinguish direct exposure through personal media consumption (proxied by individual Mandarin comprehension) from indirect exposure through interaction with peers (proxied by living in a predominantly Mandarin speaking county). This allows us to estimate the effect of the interaction of speaking Mandarin at home with residing in a primarily Mandarin speaking county, on participation in the Send Down Movement. We find a positive and significant spillover effect, which is consistent with greater local government enforcement in Mandarin speaking areas, as well as a social multiplier effect. The effect of media is amplified in areas where more people are able to comprehend. The regression specification we run, restricted to members of the CR cohort, is as follows:

$$(3.4) \quad \begin{aligned} y_{ic} = & \alpha \text{Mandarin speaker}_i + \gamma \text{Mandarin county}_c \\ & + \delta \text{Mandarin speaker}_i * \text{Mandarin county}_c + X_i \beta + \rho_p + \epsilon_{ic} \end{aligned}$$

The left hand side variable, y , is an indicator for having participated in the Send Down Movement. The coefficient δ of the interaction term between *Mandarin county* and *Mandarin speaker* shows how living in a Mandarin speaking county and speaking Mandarin at home interact to affect the outcome of interest. X_i is a vector of individual level controls, which is the same as before. Here, we include province fixed effects instead of county fixed effects.

The results are shown in column 2 of Table 3.11. The coefficient on the interaction of *Mandarin speaker* \times *Mandarin county* represents the effect of speaking Mandarin and living in a Mandarin-speaking county on participation in the Send Down Movement. The coefficient indicates a 7.2 percentage point increase in likelihood in joining the Send Down Movement if one spoke Mandarin and lived in a Mandarin speaking county, above the effect of either of the two factors alone for members of the CR Cohort.

The results provide evidence that social interactions effects are non-trivial. The spillover effect from simply living in a Mandarin speaking county is positive while the coefficient on just speaking Mandarin in a non-Mandarin county is close to null. This suggests the effect of individual radio comprehension is driven by radio as a coordination device. Common or community level beliefs affect individual responses.

The multiplier effect is consistent with the historical setting of the Send Down Movement, where volunteering was a very public way to demonstrate one's loyalty to the Communist Party. Individuals were also more likely to volunteer if their friends volunteered as well. Other examples of social multiplier in the literature include peer effects on university grades (Sacerdote, 2001), crime (Glaeser et al., 1996), and returns to education (Acemoglu and Angrist, 2000).

3.5.3. Effects on Party Membership

In this section, we move beyond contemporaneous outcomes, to examine the long-term effects of media on individual behavior. Living in a Communist regime has been found to have a persistent effect on preferences and attitudes (Alesina and Fuchs-Schündeln, 2007). Our paper provides context for this persistence by studying the effect of media exposure at a critical age juncture. We consider Communist Party membership as a proxy for ideological behavior. Joining the Communist party is a rigorous process which involves a lengthy application process and scrutiny. Membership to the Communist Party is exclusive, consisting of 5% of the population. Individuals must submit multiple applications and endure several rounds of evaluations.⁴⁰ The results here contributes to a strand of literature that tries to understand the determinants of Communist party membership.⁴¹

We measure Communist Party membership through the following question in the CFPS questionnaire:

- “Are you a member of the Communist party of China?” We refer to joining the Communist party within 45 years of birth as *Communist*.⁴²

Using the same difference-in-differences strategy as in Section 5.1, we find that Mandarin speakers who were of the impressionable age cohort during the Cultural Revolution were more likely to be members of the Communist Party. The estimates of Equation 3.2, for Communist membership outcome, are shown in the last two columns of Table 3.10.

As with before, we control for family and individual characteristics that can determine party status and are possibly correlated with language use. These include parental occupation, parental party affiliation, education, class status, birth cohorts, hukou and urban status, etc. In the second column, we also the interact the controls with the CR cohort

⁴⁰www.nytimes.com/ref/college/coll-china-politics-002.html

⁴¹See Appleton et al. (2009) for a comprehensive literature review

⁴²We exclude joining the Communist party after this age because we focus on ideological motivation from one’s youth. The motivations to join in older age is less plausibly a result of this. Only 5% of individuals join after 45.

dummy. This allows the control variables to affect party membership differently for the CR cohort.

We find that individuals who speak Mandarin and belong to the Cultural Revolution Cohort are 6.0 percentage points more likely to have joined the Communist party, compared to a baseline mean value of 7.1 percent mean probability of joining the party. The mean year of joining the party for the impressionable CR cohort is 1979. We note that 19% of those who participated in the Send Down Movement also joined the Communist party. Performing the same counterfactual as with the Send Down Movement, we simulate a situation in which the treatment effect is zero. We find that in the absence of treatment, the rate of participation in the Communist party would be 7.0% lower, among those in the Cultural Revolution Cohort. This suggests that 7.0% of party members born between 1944 and 1956 joined the Communist Party under the influence of media.

In the appendix, we use binned birth cohorts to flexibly plot out the trends for joining the Communist party for Mandarin speakers and non-Mandarin speakers over time. We display the set of coefficients δ_j in relation to γ_j from Equation 3.3 in Figure 3A.1. The Cultural Revolution Cohort is highlighted between the two vertical red lines. We find that for cohorts outside of the Cultural Revolution period, there is no significant difference in Communist party membership between Mandarin speakers and non-Mandarin speakers.⁴³ This lends credibility to the research design and suggests that the differential outcomes of Mandarin speakers belonging to the CR cohort is driven by exposure to Mandarin-language propaganda at a critical age juncture.

⁴³We also perform a falsification test, shown in Appendix Table 3A.4. Instead of using the Cultural Revolution Cohort as the treated cohort, we examine other cohorts to see if the effect of speaking Mandarin for these cohorts is significantly different from zero for these cohorts. We use the baseline specification, Equation 3.2. The birth cohorts are 0-9, 10-21 (Cultural Revolution Cohort), 22-32, 33+, and born after the Cultural Revolution. These birth cohorts are constructed to be roughly the same length of time as the Cultural Revolution Cohort.

3.6. Conclusion

The difficulty associated with projecting state influence across space is a recurring theme in the setting of developing nations. To this end, mass communication technologies are frequently invoked as instruments which potentially augment state capacity. In this paper, we study how mass media enabled civilian participation and compliance in the context of political campaigns directed by the Communist Party in China.

We utilize a previously unexplored institutional detail of a widely studied historical movement, the Chinese Cultural Revolution, to study the causal impact of state-sponsored propaganda on individual and collective behavior during and after the time period. Our identification is based on the interaction of the geographic variation in radio access and variation arising from the degree of intelligibility between the local dialects and the language used in broadcasts.

Contemporaneously, we find that state radio broadcasts had positive and significant effects on the incidences of conflict during the Cultural Revolution. Localities where both radios were readily available and Mandarin was reasonably well understood experienced a greater intensity of conflict as proxied by the number of individuals killed and total persecuted victims.

The results suggest that linguistic diversity constrained the state's ability to conduct persuasion, highlighting an important tension faced by the state between standardization of policy and effective administration. This provides a context for why linguistic standardization is often an implicit or explicit policy of the nation building strategy pursued by the state.

We then investigate the mechanisms through which media mobilized participation in conflict and revolutionary behavior. Through studying the Send Down Movement, we provide evidence that one channel through which propaganda operated was through a bottom-up process, driven by a differential response by individuals who were exposed. This is in contrast to the channel of differential coercion, or top-down enforcement on

the part of local elites or bureaucrats. The media provided a channel through which the central government could influence individual action directly instead of relying solely on the enforcement of local bureaucrats. This suggests that media can resolve agency problems inherent in the implementation of central policy.

Our unique historical setting also allows us to examine the persistence of ideology. In contrast to insights from the political communication literature, we find that radio propaganda had enduring long term consequences through the recruitment of Communist party members. Individuals who both were able to understand and were more exposed to the radio messages were significantly more likely to gain membership to the Communist party later in life. This suggest that exposure to propaganda at critical junctures in life has long-lasting implications for life-trajectories and choices.

Our paper sheds light on the role of propaganda in building state capacity in both the short and long-run. In the short-term, authoritarian regimes utilize media for the successful implementation of mass campaigns that would be otherwise difficult to coordinate. In the long-run, media cultivates permanent support for the regime through the supply of future party members. The ability of propaganda to influence individual behavior contemporaneously as well as through time is central to the administrative capacity of authoritarian governments.

Tables and Figures

Tables and Figures for Chapter 1

Table 1.1. Demographics of Movers and Non-Movers

Years in Neighborhood	≤ 2 years	2-5 years	5-10 years	≥ 10 years	non-mover
age	46.1	50.96	60.0	54.5	63.6
income	56.915	67.862	75.271	78.246	69.352
percent with college degree	35.60	39.73	41.81	42.88	64.37

Note: This table shows summary statistics of the demographics of movers and non-movers, broken down by years in current neighborhood. Income is measured in thousands.

Table 1.2. Individual Level Regressions, Selected Regressors

	BMI
Age	-.0243*** (.0018969)
Female	-.999*** (.0405174)
Household Income	-3.06e-06*** (2.69e-07)
Observations	66698
R-squared	0.0662
Standard errors in parentheses	
*** p<0.01, ** p<0.05, * p<0.1	

Note: This table shows the estimated coefficients of selected regressors from a regression of individual BMI on individual covariates for non-movers. The individual covariates are age, education, household income, race, gender, marital status, employment status, size of family, country of birth and survey-year fixed effects.

Table 1.3. Median BMI in High and Low Residualized BMI Zip Codes

	5th percentile	95th percentile
<1 year since move	24.09	26.39
non-mover	24.23	27.92

Note: This table shows the median BMI for recent movers and non-movers in areas with the lowest 5th percentile and highest 5th percentile of residualized BMI. Residualized BMI is calculated as the median zip code level residual of a regression of individual BMI on individual covariates. Individual covariates are age, education, household income, race, gender, marital status, employment status, size of family, country of birth and survey-year fixed effects.

Table 1.4. Demographics of Movers to High Residualized BMI Zip Codes

	≤ 2 years	2-5 years	5-10 years	≥ 10 years	non-mover
Age	40.7	44.1	49.1	53.7	62.7
Income	42.225	51.335	61.496	61.161	55.413
Percent with college degree	19.38	22.97	26.04	25.38	24.24

Note: This table shows the demographics of movers to high residualized BMI zip codes, broken down by number of years in the neighborhood. Income is measured in thousands.

Table 1.5. Median BMI in High and Low BMI Zip Codes

	5th percentile	95th percentile
<1 year since move	23.31	27.17
non-mover	24.09	27.85

Note: This table shows the median BMI for recent movers and non-movers in the zip codes with the lowest 5th percentile and highest 5th percentile of BMI, calculated as the median BMI in each zip code.

Table 1.6. Treatment Effects

	(1)	(2)	(3)	(4)
BMI	Base	Parametric	10-year cutoff	Restricted
2-5 years	0.032 (0.150)		-0.082 (0.171)	-0.052 (0.148)
5-10 years	0.120 (0.149)			-0.014 (0.154)
10-15 years	-0.134 (0.213)			-0.189 (0.217)
log(year)		0.029 (0.080)		
5-7 years			0.033 (0.268)	
7-10 years			-0.092 (0.285)	
Non-Mover				0.656 (0.189)
Observations	5493	5493	4420	8187
R-squared	0.0810	0.0810	0.1032	0.1016

Note: This table shows the results from a regression of BMI on time in neighborhood. Columns (1)-(3) show the the analogous estimates from Figure 1.10a. Column (4) shows the estimates from a regression on both non-movers and movers, where the coefficients on the individual level covariates are restricted to be the same. Individual level controls are age, education, household income, race, gender, marital status, employment status, size of family, country of birth and survey-year fixed effects. *** p<0.01, ** p<0.05, * p<0.1

Table 1.7. Explaining Zip Code Residuals

	(1) Residual
walk score	-0.0381** (0.0185)
transit score	0.0228 (0.0158)
bike score	-0.00584 (0.00807)
Observations	119
R-squared	0.142

Robust standard errors in parentheses

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


Note: This table shows the estimated coefficients from a regression of zip code level residuals (calculated as the median zip code level residual from a regression of individual BMI on individual covariates for non-movers) on neighborhood physical characteristics: walk score, transit score, and bike score. These scores are on a scale of 0 to 100. Residuals have a median value of -0.69, with a standard deviation of 0.8.

Table 1.8. Quantile Regressions

Quantiles	10	20	50	80	90
2-5 years	-0.08 (0.606)	-0.129 (0.181)	0.148 (0.355)	0.048 (0.309)	0.511 (0.368)
5-10 years	0.25 (0.152)	0.049 (0.290)	0.37 (0.232)	-0.062 (0.409)	0.350 (0.518)
10-15 years	0.072 (0.791)	0.061 (0.229)	-0.220 (0.454)	-0.417 (0.486)	0.350 (0.540)


Note: This table shows the results of quantile regressions, regressing mover individual BMI on individual characteristics including age, education, household income, race, gender, marital status, employment status, size of family, country of birth and year fixed effects, as well as time in neighborhood categories (specified as 2-5 years, 5-10 years, and 10-15 years). The time in neighborhood coefficients presented show how the BMI of movers converges to the BMI of non-movers. Each column is a different regression conducted at the quantile specified. Robust standard errors are in parenthesis. There are 5493 total observations.

Figure 1.1. Histogram of Years Living in Current Neighborhood




Note: This figure shows a histogram of the number of years since last move.

Figure 1.2. Individual BMI Distribution




Note: This figure shows a histogram of individual BMI.

Figure 1.3. Zip Code BMI Distribution



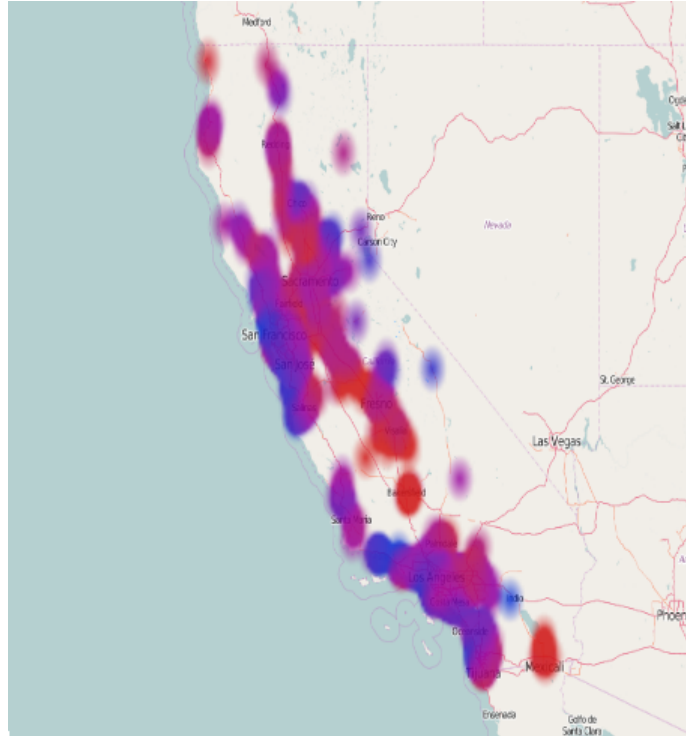
Note: This figure shows a histogram of zip code level BMI, calculated as the median BMI of all residents in a particular zip code.

Figure 1.4. Zip Code BMI Distribution for non-movers



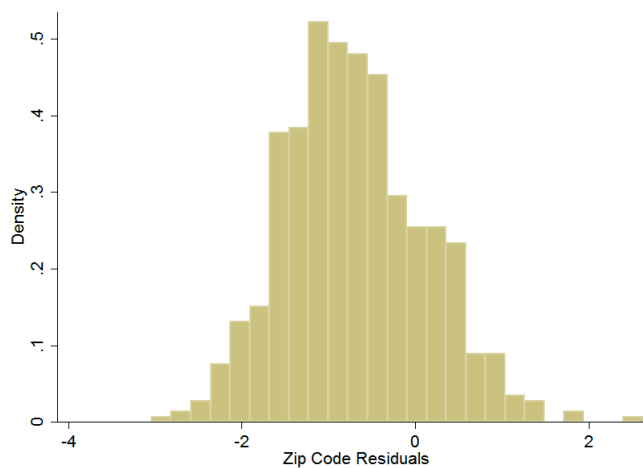
Note: This figure shows a histogram of zip code level BMI, calculated as the median BMI of all non-movers in a particular zip code.

Figure 1.5. California BMI Distribution



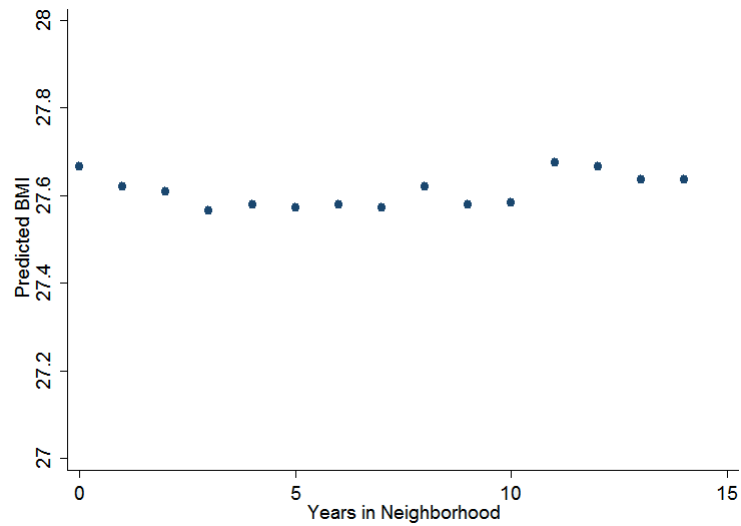
Note: This figure shows the geographical distribution of BMI across zip codes in California. Blue shades indicate low BMI while red shades indicate high BMI. Zip code BMI is calculated as the median individual BMI of residents living in a particular zip code.

Figure 1.6. Zip Code Residual BMI Histogram




Note: This figure shows a histogram of residualized BMI by zip code. Residualized BMI is calculated as the median residual aggregated on a zip code level of a regression of individual BMI on individual covariates for non-movers. Individual covariates are age, education, household income, race, gender, marital status, employment status, size of family, country of birth and survey-year fixed effects.

Figure 1.8. Covariate Stability




Note: This figure shows the predicted BMI from a regression of BMI on time invariant control covariates, for movers to high residualized BMI areas. The time invariant control covariates are education, race, gender, marital status, employment status, size of family, country of birth and survey-year fixed effects (all covariates excluding age and income). Since the plotted points are approximately level, there is evidence of covariate balance among movers.

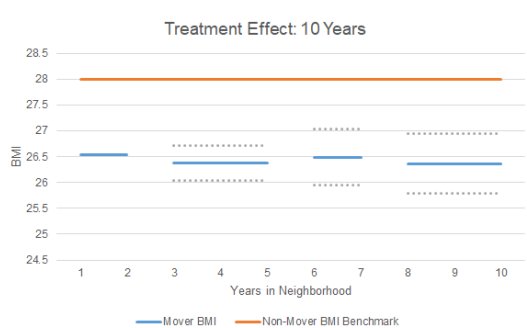
Figure 1.9. Effect of Location on BMI, 95th Percentile



(a) Effect of Location on BMI, 95th Percentile



(b) Non Parametric Treatment Effect



(c) 10-Year Cutoff

Note: These figures show graphical results from regressions of individual BMI on individual observables and time in neighborhood for movers to high residualized BMI zip codes. Individual level controls are age, education, household income, race, gender, marital status, employment status, size of family, country of birth and survey-year fixed effects. Movers are defined as those who have moved within the last 15 years. In (a), the base specification is shown. The time variable is parameterized as 2-5 years, 5-10 years, and 10-15 years. The solid straight line is the median BMI in high residualized BMI areas, while the broken solid line shows the point estimates for the coefficients of each bin of time category. These estimates are in reference to the solid line at 0-2 years, which is the median BMI for movers into high residualized BMI areas who moved within the last 2 years. The dotted lines are the 95% confidence intervals. In (b), the time variable is parameterized as the log of time in neighborhood. In (c), moving redefined as moving within the last 10 years. The time variable is parameterized as 2-5 years, 5-7 years, and 7-10 years. All standard errors are clustered at the zip code level. The table version of these figures is in Table 1.6.

Tables and Figures for Chapter 2

Table 2.1. Market Share by Revenue

Brand	Revenue Share	Average Price Per Ounce
Smart Balance	0.72%	\$0.21
Jolly Time	3.96%	\$0.12
Pop Weaver	6.61%	\$0.07
Pop Secret	12.60%	\$0.16
other	13.44%	\$0.11
lowfat	17.28%	\$0.21
Act II	18.24%	\$0.11
Orville Redenbacher's	27.13%	\$0.17

Table 2.2. Popcorn Categories

<i>Never TF</i>	<i>Reformulators</i>	<i>Non Reformulators</i>
Smart Balance	ACT II	Pop Secret
Lowfat varieties	Orville Redenbacher's	Jolly Time
		Pop Weaver
		Other

Note: See definitions of categories in text.

Table 2.3. Summary Statistics

Variable	Mean	Std. Dev.	N
Popcorn			
price per ounce	0.142	0.112	253102
average ounces per purchase	25.24	25.96	253102
total ounces by household year	108	141	59230
unique households			19829

Note: The price per ounce and average weight per purchase variables are calculated from a sample where each observation is a purchase. The variable “total ounces by household year” is calculated from a dataset where each observation is a household - year.

Table 2.4. Demographics: Popcorn

Variable	Mean	Std. Dev.	Mean	Std. Dev.
Reformulators		Before		After
Household Income	56,389	29507	56,366	29516
Some College	0.650	0.477	0.645	0.479
Child	0.266	0.442	0.267	0.442
Non Reformulators				
Household Income	54,229	29320	53,995	29338
Some College	0.632	0.482	0.631	0.483
Child	0.3	0.458	0.298	0.457

Note: Household Income is measured in dollars per year. Some College is an indicator for if the mean education attainment of the heads of household is least some college. Child is an indicator equal to 1 if a child under the age of 18 is present.

Table 2.5. Logit Parameter Estimates

	(1)	(2)	(3)	(4)
	Logit	Logit with Controls	Random Coefficients Logit	IV
price	4.353*** (0.188)	5.554*** (0.206)	–	–
$TF \cdot label$	0.001 (0.027)	0.536*** (0.037)	–	–
TF	0.037 (0.027)	0.036 (0.027)	–	–
price · income	–	0.090*** (0.003)	0.260*** (0.008)	0.174 [0.158, 0.236]
price · children	–	1.495*** (0.041)	0.111 (0.116)	0.651 [-0.255, 0.773]
price · educ	–	0.917*** (0.019)	0.917*** (0.048)	0.251 [0.009, 0.533]
TF · label · income	–	-0.023*** (0.001)	-0.015*** (0.002)	-0.011 [-0.023, -0.007]
TF · label · educ	–	-0.047*** (0.006)	-0.005 (0.011)	-0.002 [-0.047, 0.022]
TF · label · children	–	0.373*** (0.013)	0.212*** (0.024)	0.200 [0.116, 0.313]
Random Coefficients				
$TF \cdot label$ (normal)	–	–	-0.109*** (0.054)	-0.254 [-0.510, -0.122]
TF (normal)	–	–	-0.016 (0.028)	0.042 [-0.111, 0.213]
avgprice (lognormal)	–	–	3.064*** (0.010)	3.63 [3.381, 3.847]
Standard Deviations				
$TF \cdot label$ (normal)	–	–	1.19*** (0.009)	1.15 [1.393, 1.668]
TF (normal)	–	–	1.19*** (0.009)	1.09 [1.302, 1.478]
avgprice (lognormal)	–	–	0.398*** (0.004)	0.210 [0.202, 0.337]
WTP for $TF \cdot label$	\$.00026	-\$.0033	-\$.028	-\$0.013
WTP for TF	\$.0084	\$.0084	\$0.0012	\$.002
Implied Own Price Elasticities	0.27-0.96	0.27-0.97	0.7-2.5	1.7-5.0
Observations	8660620	8660620	8660620	8657853

Note: This table displays the full model estimation results for the random coefficients logit model in Equation 2.2. Each regression also includes a year dummy for the inside options, year trend on the never TF group, and popcorn brand fixed effects. For the random coefficients specifications in (3) and (4), the price coefficient is distributed $-\text{lognormal}(\mu_a, \sigma_a)$ where μ_a and σ_a is the estimated coefficients for the mean and standard deviation, respectively. In the IV specification, the residual and squared residuals from the first stage are added as explanatory variables in the second stage. In the first three columns, standard errors are in parenthesis. In the last column, bootstrapped 95% confidence intervals from 300 draws are in square brackets. Significance level stars are not shown for the bootstrapped estimates. There is a different number of observations due to the instrument, because there is not a fixed number of products in all markets (Smart Balance does not appear in all markets). In the bottom panel, willingness to pay estimates for average demographics are calculated. The first row shows the parameter of interest, willingness to pay for the label, $\frac{\delta_i}{\alpha_i}$, for a consumer with average demographic characteristics. The convention is that a negative sign means that the explanatory variable is a “bad” trait, and that the individual must be compensated for the increase in the “bad trait.” A positive sign means that the consumer is willing to pay for an increase in the desired trait. The second row shows the willingness to pay for taste, $\frac{\gamma_i}{\alpha_i}$, and the third row shows the implied own price elasticities. These elasticities are simulated in columns (3) and (4). * $p < .10$, ** $p < .05$, *** $p < .01$

Table 2.6. Heterogeneity

Criteria	WTP for TF x Label	WTP for TF
Above Median Income (\$47,000)	-0.015	.003
Below Median Income	-0.010	-0.0006
Above Median Education (some college)	-0.015	0.0026
Below Median Education	-0.013	0.0013
Above 75th pctl Frequency of Purchase	-0.030	0.020
Below 75th pctl	-0.012	-0.0034
Smokers	-0.016	0.0023

Note: The table reports the heterogeneity results from the demand model described in equation 2.2 where the sample is restricted by the criteria described in the first column. The second column shows the parameter of interest, willingness to pay for the label, or $\frac{\delta_i}{\alpha_i}$, for a consumer with average demographic characteristics. The willingness to pay for taste, $\frac{\gamma_i}{\alpha_i}$, is shown in the third column. The convention is that a negative sign means that the explanatory variable is a “bad” trait, and that the individual must be compensated for the increase in the “bad trait.” A positive sign means that the consumer is willing to pay for an increase in the desired trait.

Table 2.7. Robustness

Robustness	WTP for Label	WTP for Taste
Omit one month pre and post	-0.011	-0.0015
Outside time trend and never TF time trend	-0.014	0.002
No never TF time trend; outside dummy	-0.033	0.019
Shares Logit	-0.0005	0.005
Cookies	-0.08	-0.0019

Note: The table reports the robustness results from the demand model described in Equation 2.2 where the criteria is described in the first column. The second column shows the parameter of interest, willingness to pay for the label, or $\frac{\delta_i}{\alpha_i}$, for a consumer with average demographic characteristics. The third column shows the willingness to pay for taste, or $\frac{\gamma_i}{\alpha_i}$. The convention is that a negative sign means that the explanatory variable is a “bad” trait, and that the individual must be compensated for the increase in the “bad trait.” A positive sign means that the consumer is willing to pay for an increase in the desired trait. The estimate for cookies is a random coefficients specification without instrumental variables.

Table 2.8. Welfare Gains from Different Regimes

	Label	Ban
Prefs as Given	\$0.60	\$0.30
Benchmark Prefs	\$9	\$21


Note: This is a summary table showing the welfare gains per year for the average popcorn buyer, under different regimes compared to a no-regulation regime. This assumes one purchase occasion per month (as in the model), of an average of 25 ounces each. The first row shows the welfare gains taking preferences as given, and the second row shows the welfare gains using the benchmark willingness to pay for the label.

Figure 2.1. Nutrition Fact Label

Nutrition Facts	
Serving Size 1 Tbsp (14g)	
Servings Per Container about 32	
Amount Per Serving	
Calories 100	Calories from Fat 100
% Daily Value*	
Total Fat 11g	17%
Saturated Fat 2.5g	12%
Trans Fat 2.5g	
Cholesterol 0mg	0%
Sodium 105mg	4%
Total Carbohydrate 0g	0%
Dietary Fiber 0g	0%
Sugars 0g	


Note: This figure shows how the trans fat label appears on the nutrition facts panel.
<https://www.fda.gov/ForConsumers/ConsumerUpdates/ucm372915.htm>

Figure 2.2. Trans Fat Content in Different Product Groups




Note: This is a figure from Otite et al. (2013) showing the average trans fat content in different product groups in the years following the 2006 trans fat labeling legislation.

Figure 2.3. Demand for Trans Fat Popcorn




(a) No Brand FE



(b) Brand FE

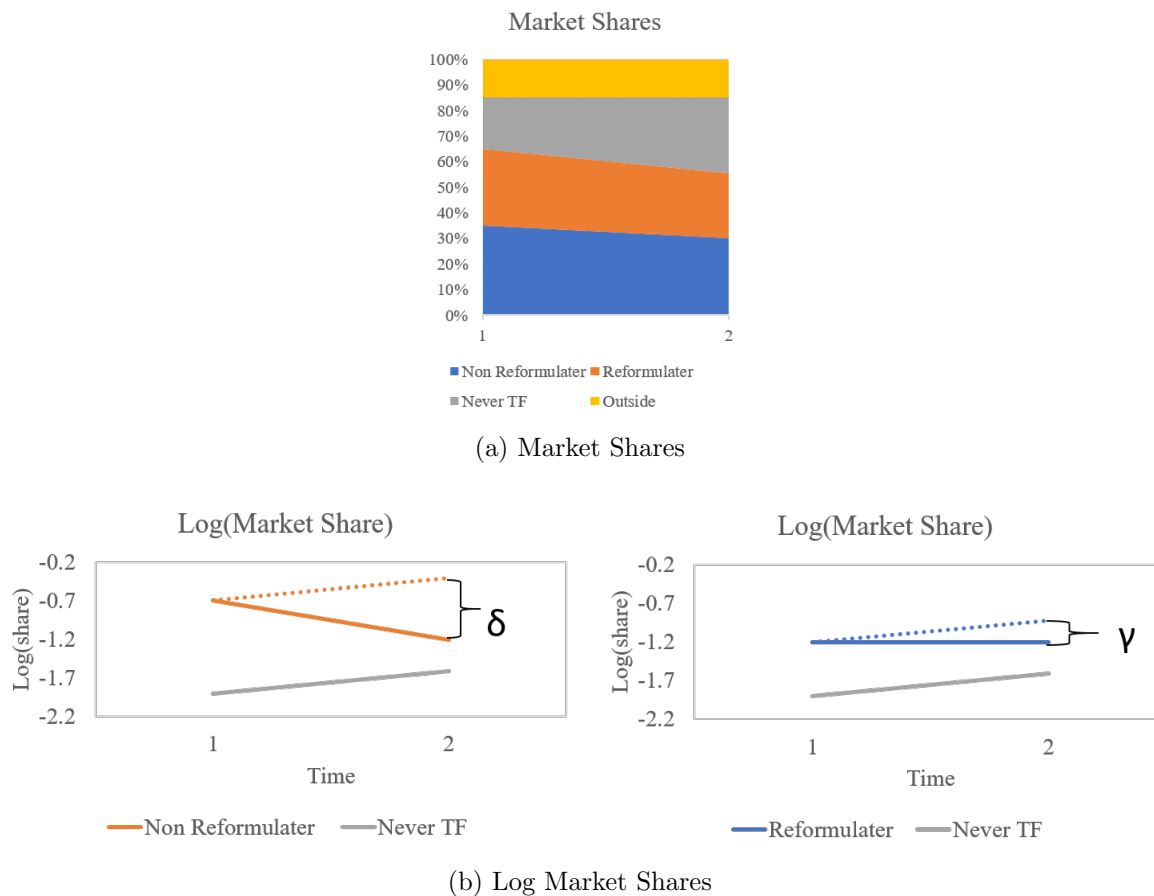
Note: This figure plots the per-household regression-adjusted quarterly demand for popcorn with trans fat, controlling for price per ounce and quarter FE. The top panel does not include brand fixed effects while the bottom panel does. The outcome variable is total weight demanded in a region-quarter-brand in ounces, divided by total households in the panel for a per-household average. The outcome is normalized to 0 in time -1. The 90% confidence intervals are in dotted lines. The horizontal red line in the post period represents the difference-in-differences estimate, which is $-2.51(0.558)$ in the top graph and $-0.615(0.176)$ in the bottom graph. The horizontal red line in the pre-period is at zero for reference.

Figure 2.4. Prices of Trans Fat Popcorn



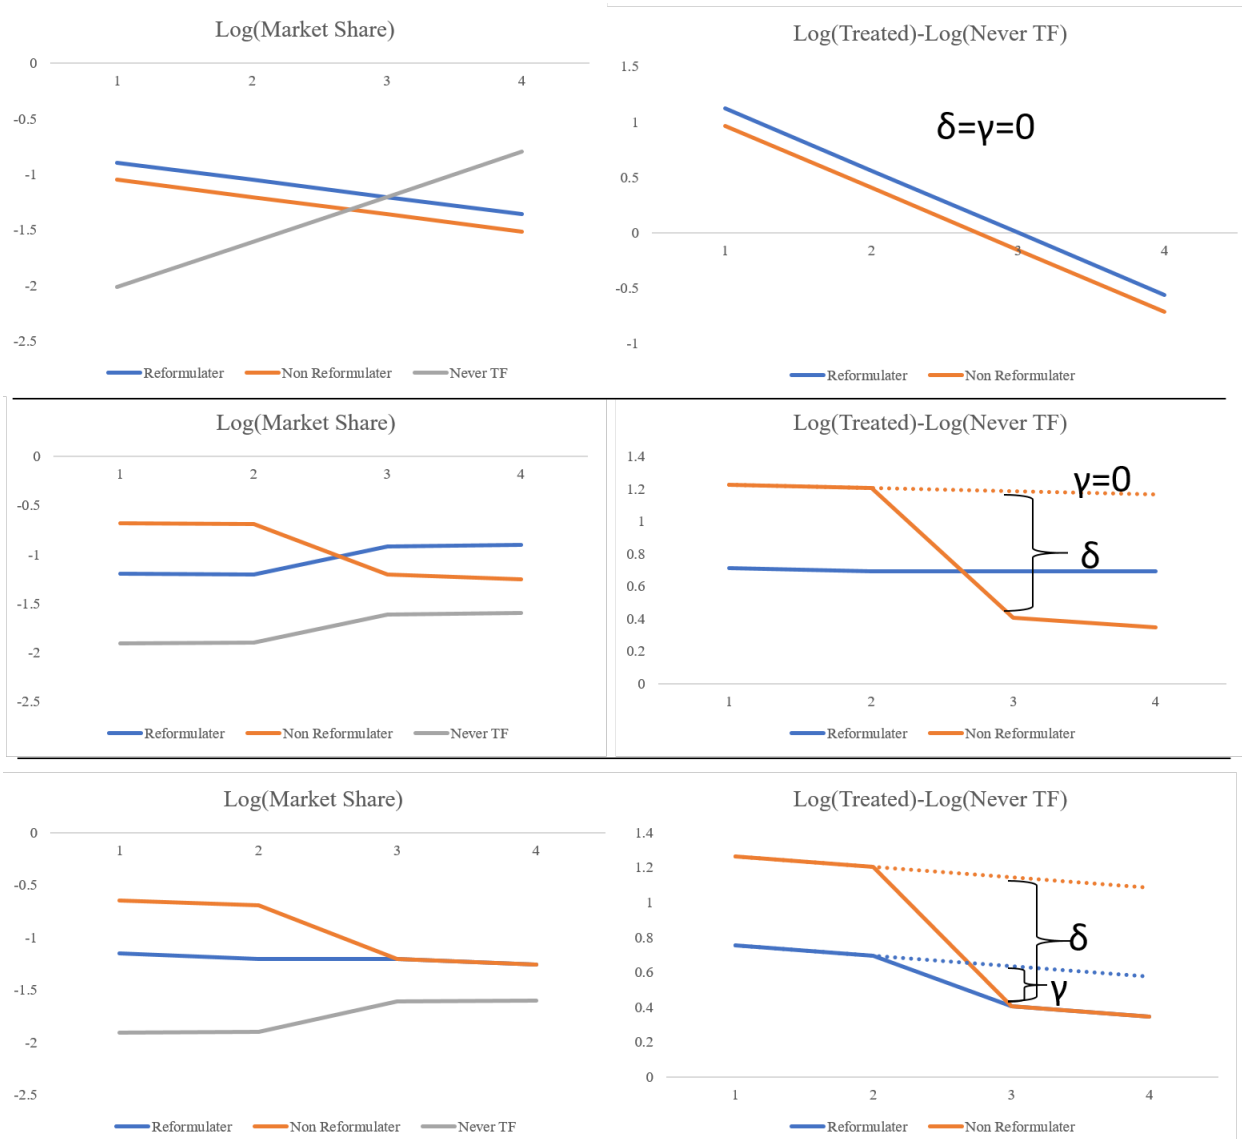
Note: This figure plots the regression-adjusted quarterly price per ounce of popcorn with trans fat, controlling for brand FE and quarter FE. The left hand side is price measured in dollars. The outcome is normalized to 0 in time -1. The 90% confidence intervals are in dotted lines.

Figure 2.5. Simple Example: 2 Periods



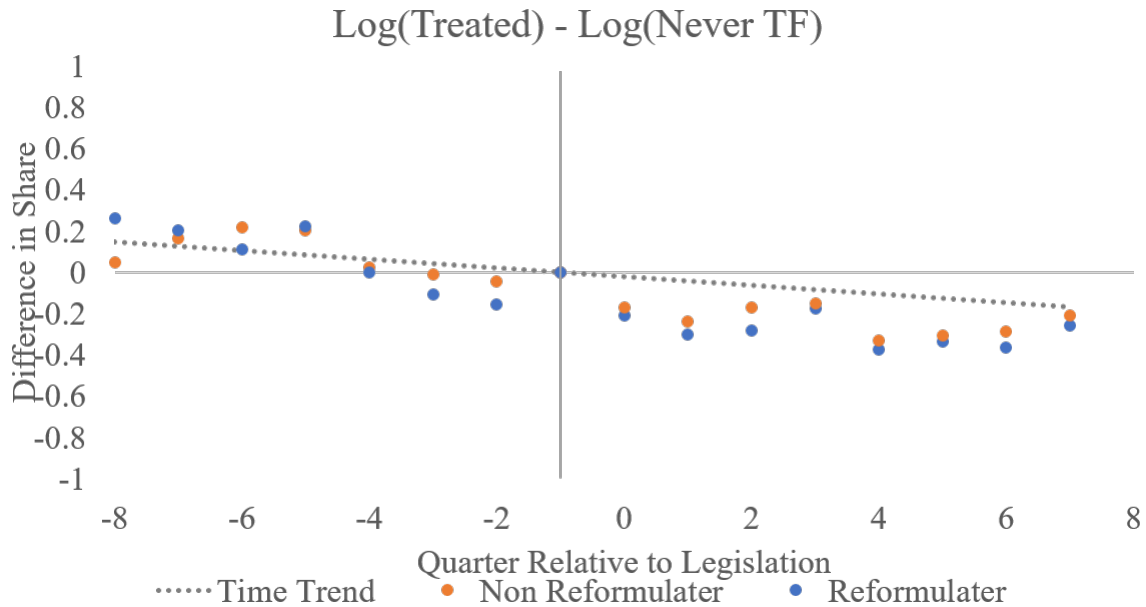
Note: These figures show a two period example to demonstrate how γ and δ are identified in a two period model. In these figures, the legislation occurs in period two. On the top, the market share levels are shown. On the two graphs on the bottom, the corresponding log of shares are plotted, for the non reformulators and reformulators respectively. The dotted lines represent the counterfactual demand for the reformulators and non reformulators, and the estimates for γ and δ are based on the difference between the counterfactual demand and the actual demand.

Figure 2.6. Simple Example: Many Periods



Note: These figures show a simple 4 period example to demonstrate the necessity of including time trends in the model. In these figures, the legislation occurs in period 3. The two graphs on the top show the log shares (left) and difference in log shares (right) between the treated group and the never TF group in a counterfactual where there are no deviations from log trends. In this case, the estimated γ and δ are zero. I note that the change in level shares is non linear. In the center, the two graphs show the log shares and difference in log shares when $\gamma = 0$ but $\delta \neq 0$. On the bottom, a deviation in log trends exists in both reformulators and non reformulators, which is where the identification for the parameters of interest comes from.

Figure 2.7. Demand for Trans Fat Popcorn



Note: This figure plots the regression-adjusted estimates of the difference in log shares between reformulators and the never TF group, and non reformulators and the never TF group, from the data. The difference in log shares for both groups is normalized to 0 in time -1. Evidence of two assumptions are shown here. First, the trends between reformulators and non reformulators relative to the control group are very similar in the pre-period, showing evidence that there are no time varying unobservables correlated with the reformulation decision. Second, the pre-trends for both groups are approximately linear, motivating the use of the linear time trend ρ_1 . The dotted line shows the estimated time trend from the data. In the pre-period, the estimated trend approximates the data, and in the post period, it shows the counterfactual trends in the absence of legislation. The γ and δ estimates are the deviation from this time trend in the post period. I note that this is not the exact same way that the full logit specification is estimated because this graph shows aggregate trends, which is for visual purposes only. Aggregated data are easier to visualize, but masks individual heterogeneity.

Tables and Figures for Chapter 3

Table 3.1. Summary Statistics

Variable	Obs	Mean	Std. Dev.
Percent Killed per County	1985	.048	.164
Signal	1985	-.455	.979
Mandarin	1985	.666	.472
Signal Free	1985	3.548	.981
Distance from Beijing	1985	1201.898	649.694
Altitude	1985	.345	.428
Area	1985	3566.197	8725.911
Population	1985	310695.8	287279.8
River Dummy	1985	.455	.498
Coast Dummy	1985	.12	.326
Ruggedness	1985	11.398	1.292
Distance to Closest Provincial Capital	1985	120.973	100.877
Railroad Access	1985	.325	.469
Non Agricultural Population	1985	.133	.248
Male Female Ratio	1985	1.071	.116
Households	1985	70587.17	64245.61
Township Dummy	1985	.101	.301
Linguistic Fractionalization	1985	.038	.139
Buddhist Temples	1985	1.045	2.361

Note: All distance measurements are in km. Area is in square km. Mandarin is an indicator for if the main dialect of a particular county belongs to the Mandarin family of dialects.

Table 3.2. Correlates of Signal Strength

	(1) Signal Indicator	(2) Signal Indicator with Geographic Controls
Mandarin	0.043** (0.021)	0.029 (0.028)
Ln(Population)	0.573*** (0.039)	0.010 (0.046)
Male/Female Ratio	-0.008 (0.005)	0.001 (0.007)
Non-Agricultural Population	0.052*** (0.011)	-0.079*** (0.015)
Ln(households)	0.569*** (0.039)	0.032 (0.046)
Township	0.039*** (0.013)	-0.122*** (0.018)
Linguistic Fractionalization	-0.006 (0.006)	0.002 (0.009)
Buddhist Temples	0.150*** (0.027)	-0.003 (0.037)
Observations	1985	1985

Note: The first column reports the univariate correlation coefficient between $Signal_c$ and the correlates on the left hand side of the table. The second column reports the coefficient on $Signal_c$ in the regression of the correlated variable on radio signal and geographical controls. Each row is a separate regression. The geographic controls are free space, altitude, geographic coordinates, an indicator for river, an indicator for coast, ruggedness, and railroad access. The correlates are at the county level. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 3.3. Media and Intensity of Conflict: Mandarin Indicator

	Dep Var: Percent of Population Killed							
	Mandarin Indicator				Exp. Intelligibility			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}\{\text{Signal}\} \times \mathbb{1}\{\text{Mandarin}\}$	0.045** (0.018) [0.275]	0.045** (0.019) [0.279]	0.049** (0.020) [0.303]	0.039** (0.018) [0.224]	-	-	-	-
$\mathbb{1}\{\text{Signal}\} \times \text{Experimental Intelligibility}$	-	-	-	-	0.187*** (0.065) [0.159]	0.169** (0.075) [0.144]	0.181** (0.079) [0.155]	0.149** (0.071) [0.125]
$\mathbb{1}\{\text{Signal}\}$	-0.058*** (0.018)	-0.030* (0.017)	-0.030 (0.019)	-0.030 (0.019)	-0.166*** (0.051)	-0.124** (0.056)	-0.132** (0.061)	-0.110* (0.056)
$\mathbb{1}\{\text{Mandarin}\}$	-0.041** (0.017)	-0.027 (0.020)	-0.033 (0.021)	-0.006 (0.020)	-	-	-	-
Experimental Intelligibility (% of Mandarin Words Comprehended)	-	-	-	-	-0.139** (0.061)	-0.155* (0.087)	-0.156* (0.088)	-0.096 (0.089)
<i>Control Variables:</i>								
Free Space Signal Propagation	-	X	X	X	-	X	X	X
Geographic Controls	-	-	X	X	-	-	X	X
Socioeconomic & Demographic Controls	-	-	-	X	-	-	-	X
Province FE	-	X	X	X	-	X	X	X
Observations	1985	1985	1985	1985	1969	1969	1969	1969
Clusters	102	102	102	102	95	95	95	95
R ²	0.015	0.194	0.201	0.228	0.017	0.194	0.200	0.233

Note: This table reports OLS estimates of the effect of radio exposure on violence. A unit of observation is a county in the 1964 census. The dependent variable in each column is the percent of the population killed in each county due to the Cultural Revolution. $\mathbb{1}\{\text{Signal}\}$ represents Signal_c , which is a binary indicator for if the simulated quality of radio signal in 1964 under real conditions is above the median signal strength. *Mandarin* is an indicator that is one if the primary dialects spoken in a county belongs to the Mandarin family. *Experimental Intelligibility* is a percentage indicating the percent of Mandarin words correctly identified by the sample of 160 respondents for each dialect group from the linguistic experiment conducted by Tang and Van Heuven (2009). The coefficient of interest is on the interaction of signal and linguistic distance each column. The Free Space Signal Propagation controls refer to radio signal in the absence of geomorphological obstacles. The full set of controls in the preferred specification, column (4), include railroad access, ruggedness of the terrain, river and coastal access, area, treaty port status, distance to major cities, distance to Beijing (geographic controls), historical Buddhist temples, 1964 population, 1964 gender ratio, 1964 fraction of non agricultural population, number of households, and ethnolinguistic fragmentation (socioeconomic & demographic controls). Standard errors clustered at the dialect group level are in parentheses. The beta coefficient in squared brackets reports by how many standard deviations the dependent variable changes due to a one deviation increase in the explanatory variable. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 3.4. Alternative Dependent and Independent Variable Definitions

Dep Var:	Number of Deaths (persons)				Percent of Population Persecuted				Percent of Population Killed			
	Mandarin Indicator		Exp. Intelligibility		Mandarin Indicator		Exp. Intelligibility		Mandarin Indicator		Exp. Intelligibility	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\mathbb{1}\{\text{Signal}\} \times \mathbb{1}\{\text{Mandarin}\}$	97.046** (40.323)	88.906** (37.774)	-	-	0.041** (0.021)	0.038* (0.022)	-	-	-	-	-	-
$\mathbb{1}\{\text{Signal}\} \times \text{Experimental Intelligibility}$	-	-	420.897*** (155.481)	393.700*** (149.532)	-	-	0.292* (0.164)	0.263 (0.164)	-	-	-	-
Signal continuous $\times \mathbb{1}\{\text{Mandarin}\}$	-	-	-	-	-	-	-	-	0.024** (0.011)	0.025** (0.011)	-	-
Signal continuous $\times \mathbb{1}\{\text{Experimental Intelligibility}\}$	-	-	-	-	-	-	-	-	-	-	0.072** (0.042)	0.050* (0.037)
Free Space Signal Propagation	X	X	X	X	X	X	X	X	X	X	X	X
Geographic Controls	X	X	X	X	X	X	X	X	X	X	X	X
Socioeconomic & Demographic Controls	-	X	-	X	-	X	-	X	-	X	-	X
Observations	1985	1969	1985	1969	1985	1969	1985	1969	1985	1969	1985	1969
Clusters	102	102	102	102	102	102	102	102	95	95	95	95
R ²	0.015	0.194	0.199	0.223	0.177	0.190	0.193	0.210	0.017	0.195	0.199	0.222

Note: This table reports OLS estimates of the effect of radio exposure on violence measured with alternative outcome and explanatory variables. A unit of observation remains a county in the 1964 census. In columns (1)-(4) the outcome variable is the raw number of people killed in a county rather than fraction of population killed. In columns (5)-(8) the outcome variable is the percent persecuted in a county due to the Cultural Revolution. In columns (9) - (12) the explanatory variable is *Signal continuous*, the actual predicted signal strength in a county under real conditions instead of an indicator variable. The Free Space Signal Propagation controls refer to radio signal in the absence of geomorphological obstacles. Geographic controls are railroad access, ruggedness of the terrain, river and coastal access, area, treaty port status, distance to major cities, distance to Beijing, and socioeconomic & demographic controls refer to historical Buddhist temples, 1964 population, 1964 gender ratio, 1964 fraction of non agricultural population, number of households, and ethnolinguistic fragmentation. All specifications include a province fixed effect. Standard errors are clustered at the dialect level. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 3.5. Restricted Samples

Sample Excludes:	Dep Var: Percent of Population Killed							
	Northwest provinces		Southwest provinces		Northeast provinces		All border provinces	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}\{\text{Signal}\} \times \mathbb{1}\{\text{Mandarin}\}$	0.037** (0.018)	–	0.038* (0.019)	–	0.042** (0.018)	–	0.051** (0.025)	–
$\mathbb{1}\{\text{Signal}\} \times \text{Experimental Intelligibility}$	–	0.168** (0.077)	–	0.117** (0.051)	–	0.167** (0.078)	–	0.253** (0.123)
Observations	1792	1776	1627	1615	1832	1816	1330	1314
Clusters	86	86	91	91	92	92	88	88
R^2	0.015	0.194	0.199	0.223	0.017	0.195	0.199	0.222

Note: This table reports results on restricted samples. Columns (1) and (2) exclude Gansu, Qinghai, Ningxia, Xizang (Tibet), and Xinjiang. Columns (3) and (4) exclude Guangxi and Yunnan. Columns (5) and (6) exclude Inner Mongolia, Heilongjiang, and Jilin. Columns (7) and (8) exclude all aforementioned provinces. Standard errors are clustered at the dialect level. All specifications include the full set of baseline controls: Free Space Signal Propagation (radio signal in the absence of geomorphological obstacles), railroad access, ruggedness of the terrain, river and coastal access, area, treaty port status, distance to major cities, distance to Beijing (geographic controls), historical Buddhist temples, 1964 population, 1964 gender ratio, 1964 fraction of non agricultural population, number of households, and ethnolinguistic fragmentation (socioeconomic & demographic controls), and a province fixed effect. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 3.6. Clustering at Different Geographic Units

Specifications:	Dep Var: Percent of Population Killed					
	Cluster on nearest station		Cluster on station by dialect		Conley's SEs w/ 150km cutoff	
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbb{1}\{\text{Signal}\} \times \mathbb{1}\{\text{Mandarin}\}$	0.039* (0.020)	–	0.039** (0.017)	–	0.039** (0.017)	–
$\mathbb{1}\{\text{Signal}\} \times \text{Experimental Intelligibility}$	–	0.152** (0.074)	–	0.152** (0.064)	–	0.152*** (0.058)
Observations	1985	1969	1985	1969	1985	1969
Clusters	114	114	221	221		
R^2	0.199	0.223	0.017	0.195	0.199	0.222

Note: The table reports the impact of radio exposure on percent killed, with standard errors clustered at different geographic levels. Each pair of columns corresponds to columns (4) and (8) of Table 3.3. In columns (1) and (2), standard errors are clustered at the nearest radio station level. Columns (3) and (4) clusters at the radio station-by-dialect level. Columns (5) and (6) report the Conley (1999) standard errors, correcting for two-dimensional spatial correlation with a 150km cutoff. All specifications include the full set of baseline controls: Free Space Signal Propagation (radio signal in the absence of geomorphological obstacles), railroad access, ruggedness of the terrain, river and coastal access, area, treaty port status, distance to major cities, distance to Beijing (geographic controls), historical Buddhist temples, 1964 population, 1964 gender ratio, 1964 fraction of non agricultural population, number of households, ethnolinguistic fragmentation (socioeconomic & demographic controls), and a province fixed effect. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 3.7. Additional Controls

	Dep Var: Percent of Population Killed					
	Baseline controls (1)	Baseline controls × Mandarin (2)	Nearest radio station FE (3)	Agricultural suitability (4)	Historical development (5)	Historical conflict (6)
<i>Panel A: Mandarin Dummy</i>						
$\mathbb{1}\{\text{Signal}\} \times \mathbb{1}\{\text{Mandarin}\}$	0.039** (0.018)	0.048*** (0.018)	0.037* (0.019)	0.041* (0.023)	0.037* (0.021)	0.035** (0.017)
Observations	1985	1985	1985	1338	1338	1985
<i>Panel B: % of Mandarin Words Comprehended</i>						
$\mathbb{1}\{\text{Signal}\} \times \text{Experimental Intelligibility}$	0.151** (0.076)	0.164** (0.069)	0.121* (0.066)	0.168** (0.079)	0.158** (0.077)	0.137* (0.073)
Observations	1969	1969	1969	1323	1323	1969

Note: This table examines the stability and robustness of the results to including additional sets of control variables. Panel A corresponds to the specification using *Mandarin* as an explanatory variable and Panel B corresponds to the specification using *Experimental Intelligibility* as an explanatory variable. In the first column, we present the baseline specification which corresponds to columns (4) and (8) of Table 3.3. Then in the remaining specifications we use different additional sets of controls as described in the text. All specifications include the full set of baseline controls: Free Space Signal Propagation (radio signal in the absence of geomorphological obstacles), railroad access, ruggedness of the terrain, river and coastal access, area, treaty port status, distance to major cities, distance to Beijing (geographic controls), historical Buddhist temples, 1964 population, 1964 gender ratio, 1964 fraction of non agricultural population, number of households, and ethnolinguistic fragmentation (socioeconomic & demographic controls), and a province fixed effect. In column (2), we interact the respective definitions of the *Mandarin* variable (binary indicator and *Experimental Intelligibility*) with baseline control variables. In column (4), we control for crop suitability of the local soil (this includes suitability of grain, wheat, rice, and millet). In column (5) we include controls for historical development variables: the number of civil service entrants and imperial exam qualifiers. In column (6), we control for conflict during the Taiping and Boxer Rebellion as well as revolutionaries in the initial Republican revolution. Standard errors are clustered at the dialect level. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 3.8. Heterogeneity

	Dep Var: Percent of Population Killed							
	Primary Schools		Communist Party Membership		Great Famine Intensity		Linguistic Fragmentation	
	Median: 423 Schools		Median: 1.08%		Median: 15.4%			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Radio Signal x Mandarin	Below	Above	Below	Above	Below	Above	No frag.	Frag.
	0.060** (0.023)	0.015 (0.023)	0.062* (0.032)	0.041** (0.019)	0.029 (0.026)	0.044* (0.024)	0.039** (0.019)	0.088* (0.052)
Observations	943	1042	895	1090	808	1177	1846	139
R^2	0.307	0.113	0.261	0.194	0.123	0.327	0.212	0.479

Note: Heterogeneous variables are measured at the province level, except for linguistic fragmentation, which is measured at the county level. In each column, we restrict the sample to only observations described by the column headers. In columns (1) - (6), the restrictions are below or above the median of the specified variable. In columns (7) and (8), the data restrictions are no linguistic heterogeneity (7), and presence of linguistic heterogeneity (8). The median of linguistic heterogeneity is 0. All specifications include the full set of controls as described in Table 3.3 and a province fixed effect. All specifications include the full set of baseline controls: Free Space Signal Propagation (radio signal in the absence of geomorphological obstacles), railroad access, ruggedness of the terrain, river and coastal access, area, treaty port status, distance to major cities, distance to Beijing (geographic controls), historical Buddhist temples, 1964 population, 1964 gender ratio, 1964 fraction of non agricultural population, number of households, and ethnolinguistic fragmentation (socioeconomic & demographic controls), and a province fixed effect. Standard errors are clustered at the dialect level. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 3.9. Summary Statistics

Variable	Obs	Mean	Std. Dev.
Dependent Variables			
Sent Down	13350	.014	.119
Communist Party	13350	.071	.256
Independent Variables			
Mandarin	13350	.145	.353
CR Cohort	13350	.217	.412
Education	13350	5.909	4.795
Age	13350	51.244	12.311
Gender	13350	.496	.5
Urban area (Census Bureau's definition)	13350	.382	.486
Father's level of education	13350	1.763	1.014
Mother's level of education	13350	1.355	.734

Note: Sent Down and Communist Party are binary variables. Mandarin is an indicator for speaking Mandarin at home. CR Cohort is an indicator for belonging to the Cultural Revolution Cohort, defined as being ages 10-21 in 1966. Education values range from 1 to 22, in increasing level of education. Gender is percent male.

Table 3.10. Individual Behavior

Dep var:	Send Down		Communist Party	
	Mean value: 0.014		Mean value: 0.071	
	(1)	(2)	(3)	(4)
CR Cohort x Mandarin	0.043** (0.019)	0.058** (0.027)	0.058*** (0.019)	0.060** (0.027)
Mandarin	-0.005 (0.009)	-0.006 (0.009)	0.016	0.012
CR Cohort	0.013*** (0.005)	-0.032 (0.115)	0.012 (0.007)	0.015 (0.109)
Observations	13350	13350	13350	13350
R^2	0.297	0.323	0.215	0.238
Additional Interactions	–	X	–	X

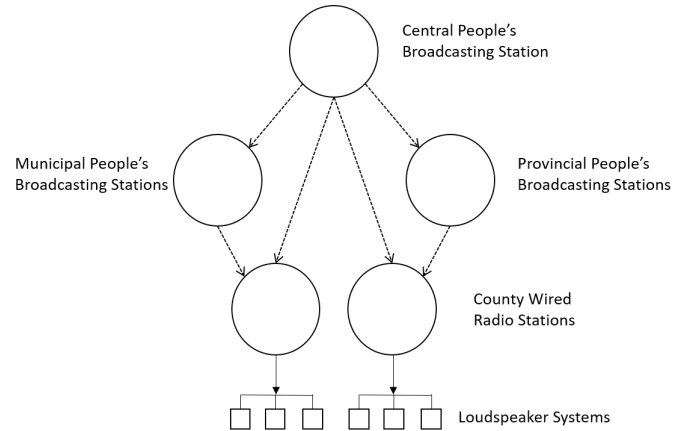
Note: The dependent variable in the first two columns is participation in the Send Down Movement, and in the second two columns is joining the Communist Party. Mandarin is an indicator for speaking Mandarin at home. CR cohort is an indicator for being age 10 to 21 in 1966. Controls include gender, education, age, age squared, father's education, mother's education, father's political party, mother's political party, father's occupation, mother's occupation, birth county, urban area of residence dummy, father's birth year, mother's birth year, own birth year, ethnicity, and parents' hukou status. Additional interaction controls include interactions between CR cohort and education, gender, urban dummy, birth province, and ethnicity. The coefficient of interest is on the interaction term, which measures the effect of speaking Mandarin for the Cultural Revolution cohort on participation in the Send Down Movement. Standard errors are clustered on the last known county of residence before the Cultural Revolution, and county of birth if born after the Cultural Revolution. * $p < .10$, ** $p < .05$, *** $p < .01$

Table 3.11. Send Down Movement: Mandarin Speakers in Mandarin-Speaking Counties

	(1)	(2)
Sample restriction:	Full sample	CR Cohort
Mandarin Speaker x Mandarin County x CR Cohort	0.066 (0.042)	–
Mandarin Speaker x Mandarin County	0.018* (0.010)	0.072* (0.040)
Sum of coefficients	0.073 (0.041)	
p-value		
Observations	16487	2893
R^2	0.151	0.547


Note: The first column uses Send Down status as the dependent variable, including individuals of all age cohorts in a triple interaction specification. The second column restricts the regression to the CR cohort. In the first column, the sum of the coefficients of the triple interaction and *Mandarin Speaker x Mandarin County* and the associated p-value are presented. Controls include an indicator for Mandarin speaker, indicator for Mandarin-speaking county, an indicator for being part of the CR cohort, and their double interactions. The baseline controls are gender, education, age, age squared, father's education, mother's education, father's political party, mother's political party, father's occupation, mother's occupation, birth county, urban area of residence dummy, father's birth year, mother's birth year, own birth year, ethnicity, and parents' hukou status, and interactions between the CR cohort indicator and education, gender, urban dummy, birth province, and ethnicity. Standard errors are clustered on the last known county of residence before the Cultural Revolution, and county of birth if born after the Cultural Revolution. * $p < .10$, ** $p < .05$, *** $p < .01$

Figure 3.1. Radio Diffusion Exchange




Note: This is a reproduction of a figure in Liu (1964). It shows the overall structure of the radio diffusion exchanges. The dotted lines represent signals sent over the airwaves and solid lines represent wired connections.

Figure 3.2. Extent of the Violence during the Cultural Revolution




Note: In this figure, more darkly-shaded regions correspond to areas with a higher number of fatalities directly attributed to the Cultural Revolution.

Figure 3.3. Predicted Radio Signal



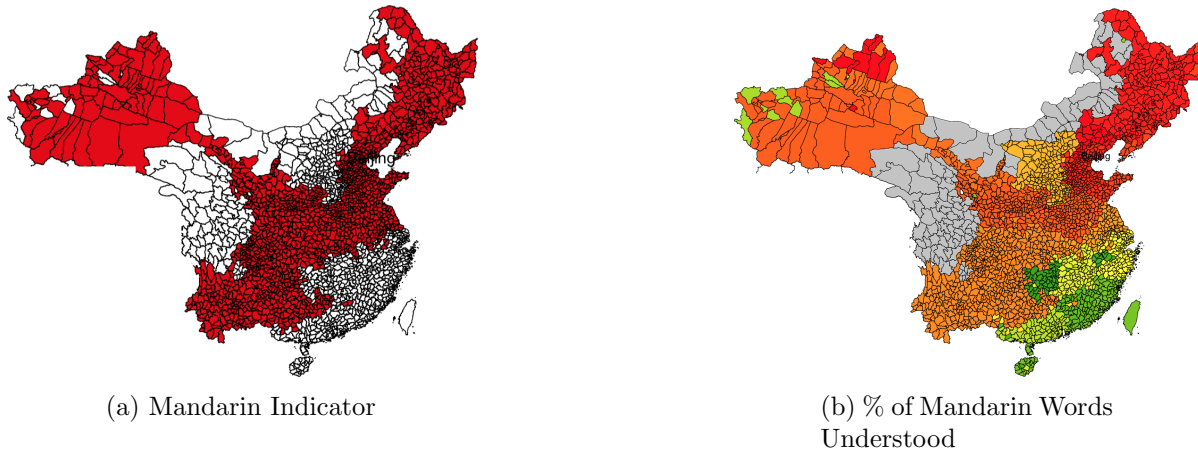
(a) Signal Strength



(b) Free Space Signal Strength

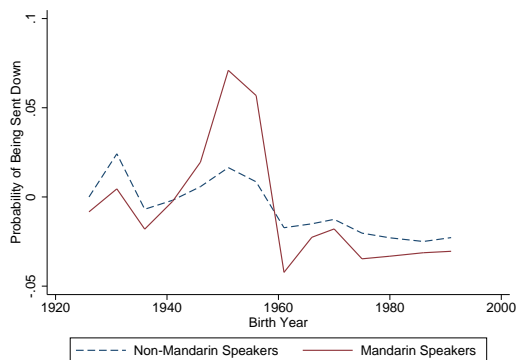
Note: The maps represent the geographic distribution of the simulated intensity of radio signal in 1964, under real conditions (top) and in the absence of geomorphological obstacles (bottom). The geographic boundaries of the map of China are dictated by the availability of violence data from Figure 3.2.

Figure 3.4. Linguistic Distance to Mandarin



Note: The maps present the geographic distribution of linguistic distance to Mandarin as it is operationalized in the paper. The geographic boundaries of the map of China are dictated by the availability of violence data from Figure 3.2. On the left, the red region indicates areas whose main dialect belongs to the Mandarin family while the white region indicates areas whose main dialect does not belong to the Mandarin family. On the right, redder shades represent a higher fraction of Beijing Mandarin words identified by speakers (of the dialect of vernacular Chinese) in that region, while greener shades represent a smaller fraction identified. The gray regions indicate missing data from the linguistic study, Tang and Van Heuven (2009). We note that there is variation in comprehension of Beijing Mandarin even within the Mandarin family of dialects.

Figure 3.5. Participation in the Send Down Movement: Coefficients of Mandarin x Age Cohort



Note: This figure plots regression-adjusted trends in Send Down Movement participation between Mandarin and non-Mandarin speakers. The participation rate for non-Mandarin speakers in 1921 is normalized to 0. The outcome is an indicator variable for participation in the Send Down Movement. The regression includes a full set of province and county fixed effects. Additional controls include parental occupation, parental income, migration history, hukou status, parental hukou status, ethnicity, parent political party affiliations, and interactions between birth cohort and Han ethnicity. Standard errors are clustered on the last known county of residence before the Cultural Revolution, and county of birth if born after the Cultural Revolution.

References

- Aaronson, Daniel (1998), “Using sibling data to estimate the impact of neighborhoods on children’s educational outcomes.” *Journal of Human Resources*, 915–946.
- Abaluck, Jason (2011), “What would we eat if we knew more: The implications of a large-scale change in nutrition labeling.” *Work. Pap., Mass. Inst. Technol.*
- Acemoglu, Daron and Joshua Angrist (2000), “How large are human-capital externalities? evidence from compulsory schooling laws.” *NBER macroeconomics annual*, 15, 9–59.
- Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa, and Ekaterina Zhuravskaya (2015), “Radio and the rise of the nazis in prewar germany.” *The Quarterly Journal of Economics*, 130, 1885–1939.
- Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel (2014), “Regulating consumer financial products: Evidence from credit cards.” *The Quarterly Journal of Economics*, 130, 111–164.
- Alesina, Alberto and Nicola Fuchs-Schündeln (2007), “Good-bye lenin (or not?): The effect of communism on people’s preferences.” *The American Economic Review*, 97, 1507–1528.
- Alesina, Alberto and Eliana La Ferrara (2000), “Participation in heterogeneous communities.” *The quarterly journal of economics*, 115, 847–904.
- Allcott, Hunt, Rebecca Diamond, and Jean-Pierre Dubé (2017), “The geography of poverty and nutrition: Food deserts and food choices across the united states.” Technical report, National Bureau of Economic Research.
- Alwin, Duane F and Jon A Krosnick (1991), “Aging, cohorts, and the stability of sociopolitical orientations over the life span.” *American Journal of Sociology*, 169–195.

- Anderson, Michael L. and David A. Matsa (2011), “Are restaurants really supersizing america?” *American Economic Journal: Applied Economics*, 3, 152–88, URL <http://www.aeaweb.org/articles.php?doi=10.1257/app.3.1.152>.
- Appleton, Simon, John Knight, Lina Song, and Qingjie Xia (2009), “The economics of communist party membership: the curious case of rising numbers and wage premium during china’s transition.” *The Journal of Development Studies*, 45, 256–275.
- Bai, Liang (2014), “Economic legacies of the cultural revolution.”
- Bai, Ying and Ruixue Jia (2016), “Elite recruitment and political stability: The impact of the abolition of china’s civil service exam.” *Econometrica*, 84, 677–733.
- Barnes, Dayle (1982), “Nationalism and the mandarin movement: the first half-century.” *Language spread: Studies in diffusion and social change*, 260–290.
- Baum-Snow, Nathaniel, Loren Brandt, J Vernon Henderson, Matthew A Turner, and Qinghua Zhang (2017), “Roads, railroads, and decentralization of chinese cities.” *Review of Economics and Statistics*, 99, 435–448.
- Bazzi, Samuel, Arya Gaduh, Alexander D Rothenberg, and Maisy Wong (2017), “Unity in diversity? ethnicity, migration, and nation building in indonesia.”
- Bhargava, Saurabh and Dayanand Manoli (2015), “Psychological frictions and the incomplete take-up of social benefits: Evidence from an irs field experiment.” *The American Economic Review*, 105, 3489–3529.
- Bollinger, Bryan, Phillip Leslie, and Alan Sorensen (2011), “Calorie posting in chain restaurants.” *American Economic Journal: Economic Policy*, 3, 91–128.
- Bronnenberg, Bart J, Jean-Pierre H Dubé, and Matthew Gentzkow (2012), “The evolution of brand preferences: Evidence from consumer migration.” *American Economic Review*, 102, 2472–2508.
- Bursztyn, Leonardo and Davide Cantoni (2016), “A tear in the iron curtain: The impact of western television on consumption behavior.” *Review of Economics and Statistics*, 98, 25–41.

- Cantoni, Davide, Yuyu Chen, David Y Yang, Noam Yuchtman, and Y Jane Zhang (2017), "Curriculum and ideology." *Journal of Political Economy*, 125, 338–392.
- Card, David, Andrew Johnston, Pauline Leung, Alexandre Mas, and Zhuan Pei (2015), "The effect of unemployment benefits on the duration of unemployment insurance receipt: New evidence from a regression kink design in missouri, 2003-2013." Technical report, National Bureau of Economic Research.
- Carrell, Scott E, Mark Hoekstra, and James E West (2011), "Is poor fitness contagious?: Evidence from randomly assigned friends." *Journal of Public Economics*, 95, 657–663.
- Chen, Yuyu and David Y Yang (2017), "1984 or the brave new world? evidence from a field experiment on media censorship in china."
- Chetty, Raj and Nathaniel Hendren (2015), "The impacts of neighborhoods on intergenerational mobility: Childhood exposure effects and county-level estimates." Technical report, Working Paper, May 2015, <http://www.equality-of-opportunity.org/images/nbhds-paper.pdf> (accessed May 5, 2015).
- Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler (2016), "The association between income and life expectancy in the united states, 2001-2014." *JAMA*, 315, 1750–1766.
- Chiang, Chun-Fang and Brian Knight (2011), "Media bias and influence: Evidence from newspaper endorsements." *The Review of Economic Studies*, rdq037.
- Christakis, Nicholas A and James H Fowler (2007), "The spread of obesity in a large social network over 32 years." *New England journal of medicine*, 357, 370–379.
- Cook, Alexander C (2014), *Mao's Little Red Book: A Global History*. Cambridge University Press.
- Currie, Janet, Stefano DellaVigna, Enrico Moretti, and Vikram Pathania (2010), "The effect of fast food restaurants on obesity and weight gain." *American Economic Journal: Economic Policy*, 2, 32–63.
- DellaVigna, Stefano, Ruben Enikolopov, Vera Mironova, Maria Petrova, and Ekaterina Zhuravskaya (2014), "Cross-border media and nationalism: Evidence from serbian radio

- in croatia.” *American Economic Journal: Applied Economics*, 6, 103–132.
- DellaVigna, Stefano and Matthew Gentzkow (2010), “Persuasion: empirical evidence.” *Annu. Rev. Econ.*, 2, 643–669.
- DellaVigna, Stefano and Ethan Kaplan (2007), “The Fox News effect: Media bias and voting.” *The Quarterly Journal of Economics*, 122, 1187–1234.
- Department of Health and Human Services (2003), *21 CFR Part 101 Food Labeling; Trans Fatty Acids in Nutrition Labeling; Consumer Research to Consider Nutrient Content and Health Claims and Possible Footnote or Disclosure Statements; Final Rule and Proposed Rule*, 133 edition, volume 68. Federal Register.
- Diez Roux, Ana V and Christina Mair (2010), “Neighborhoods and health.” *Annals of the New York Academy of Sciences*, 1186, 125–145.
- Dranove, D and GZ Jin (2011), “Quality disclosure and certification.” *the Journal of Economic Literature*.
- Dubois, Pierre, Rachel Griffith, and Martin OConnell (2017), “The effects of banning advertising in junk food markets.” *The Review of Economic Studies*, rdx025.
- Durand, Casey P, Mohammad Andalib, Genevieve F Dunton, Jennifer Wolch, and Mary Ann Pentz (2011), “A systematic review of built environment factors related to physical activity and obesity risk: implications for smart growth urban planning.” *Obesity Reviews*, 12, e173–e182.
- Durante, Ruben, Paolo Pinotti, Andrea Tesei, et al. (2015), *The Political Legacy of Entertainment TV*. Centre for Economic Policy Research.
- Easterly, William and Ross Levine (1997), “Africa’s growth tragedy: policies and ethnic divisions.” *The Quarterly Journal of Economics*, 1203–1250.
- Eckel, Robert H, Penny Kris-Etherton, Alice H Lichtenstein, Judith Wylie-Rosett, Allison Groom, Kimberly F Stitzel, and Shirley Yin-Piazza (2009), “Americans’ awareness, knowledge, and behaviors regarding fats: 2006-2007.” *Journal of the American Dietetic Association*, 109, 288–296.

- Einav, Liran, Ephraim Leibtag, and Aviv Nevo (2010), “Recording discrepancies in nielsen homescan data: Are they present and do they matter?” *QME*, 8, 207–239.
- Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya (2011), “Media and political persuasion: Evidence from russia.” *The American Economic Review*, 101, 3253–3285.
- Escobar-Chaves, Soledad Liliana and Craig A Anderson (2008), “Media and risky behaviors.” *The Future of Children*, 18, 147–180.
- Esteban, Joan, Laura Mayoral, and Debraj Ray (2012), “Ethnicity and conflict: An empirical study.” *The American Economic Review*, 102, 1310–1342.
- Esteban, Joan and Debraj Ray (2008), “On the salience of ethnic conflict.” *The American Economic Review*, 98, 2185–2202.
- Figlio, David N and Maurice E Lucas (2004), “What’s in a grade? school report cards and the housing market.” *The American Economic Review*, 94, 591–604.
- Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams (2014), “Sources of geographic variation in health care: Evidence from patient migration.” Technical report, National Bureau of Economic Research.
- Flanagan, Constance A and Lonnie R Sherrod (1998), “Youth political development: An introduction.” *Journal of Social Issues*, 54, 447–456.
- Gerber, Alan S, Dean Karlan, and Daniel Bergan (2009), “Does the media matter? a field experiment measuring the effect of newspapers on voting behavior and political opinions.” *American Economic Journal: Applied Economics*, 1, 35–52.
- Giuliano, Paola and Antonio Spilimbergo (2014), “Growing up in a recession.” *The Review of Economic Studies*, 81, 787–817.
- Glaeser, Edward L, Bruce Sacerdote, and Jose A Scheinkman (1996), “Crime and social interactions.” *The Quarterly Journal of Economics*, 111, 507–548.
- Gong, Jie, Yi Lu, and Huihua Xie (2014), “Adolescent adversity and long-run health.”
- Gong, Jie, Yi Lu, and Huihua Xie (2015), “Adolescent environment and noncognitive skills.”

- Greenstone, Michael, Paul Oyer, and Annette Vissing-Jorgensen (2006), “Mandated disclosure, stock returns, and the 1964 securities acts amendments.” *The Quarterly Journal of Economics*, 121, 399–460.
- Guyton, John, Dayanand S Manoli, Brenda Schafer, and Michael Sebastiani (2016), “Reminders & recidivism: Evidence from tax filing & eicr participation among low-income nonfilers.” Technical report, National Bureau of Economic Research.
- Handbury, Jessie, Ilya Rahkovsky, and Molly Schnell (2015), “Is the focus on food deserts fruitless? retail access and food purchases across the socioeconomic spectrum.” Technical report, National Bureau of Economic Research.
- Hastings, Justine, Christopher A Neilson, and Seth D Zimmerman (2015), “The effects of earnings disclosure on college enrollment decisions.” Technical report, National Bureau of Economic Research.
- Hersch, Joni and W Kip Viscusi (1990), “Cigarette smoking, seatbelt use, and differences in wage-risk tradeoffs.” *Journal of Human Resources*, 202–227.
- Hill, Seth J, James Lo, Lynn Vavreck, and John Zaller (2013), “How quickly we forget: The duration of persuasion effects from mass communication.” *Political Communication*, 30, 521–547.
- Hoffmann, Hilmar (1996), *The triumph of propaganda: Film and national socialism, 1933-1945*, volume 1. Berghahn Books.
- Houn, Franklin W (1957), “Radio broadcasting and propaganda in communist china.” *Journalism & Mass Communication Quarterly*, 34, 366–377.
- Howse, Hugh (1960), “The use of radio in china.” *The China Quarterly*, 2, 59–68.
- Hu, Frank B., Meir J. Stampfer, JoAnn E. Manson, Eric Rimm, Graham A. Colditz, Bernard A. Rosner, Charles H. Hennekens, and Walter C. Willett (1997), “Dietary fat intake and the risk of coronary heart disease in women.” *New England Journal of Medicine*, 337, 1491–1499, URL <http://dx.doi.org/10.1056/NEJM199711203372102>. PMID: 9366580.

- Hufford, George A (2002), “The its irregular terrain model, version 1.2. 2 the algorithm.” *Institute for Telecommunication Sciences, National Telecommunications and Information Administration, US Department of Commerce.* <http://flattop.its.blrdoc.gov/itm.html>.
- Hut, Stefan and Emily Oster (2018), “Changes in household diet: Determinants and predictability.”
- Ippolito, Pauline M and Richard A Ippolito (1984), “Measuring the value of life saving from consumer reactions to new information.” *Journal of Public Economics*, 25, 53–81.
- Ippolito, Pauline M and Alan D Mathios (1991), “Information, advertising, and health choices: a study of the cereal market.” In *Economics of Food Safety*, 211–246, Springer.
- Jan, George P. (1967), “Radio propaganda in chinese villages.” *Asian Survey*, 7, 305–315, URL <http://www.jstor.org/stable/2642659>.
- Jin, Ginger Zhe and Phillip Leslie (2003), “The effect of information on product quality: Evidence from restaurant hygiene grade cards.” *The Quarterly Journal of Economics*, 118, 409–451.
- Jin, Ginger Zhe and Alan T Sorensen (2006), “Information and consumer choice: the value of publicized health plan ratings.” *Journal of Health Economics*, 25, 248–275.
- Jing, Lin (1991), *The Red Guards’ path to violence: political, educational, and psychological factors*. Praeger Publishers.
- Kahneman, Daniel, Peter P Wakker, and Rakesh Sarin (1997), “Back to bentham? explorations of experienced utility.” *The Quarterly Journal of Economics*, 112, 375–406.
- Katz, Elihu and F Paul (1955), “Lazarsfeld (1955), personal influence.” *The Part Played by People in the Flow of Mass Communication*. New York.
- Kiesel, Kristin and Sofia B Villas-Boas (2013), “Can information costs affect consumer choice? nutritional labels in a supermarket experiment.” *International Journal of Industrial Organization*, 31, 153–163.
- Kinnan, Cynthia, Shing-Yi Wang, and Yongxiang Wang (2015), “Relaxing migration constraints for rural households.”

- Klein, Jonathan D, Jane D Brown, Carol Dykers, Kim Walsh Childers, Janice Oliveri, and Carol Porter (1993), "Adolescents' risky behavior and mass media use." *Pediatrics*, 92, 24–31.
- Kling, Jeffrey R, Jeffrey B Liebman, and Lawrence F Katz (2007), "Experimental analysis of neighborhood effects." *Econometrica*, 75, 83–119.
- Kung, James Kai-sing and Justin Yifu Lin (2003), "The causes of china's great leap famine, 1959–1961." *Economic Development and Cultural Change*, 52, 51–73.
- La Porta, Rafael, Florencio Lopez-de Silanes, Andrei Shleifer, and Robert Vishny (1999), "The quality of government." *Journal of Law, Economics, and Organization*, 15, 222–279.
- Lally, Philippa, Cornelia HM Van Jaarsveld, Henry WW Potts, and Jane Wardle (2010), "How are habits formed: Modelling habit formation in the real world." *European Journal of Social Psychology*, 40, 998–1009.
- Latham, Kevin (2007), *Pop culture China!: media, arts, and lifestyle*. ABC-CLIO.
- Lazarsfeld, Paul F, Bernard Berelson, and Hazel Gaudet (1944), "The peoples choice: How the voter makes up his mind in a presidential election." *New York: Duell, Sloan and Pearce*.
- Lin, Chung-Tung J and Steven T Yen (2010), "Knowledge of dietary fats among us consumers." *Journal of the American Dietetic Association*, 110, 613–618.
- Liu, Alan PL (1964), "Radio broadcasting in communist china." Technical report, DTIC Document.
- Liu, Alan PL (1971), *Communications and national integration in Communist China*. 1, Univ of California Press.
- Ludwig, Jens, Greg J Duncan, Lisa A Gennetian, Lawrence F Katz, Ronald C Kessler, Jeffrey R Kling, and Lisa Sanbonmatsu (2013), "Documentation of the mto adult public use datasets for the american economic review papers & proceedings article:long-term neighborhood effects on low-income families: Evidence from moving to opportunity."

- Macintyre, Sally, Anne Ellaway, and Steven Cummins (2002), “Place effects on health: how can we conceptualise, operationalise and measure them?” *Social science & medicine*, 55, 125–139.
- Manoli, Dayanand S and Nicholas Turner (2014), “Nudges and learning: evidence from informational interventions for low-income taxpayers.” Technical report, National Bureau of Economic Research.
- Markham, James W and Alan PL Liu (1969), “Mass communication and media in china’s cultural revolution.” *Journalism & Mass Communication Quarterly*, 46, 314–319.
- Mathios, Alan D (2000), “The impact of mandatory disclosure laws on product choices: An analysis of the salad dressing market.” *The Journal of Law and Economics*, 43, 651–678.
- McCormack, Gavin R, Melanie Rock, Ann M Toohey, and Danica Hignell (2010), “Characteristics of urban parks associated with park use and physical activity: a review of qualitative research.” *Health & place*, 16, 712–726.
- Meng, Xin and Robert Gregory (2007), “Exploring the impact of interrupted education on earnings: the educational cost of the chinese cultural revolution.”
- Meng, Xin, Nancy Qian, and Pierre Yared (2015), “The institutional causes of china’s great famine, 1959–1961.” *The Review of Economic Studies*, 82, 1568–1611.
- Michalopoulos, Stelios (2012), “The origins of ethnolinguistic diversity.” *The American Economic Review*, 102, 1508–1539.
- Mohapatra, Desi Prasad, Nadia Streletskaya, Emily Wang, and Hongli Wei (2017), “Welfare analysis of trans fat ban: The microwaveable popcorn market.”
- Mozaffarian, Dariush, Antti Aro, and Walter C Willett (2009), “Health effects of trans-fatty acids: experimental and observational evidence.” *European journal of clinical nutrition*, 63, S5–S21.
- Mozaffarian, Dariush, Martijn B. Katan, Alberto Ascherio, Meir J. Stampfer, and Walter C. Willett (2006), “Trans fatty acids and cardiovascular disease.” *New England Journal of Medicine*, 354, 1601–1613, URL <http://dx.doi.org/10.1056/NEJMra054035>.

PMID: 16611951.

- Nevo, Aviv (2001), “Measuring market power in the ready-to-eat cereal industry.” *Econometrica*, 69, 307–342.
- Niederdeppe, Jeff and Dominick L Frosch (2009), “News coverage and sales of products with trans fat: effects before and after changes in federal labeling policy.” *American journal of preventive medicine*, 36, 395–401.
- Nunn, Nathan and Diego Puga (2012), “Ruggedness: The blessing of bad geography in africa.” *Review of Economics and Statistics*, 94, 20–36.
- Oh, Kyungwon, Frank B Hu, JoAnn E Manson, Meir J Stampfer, and Walter C Willett (2005), “Dietary fat intake and risk of coronary heart disease in women: 20 years of follow-up of the nurses’ health study.” *American Journal of Epidemiology*, 161, 672–679.
- Olken, Benjamin A (2009), “Do television and radio destroy social capital? evidence from indonesian villages.” *American Economic Journal: Applied Economics*, 1, 1–33.
- Oster, Emily (2017), “Diabetes and diet: Purchasing behavior change in response to health information.” *American Economic Journal: Applied Economics*.
- Otite, Fadar O, Michael F Jacobson, Aspan Dahmubed, and Dariush Mozaffarian (2013), “Peer reviewed: Trends in trans fatty acids reformulations of us supermarket and brand-name foods from 2007 through 2011.” *Preventing chronic disease*, 10.
- Pan, Yihong (2003), *Tempered in the revolutionary furnace: China’s youth in the rustication movement*. Lexington books.
- Phillips, Caleb, Douglas Sicker, and Dirk Grunwald (2011), “The stability of the longley-rice irregular terrain model for typical problems.” *arXiv preprint arXiv:1109.1843*.
- Pope, Devin G (2009), “Reacting to rankings: evidence from america’s best hospitals.” *Journal of Health Economics*, 28, 1154–1165.
- Pye, Lucian W (2015), *Communications and Political Development.(SPD-1)*. Princeton University Press.
- Ramsey, S Robert (1987), *The languages of China*. Princeton University Press.

- Rene, Helena K (2013), *China's Sent-down Generation: Public Administration and the Legacies of Mao's Rustication Program*. Georgetown University Press.
- Restrepo, Brandon J and Matthias Rieger (2016a), "Denmarks policy on artificial trans fat and cardiovascular disease." *American journal of preventive medicine*, 50, 69–76.
- Restrepo, Brandon J and Matthias Rieger (2016b), "Trans fat and cardiovascular disease mortality: Evidence from bans in restaurants in new york." *Journal of Health Economics*, 45, 176–196.
- Rogall, Thorsten (2014), "Mobilizing the masses for genocide."
- Sacerdote, Bruce (2001), "Peer effects with random assignment: Results for dartmouth roommates." *The Quarterly journal of economics*, 116, 681–704.
- Sanbonmatsu, Lisa, Nicholas A Potter, Emma Adam, Greg J Duncan, Lawrence F Katz, Ronald C Kessler, Jens Ludwig, Jordan Marvakov, Fanghua Yang, William J Congdon, et al. (2012), "The long-term effects of moving to opportunity on adult health and economic self-sufficiency." *Cityscape*, 109–136.
- Scanlon, Dennis P, Michael Chernew, Catherine McLaughlin, and Gary Solon (2002), "The impact of health plan report cards on managed care enrollment." *Journal of Health Economics*, 21, 19–41.
- Stamler, Jeremiah, Deborah Wentworth, and James D Neaton (1986), "Is relationship between serum cholesterol and risk of premature death from coronary heart disease continuous and graded?: findings in 356 222 primary screenees of the multiple risk factor intervention trial (mrfit)." *Jama*, 256, 2823–2828.
- Su, Yang (2011), *Collective killings in rural China during the cultural revolution*. Cambridge University Press.
- Sunstein, Cass R and Richard H Thaler (2003), "Libertarian paternalism is not an oxymoron." *The University of Chicago Law Review*, 1159–1202.
- T, Gordon and Kannel WB (1971), "Premature mortality from coronary heart disease: The framingham study." *JAMA*, 215, 1617–1625, URL +<http://dx.doi.org/10.1001/jama.1971.03180230027005>.


- Tang, Chaoju and Vincent J Van Heuven (2009), “Mutual intelligibility of chinese dialects experimentally tested.” *Lingua*, 119, 709–732.
- Train, Kenneth E (2009), *Discrete choice methods with simulation*. Cambridge university press.
- Unnevehr, Laurian J and Evelina Jagmanaite (2008), “Getting rid of trans fats in the us diet: policies, incentives and progress.” *Food Policy*, 33, 497–503.
- Verschuren, WM Monique, David R Jacobs, Bennie PM Bloemberg, Daan Kromhout, Alessandro Menotti, Christ Aravanis, Henry Blackburn, Ratko Buzina, Anastasios S Dontas, Flaminio Fidanza, et al. (1995), “Serum total cholesterol and long-term coronary heart disease mortality in different cultures: Twenty-fiveyear follow-up of the seven countries study.” *Jama*, 274, 131–136.
- Viscusi, W Kip (1990), “Do smokers underestimate risks?” *Journal of political Economy*, 98, 1253–1269.
- Viscusi, W Kip and Joseph E Aldy (2003), “The value of a statistical life: a critical review of market estimates throughout the world.” *Journal of risk and uncertainty*, 27, 5–76.
- Voigtländer, Nico and Hans-Joachim Voth (2014), “Highway to hitler.” Technical report, National Bureau of Economic Research.
- Voigtländer, Nico and Hans-Joachim Voth (2015), “Nazi indoctrination and anti-semitic beliefs in germany.” *Proceedings of the National Academy of Sciences*, 112, 7931–7936.
- Volland, Nicolai (2003), *The Control of the Media in the People’s Republic of China*. Ph.D. thesis.
- Walder, Andrew G and Yang Su (2003), “The cultural revolution in the countryside: Scope, timing and human impact.” *The China Quarterly*, 173, 74–99.
- Xu, Yanqing and Fahui Wang (2015), “Built environment and obesity by urbanicity in the us.” *Health & place*, 34, 19–29.
- Yanagizawa-Drott, David (2014), “Propaganda and conflict: Evidence from the rwandan genocide.” *The Quarterly Journal of Economics*, 129, 1947–1994.
- Zaller, John (1992), *The nature and origins of mass opinion*. Cambridge university press.

- Zhou, Weina (2013), “How does a traumatic experience during youth affect life later? the long-term impact of the send-down program during the chinese cultural revolution.” Technical report, Working Paper.
- Zhou, Xueguang (2004), *The state and life chances in urban China: Redistribution and stratification, 1949–1994*. Cambridge University Press.

Appendix

Appendix 1

Figure 1A.1. Treatment Effect, 5th Percentile of Residualized BMI



Note: This figure shows the results of a regression of individual BMI on individual observables and time in neighborhood for movers to low residualized BMI zip codes. Individual level controls are age, education, household income, race, gender, marital status, employment status, size of family, country of birth and survey-year fixed effects. The time variable is parameterized as 2-5 years, 5-10 years, and 10-15 years. The solid straight line is the median BMI in low residualized BMI areas, while the broken solid line shows the point estimates for the coefficients of each bin of time category. These estimates are in reference to the base category, which is the solid line at 0-2 years. The solid line at 0-2 years is the median BMI for movers into high residualized BMI areas who moved within the last 2 years. The dotted lines are the 95% confidence intervals. Standard errors are clustered at the zip code level.

Table 1A.1. Treatment Effects, 5th Percentile

BMI	Base
2-5 years	0.218 (0.155)
5-10 years	0.206 (0.199)
10-15 years	0.086 (0.186)
Observations	4825
R-squared	0.1150

Note: This table shows the analogous estimates from Figure 1A.1. Individual level controls are age, education, household income, race, gender, marital status, employment status, size of family, country of birth and survey-year fixed effects. *** p<0.01, ** p<0.05, * p<0.1

Appendix 2

Appendix 2A: Data Collection Process. This section provides a more detailed description of the data collection process. Purchase data comes from the Kilts Center for Marketing Nielsen Homescan Panel Data. To create the dataset of microwave popcorn, I restrict the data to the sample where the “Product Module Description” is “Popcorn - Unpopped.” I first calculate the total revenue of each brand between 2004 and 2006. Then I take the top 6 brands in terms of revenue, and aggregate the rest into an other category which includes the store brand. I then create a low fat category by scanning the UPC description for the string “LWF” or “100CP”. I abstract away from the total weight purchased, by collapsing the data into an indicator for if a brand was purchased during a particular shopping occasion. The purchases are aggregated up to a monthly level, with purchases coded as independent observations if multiple brands are purchased in a month (less than 5% of purchases). A price per ounce variable was created by dividing the total ounces purchased of a particular brand in a shopping trip by the total price paid. Households are selected to be part of the panel if they ever made a microwave popcorn purchase in the 4 - year sample period. In the logit estimation, households who make more than 50 purchases are dropped (less than 5% of households).

Trans fat data is collected through internet searches for news articles, as well as from data provided by the Center for Science in the Public Interest (CSPI). The CSPI data provides the total grams of trans fat in selected grocery store items from 2007 to 2011, which is in the post legislation period. If an item has a positive amount of trans fat post legislation, it most likely also had a positive amount of trans fat before the legislation. This data is supplemented with Google searches for news articles. For instance, a news release was published by Con Agra on December 14, 2005 that was titled: “ConAgra Foods Makes Smart-Snacking Breakthrough—Its Microwave Popcorn, Made with 100% Whole Grain, Will Also Be 0g Trans Fat.”⁴⁴ ConAgra is the parent company of both

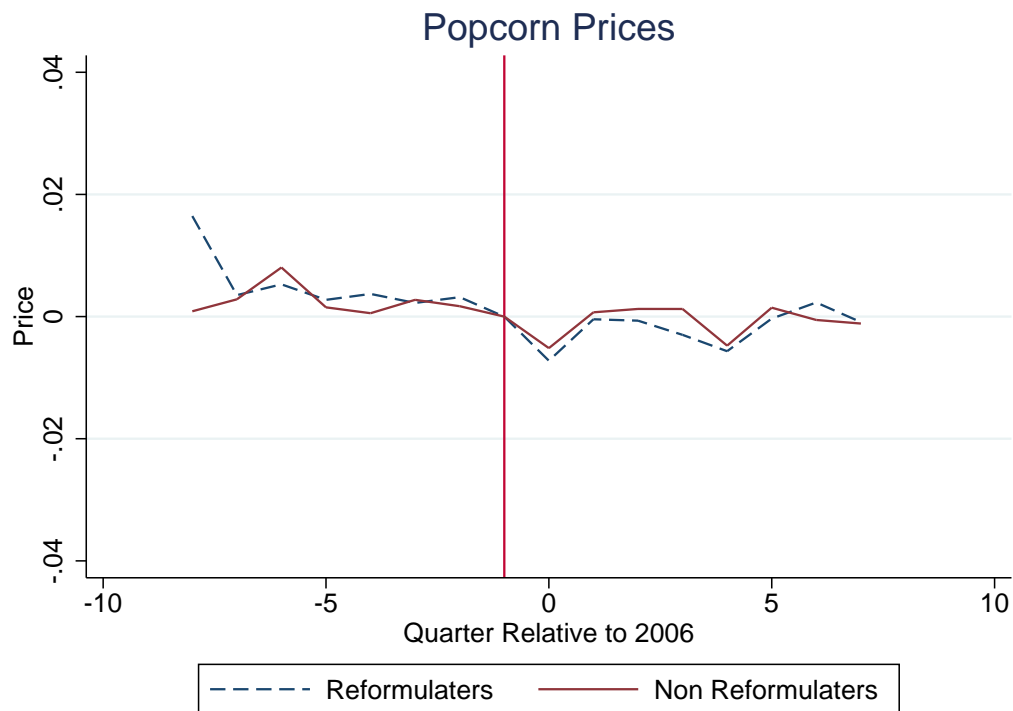
⁴⁴<http://www.conagrabrands.com/news-room/news-conagra-foods-makes-smart-snacking-breakthrough-its-microwave-popcorn-made-with-100-whole-grain-will-also-be-0g-trans-fat-1008479>

Orville Redenbacher's and Act II. Information contained within this article indicates that these brands contained trans fat before reformulation. Alternatively, Pop Weaver reformulated much later, as evidenced by a news article titled "Pop Weaver Launches First Microwavable Popcorn Made With Canola Oil," dated May 04, 2010.⁴⁵ I note that Jolly Time and Pop Weaver reformulated their products in the few years following the enactment of the labeling legislation, but Pop Secret did not reformulate until pressured to do so by the impending trans fat ban, approved by the FDA in 2015 and set to be enforced in 2018.

⁴⁵<https://www.prnewswire.com/news-releases/pop-weaver-launches-first-microwave-popcorn-made-with-canola-oil-92762939.html>

Appendix 2B: Tables and Figures.

Figure 2B.1. Popcorn Prices



Note: This figure shows the trends in prices, controlling for brand fixed effects, the trends in prices for reformulators, non reformulators, and the never TF group. The coefficients $\phi_t + quarter_t$, and $\omega_t + quarter_t$ are plotted from $price_{jt} = \sum_t \phi_t Non\ Reformulator_j * quarter_t + \sum_t \omega_t Reformulator_j * quarter_t + quarter_t + brand_j + \epsilon_{jt}$. The purpose is to show that the pre-trends for the reformulators and non reformulators are very similar, providing evidence that reformulators and non reformulators were very similar and there was no strategic reformulation. Price at time -1 is normalized to 0 for both groups.

Table 2B.1. Prices Per Ounce

	Data		Simulated	
	Pre-Prices	Label Prices	Ban Prices	Marginal Cost
Act II	0.1083	0.1088	0.1067	0.0671
Jolly Time	0.1123	0.1118	0.1132	0.0742
Orville Redenbacher's	0.1654	0.1658	0.1637	0.1197
Pop Secret	0.1632	0.1632	0.1606	0.1187
Pop Weaver	0.0625	0.0627	0.0616	0.0251
Smart Balance	0.2185	0.2191	0.2161	0.1727
Lowfat	0.2074	0.2086	0.2055	0.1613
Other	0.1098	0.1093	0.107	0.0675

Note: This table shows the average prices by brand pre-2006 as well as post-2006 in the left hand panel, as well as the simulated ban prices and marginal costs in the right hand panel.

Table 2B.2. First Stage

First Stage Endogenous Regressor	F-statistic
Price	1.00E+07
Price x Income	1.00E+07
Price x Education	1.00E+07
Price x Child	1.00E+07

Note: This table shows first stage F-statistics from the main logit specification using the Hausman instruments for the control function.

Appendix 2C: Cookies. In this section, I conduct a case study of the cookie industry, using the same framework as the analysis of the popcorn industry. I study the top 24 sub brands of cookies, which comprise around 40% of the total revenue of cookies sold. Cookies are not substitutable at the brand level (for instance, Nabisco Oreos versus Nabisco Chips Ahoy), so I use a finer level of categorization. I take the first three words of the UPC description, and the top ranked unique values of this categorization by revenue are the cookies I use in the analysis. I categorize the cookies the same way I categorize popcorn: into never TF, reformulators and non reformulators. The list of never TF include: Stauffer’s Animal Crackers, Nabisco Chips Ahoy Thin Crisps, and Nabisco Oreo Thin Crisps; reformulators: Nabisco Fig Newtons, Fruit Fig Newtons, Chunky Chips Ahoy, Chewy Chips Ahoy, Chips Ahoy, Oreos Classic, Oreos Double Stuffed, Mini Oreos, Teddy Grahams, Nilla, Nutter Butter, Snackwells; Pepperidge Farms Distinctive and Soft Baked, and non reformulators: Keebler Deluxe, Elf Sandwich, Sandies, Rainbow Chips Deluxe; Famous Amos Chocolate Chip; Little Debbie Nutty Bars. Compared to the popcorn market, the cookie market is larger, with more households buying, and making more frequent purchases. However, cookies are less of a homogeneous good than popcorn, and thus, are less substitutable. In addition, there are many different brands of cookies, and it is difficult to ascertain whether or not smaller brands, which would be grouped into the “other” category, contain trans fats, and when they reformulated. Thus, I omit this group.

I find that the willingness to pay for the trans fat label in the cookie market is \$0.08/ounce, compared to an average price of -\$0.19/ounce.⁴⁶ This is higher than the willingness to pay in the popcorn industry. The WTP for taste is negative as well, indicating that individuals prefer the taste of products without trans fat.

⁴⁶This is using the random coefficients logit, not the IV logit specification, so this number should be compared to the WTP of \$0.03 per ounce estimated for the popcorn market.

Table 2C.1. Summary Statistics

Variable	Mean	Std. Dev.	N
Cookies			
price per ounce	0.196	0.131	260850
average ounces per purchase	16.44957	13.262	1366108
total ounces by household year	262.7615	279.706	85522
unique households			22531

Note: The price per ounce and average weight per purchase variables are calculated from a sample where each observation is a purchase. The variable “total ounces by household year” is calculated from a dataset where each observation is a household - year.

Table 2C.2. Demographics: Cookies

Variable	Mean	Std. Dev.	Mean	Std. Dev.
No Trans		Before		After
Household Income	19.822	5.534	19.625	5.552
Education	4.102	0.961	4.076	0.964
Children	0.335	0.472	0.324	0.468
Trans				
Household Income	19.242	5.494	18.947	5.541
Education	4.001	0.963	3.963	0.958
Children	0.324	0.468	0.307	0.461

Note: Household Income is measured on a 27 point scale, where each increment is roughly \$5000. 3 refers to anything below \$5000 and 27 refers to anything over \$100,000. Education is on a scale of 1-6, where 1 is grade school and 6 is post grad. I use the maximum education of the household. Children is an indicator equal to 1 if there a child under the age of 18 is present.

Appendix 2D: Medical Literature on Trans Fat. In controlled trials involving trans fats, healthy subjects are administered a controlled diet, with different treatment groups receiving diets differing by the type of fat used. These fats include saturated, monounsaturated, and polyunsaturated. In a meta analysis of controlled trials, Mozaffarian et al. (2009) shows that there is an upward trend in the effect of trans fat on blood cholesterol levels for an isocaloric replacement of trans fat with saturated fat from 1% to 10% (decrease HDL, increase LDL, and increase the ratio of total cholesterol to HDL). Even an isocaloric replacement of energy from trans fats with saturated fat as small as 1% in the diet has shown to have negative effects on cholesterol. In one study cited by the meta analysis, increasing trans fatty acid consumption linearly increased LDL-C cholesterol levels. While these randomized controlled studies do not suffer from endogeneity issues, the time frame is short (around 5 weeks) and thus evaluates limited dietary changes. Longer studies would be prohibitively expensive and may be unethical.

Stamler et al. (1986) finds that the effect of blood cholesterol levels on CHD are continuous and graded, instead of following a threshold relationship. This is further evidence that the effect of trans fat on CHD may be evident even in small doses. Verschuren et al. (1995) also finds a linear effect of cholesterol levels on CHD across cultures.

For the long term effects of trans fats, the best evidence available is from prospective cohort studies. In a prospective study, diet is measured before the assessment of disease, and both fatal and non fatal diseases can be diagnosed (as opposed to a retrospective study). The most conclusive studies concerning the population effects of trans fats come from data collected from the Nurse's Health study, a cohort study following nurses since 1976. Hu et al. (1997) follow up with 80,000 women after 14 years, and find that the replacement of 2% of energy intake from trans fats with energy intake from unsaturated fats reduces the risk of CHD by 53%. They also find that a 2% increase in energy intake from trans fats, compared with carbohydrates, would increase relative risk of CHD by 16% after adjusting for age, BMI, smoking, family history of myocardial infarction, as well as saturated, monounsaturated and polyunsaturated fat and several other controls.

In a later followup study, using a horizon of 20 years, Oh et al. (2005) find that the relative risk of CHD would increase by 31% from an increase of 2% of energy intake from trans fats, adjusting for age, BMI, smoking, family history of myocardial infarction, as well as saturated, monounsaturated and polyunsaturated fat and several other controls. In a meta-analysis, Mozaffarian et al. (2009) surveys meta-analyses involving trans fats and several medical articles relating dietary trans fat consumption to health. They find that on average, a 2% increase isocaloric replacement of saturated fat with trans fat leads to an 20% increase in adjusted risk of developing CHD (after accounting for demographic controls).

These figures from the prospective cohort studies are also comparable to studies examining the effects of trans fat bans, in New York City and Denmark, for example. Restrepo and Rieger (2016a) finds that the Danish trans fat ban led to a decrease on average of about 14.2 deaths per 100,000 per year, or a 3.2% annual reduction in cardiovascular disease deaths. Restrepo and Rieger (2016b) find a 4.5% annual reduction in cardiovascular disease mortality rates from the New York City restaurant trans fat ban. If the rates found by Mozaffarian et al. (2009) are annualized, the numbers are 1.3% to 2.6%, which is in the ballpark of the estimates found by Restrepo and Rieger (2016b) and Restrepo and Rieger (2016a). The harmful effects of trans fat are cumulative and persistent, as evidenced by the effects of trans fats in randomized controlled trials in a short time span in small amounts, as well as over time in prospective cohort studies.

Appendix 3

Table 3A.1. Mandarin Comprehension

Speaker dialect	Listener dialect															Mean
	Suzhou	Wenzhou	Guangzhou	Xiamen	Fuzhou	Chaozhou	Meixian	Nanchang	Changsha	Taiyuan	Beijing	Jinan	Hankou	Chengdu	Xi'an	
Suzhou	77	7	5	18	13	5	7	13	13	20	5	18	15	15	7	16
Wenzhou	5	93	5	12	3	2	7	10	2	7	2	10	8	7	2	10
Guangzhou	5	7	92	10	20	25	55	22	13	7	3	22	8	17	7	21
Xiamen	13	5	8	97	23	28	13	18	13	3	5	15	7	17	8	18
Fuzhou	3	3	2	17	92	7	3	8	5	0	0	7	2	0	3	10
Chaozhou	7	0	3	52	13	98	3	12	3	7	2	13	10	3	5	15
Meixian	13	2	12	28	17	20	70	25	18	10	3	25	15	25	8	19
Nanchang	28	13	20	25	27	17	33	50	32	35	18	53	43	37	23	30
Changsha	12	3	8	23	17	3	17	25	93	13	13	38	53	28	2	23
Taiyuan	63	35	45	63	57	25	55	68	68	73	77	92	92	85	73	65
Beijing	87	62	90	90	93	60	80	78	92	90	98	98	97	98	93	87
Jinan	52	27	32	48	48	15	40	60	70	75	77	97	83	82	67	58
Hankou	48	32	32	52	53	27	45	53	62	58	67	95	100	73	65	57
Chengdu	47	22	40	48	72	27	48	58	62	65	62	98	95	95	68	60
Xi'an	53	33	50	58	57	30	57	58	63	68	58	82	78	70	67	59
Mean	34	22	30	43	40	26	36	37	41	35	33	51	47	43	33	

Note: This table is from Tang and Van Heuven (2009). Each entry represents the percent of words correctly understood of a particular dialect by a listener of another (or the same) dialect.

Table 3A.2. Robustness: Alternative Samples & Signal Cutoffs

	Dep Var: Percent of Population Killed							
	Censored Sample		Winsorized Sample		Alternative Signal Cutoffs			
					Greater than Median		Greater than Zero	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}\{\text{Signal} > \text{Mean}\} \times \mathbb{1}\{\text{Mandarin}\}$	0.055** (0.024)	–	0.047** (0.023)	–	–	–	–	–
$\mathbb{1}\{\text{Signal} > \text{Mean}\} \times \text{Experimental Intelligibility}$	–	0.207** (0.100)	–	0.157* (0.088)	–	–	–	–
$\mathbb{1}\{\text{Signal} > \text{Median}\} \times \mathbb{1}\{\text{Mandarin}\}$	–	–	–	–	0.038** (0.018)	–	–	–
$\mathbb{1}\{\text{Signal} > \text{Median}\} \times \text{Experimental Intelligibility}$	–	–	–	–	–	0.156** (0.076)	–	–
$\mathbb{1}\{\text{Signal} > 0\} \times \mathbb{1}\{\text{Mandarin}\}$	–	–	–	–	–	–	0.022* (0.013)	–
$\mathbb{1}\{\text{Signal} > 0\} \times \text{Experimental Intelligibility}$	–	–	–	–	–	–	–	0.147** (0.059)
Observations	1415	1403	1985	1969	1985	1969	1985	1969
R^2	0.261	0.266	0.287	0.293	0.227	0.227	0.226	0.225

Note: This table replicates the analysis from Table 3.3 for the different sample specified and alternative definition of treatment definition. In column (1) & (2), we censor the sample by dropping zero death counties and outlier counties (defined as those with above 99th percentile in outcome variable). In column (3) & (4), we replace extreme values with 99th percentile value. In the remaining columns, results with different signal thresholds are shown. Results are qualitatively similar.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 3A.3. Send Down 5-year Estimates

	(1) Send Down
Birth cohort=1926 × Mandarin speaker=0	0.000 (.)
Birth cohort=1926 × Mandarin speaker=1	-0.008 (0.048)
Birth cohort=1931 × Mandarin speaker=0	0.024 (0.019)
Birth cohort=1931 × Mandarin speaker=1	0.005 (0.032)
Birth cohort=1936 × Mandarin speaker=0	-0.007 (0.018)
Birth cohort=1936 × Mandarin speaker=1	-0.018 (0.025)
Birth cohort=1941 × Mandarin speaker=0	-0.002 (0.018)
Birth cohort=1941 × Mandarin speaker=1	-0.002 (0.023)
Birth cohort=1946 × Mandarin speaker=0	0.006 (0.018)
Birth cohort=1946 × Mandarin speaker=1	0.019 (0.021)
Birth cohort=1951 × Mandarin speaker=0	0.016 (0.018)
Birth cohort=1951 × Mandarin speaker=1	0.071*** (0.020)
Birth cohort=1956 × Mandarin speaker=0	0.008 (0.018)
Birth cohort=1956 × Mandarin speaker=1	0.057*** (0.020)
Birth cohort=1961 × Mandarin speaker=0	-0.017 (0.018)
Birth cohort=1961 × Mandarin speaker=1	-0.042** (0.020)
Birth cohort=1966 × Mandarin speaker=0	-0.015 (0.018)
Birth cohort=1966 × Mandarin speaker=1	-0.023 (0.020)
Birth cohort=1970 × Mandarin speaker=0	-0.013 (0.018)
Birth cohort=1970 × Mandarin speaker=1	-0.018 (0.020)
Birth cohort=1975 × Mandarin speaker=0	-0.020 (0.019)
Birth cohort=1975 × Mandarin speaker=1	-0.035* (0.020)
Birth cohort=1980 × Mandarin speaker=0	-0.023 (0.019)
Birth cohort=1980 × Mandarin speaker=1	-0.033 (0.020)
Birth cohort=1986 × Mandarin speaker=0	-0.025 (0.019)
Birth cohort=1986 × Mandarin speaker=1	-0.031 (0.021)
Birth cohort=1991 × Mandarin speaker=0	-0.023 (0.020)
Birth cohort=1991 × Mandarin speaker=1	-0.030 (0.021)
Observations	16487
R^2	0.277
F	


Note: This table shows the estimates and standard errors from Figure 3.5. The participation rate for non-Mandarin speakers in 1921 is normalized to 0. The outcome is a dummy for participation in the Send Down Movement. The regression includes a full set of province and county fixed effects. Additional controls include parental occupation, parental income, migration history, hukou status, parental hukou status, ethnicity, parent political party affiliations, and interactions between birth cohort and Han ethnicity. Standard errors are clustered on the last known county of residence before the Cultural Revolution, and county of birth if born after the Cultural Revolution. sym* $p < .10$, ** $p < .05$, *** $p < .01$

Table 3A.4. Falsification Test Individual Behavior

	(1)	(2)
	Communist	Sent Down
Panel A: Born After CR		
Mandarin x Birth Cohort Born After CR	-0.027 (0.071)	-0.034 (0.023)
Observations	13350	13350
R^2	0.217	0.292
Panel B: Age 0 to 9 at Start of CR		
Mandarin x Birth Cohort 0-9	0.005 (0.013)	-0.004 (0.008)
Observations	13350	13350
R^2	0.218	0.293
Panel B: Age 10 to 21 at Start of CR		
Mandarin x Birth Cohort 10-21	0.058*** (0.019)	0.043** (0.019)
Observations	13350	13350
R^2	0.218	0.296
Panel C: Age 22 to 32 at Start of CR		
Mandarin x Birth Cohort 22-32	-0.005 (0.037)	-0.000 (0.014)
Observations	13350	13350
R^2	0.217	0.293
Panel D: Age 33+ at Start of CR		
Mandarin x Birth Cohort 33+	-0.027 (0.071)	-0.034 (0.023)
Observations	13350	13350
R^2	0.217	0.292

Note: Each row represents a different regression estimation of Equation 3.2 using different age cohorts at the start of the Cultural Revolution, including the 10-21 year old Cultural Revolution Age Cohort. All regressions include province and county fixed effects. Controls include gender, education, age, age squared, father's education, mother's education, father's political party, mother's political party, father's occupation, mother's occupation, birth county, urban area of residence dummy, father's birth year, mother's birth year, own birth year, ethnicity, and parents' hukou status. Standard errors are clustered on the last known county of residence before the Cultural Revolution, and county of birth if born after the Cultural Revolution. * $p < .10$, ** $p < .05$, *** $p < .01$

Figure 3A.1. Communist Party Membership: Coefficients of Mandarin x Age Cohort



Note: This figure plots regression adjusted trends in Communist Party participation between Mandarin and non-Mandarin speakers. The outcome is a dummy for joining the Communist party within 45 years of birth. The rate for non Mandarin speakers in 1921 is normalized to 0. Controls include a county level fixed effect, and individual level controls including gender, education, age, age squared, father's education, mother's education, father's political party, mother's political party, father's occupation, mother's occupation, birth county, urban area of residence dummy, father's birth year, mother's birth year, own birth year, ethnicity, and parents' hukou status. Additional controls include interactions between birth cohort and Han ethnicity. Standard errors are clustered on the last known county of residence before the Cultural Revolution, and county of birth if born after the Cultural Revolution.