

ESSAYS IN MACROECONOMICS

By

Oscar O'Flaherty

Dissertation

Submitted to the Faculty of the  
Graduate School of Vanderbilt University  
in partial fulfillment of the requirements  
for the degree of

DOCTOR OF PHILOSOPHY

in

Economics

August 12, 2022

Nashville, Tennessee

Approved:

Mario Crucini, Ph.D.

Gregory Huffman, Ph.D.

Peter Rousseau, Ph.D.

Craig Lewis, Ph.D.

*Whatever you do in this life, it's not legendary unless your friends are there to see it. - Barney Stinson*

To my family and friends.

## ACKNOWLEDGMENTS

I have been lucky enough to be surrounded by many great people while completing my dissertation. Although it's not possible to list everyone who inspired me along the way, I am thankful for the constant love and support I received that made the journey both possible and memorable.

Vanderbilt provided me with many wonderful mentors. I am especially thankful to my advisor, Mario Crucini, not only for his support and guidance throughout the program, but most importantly for continually inspiring me as a young economist. His passion for economics frequently reminded me of the importance of having fun as a researcher. I would also like to thank my committee members Greg Huffman, Peter Rousseau, and Craig Lewis for their helpful insights and comments along the way. Special thanks to my co-author, Ben Eden, for his guidance and sharing his interesting life stories.

I am extremely grateful for the friendships I have made over the last six years. In particular, I thank Thu Tran and Elijah Coleman for being my grad school parents. Graduate school would not have been the same without their meals, the countless nights slept on their couch, and talks about my relationship problems. I thank my cohort members for being great colleagues and making grad school enjoyable, especially Craig Sylvera, Martin Schmitz, and Matt Pesner for the nights spent shooting pool and throwing darts. I also thank the elder grad students: Frank Ciarliero, Nicolas Mäder, Gracie Hao, and Trang Hoang for their continued friendship and for checking in throughout the job market. Lastly, I am grateful for my graduate school son, Andrew Rollins, for keeping me young at heart.

Thanks to my Ohio friends: Stevie Zalis, Mike Wayland, Sabrina Pieffer, Jake MacDonald, and Jonathan Fahl for keeping life interesting.

Above all, I thank my family for continually being supportive of me throughout life. The PhD journey was as new an experience for them as it was for me, and I can officially say that I “finished my homework.”

## TABLE OF CONTENTS

	Page
<b>LIST OF TABLES</b> . . . . .	<b>vii</b>
<b>LIST OF FIGURES</b> . . . . .	<b>viii</b>
<b>1 Stay-at-Home Orders in a Fiscal Union</b> . . . . .	<b>1</b>
1.1 Introduction . . . . .	1
1.2 Empirical Framework: Isolating the economic impact of stay-at-home orders . . . . .	4
1.2.1 Variable Definitions . . . . .	4
1.2.1.1 Baseline Labor Variables . . . . .	4
1.2.1.2 Supplemental Labor Variables . . . . .	5
1.2.1.3 Expenditure Variables . . . . .	5
1.2.1.4 Infection variables . . . . .	6
1.2.2 Context and timeline . . . . .	6
1.2.3 Event Study Specification . . . . .	9
1.3 Results: Effect of Stay-at-Home Orders on the Economy . . . . .	10
1.3.1 Opportunity Insights Employment Measure . . . . .	12
1.3.2 Heterogeneity across Sectors and the Anticipation Effect . . . . .	14
1.3.3 Identifying Variation induced by Treatment Timing . . . . .	15
1.3.3.1 Refined Event Study Specification . . . . .	17
1.3.3.2 Results . . . . .	17
1.3.4 Non-essential Business Closures . . . . .	17
1.3.5 Interpretation of Event Study Results . . . . .	19
1.4 Model Framework: Optimal mitigation policy with multiple locations . . . . .	21
1.4.1 Initial Conditions . . . . .	21
1.4.2 Virus Transmission . . . . .	22
1.4.3 The Economy . . . . .	23
1.4.3.1 Households . . . . .	23
1.4.3.2 Firms . . . . .	24
1.4.3.3 Government . . . . .	24
1.5 Model Implications: Multiple locations and optimal policy . . . . .	24
1.5.1 Multiple-location model vs single-location model . . . . .	25
1.5.2 Application to US Census Region Data . . . . .	27
1.5.2.1 Calibration Procedure . . . . .	27
1.5.2.2 Interpretation of Transmission Matrix . . . . .	28
1.5.2.3 Results . . . . .	29
1.5.2.4 Comparative Statics . . . . .	30
1.5.3 Optimal Mitigation Policy . . . . .	31
1.5.3.1 Planner’s Problem . . . . .	31
1.5.3.2 Results . . . . .	32
1.6 Conclusion . . . . .	32
1.7 Appendix . . . . .	34
1.7.1 Data . . . . .	34
1.7.1.1 Summary Statistics . . . . .	34
1.7.1.2 Difference-in-difference Estimates . . . . .	38
1.7.1.3 Additional Event Study Estimates . . . . .	40
1.7.2 Model . . . . .	43

1.7.2.1	Equilibrium Conditions . . . . .	43
1.7.2.2	Parameters chosen to Optimize Model . . . . .	44
1.7.2.3	Extended Model Calibration Procedure . . . . .	45
1.7.2.4	Additional Comparative Statics . . . . .	46
1.7.2.5	National Mitigation Policy . . . . .	49
<b>2</b>	<b>Selection Effects in Retail Chain Pricing . . . . .</b>	<b>50</b>
2.1	Introduction . . . . .	50
2.2	Data . . . . .	53
2.2.1	Sample Formation . . . . .	53
2.2.1.1	Stores . . . . .	53
2.2.1.2	Products . . . . .	54
2.2.2	Retail Chain Pricing . . . . .	55
2.3	Price Synchronization within Retail Chains . . . . .	56
2.3.1	Price Synchronization . . . . .	56
2.3.2	Variance Decomposition . . . . .	57
2.3.2.1	Robustness Checks . . . . .	58
2.4	Menu Cost Model with Retail Chains . . . . .	59
2.4.1	Standard Model . . . . .	59
2.4.2	Retail Chain Extension . . . . .	60
2.4.3	Calibration . . . . .	61
2.4.4	Model Intuition . . . . .	62
2.5	Selection Effects in Retail Chain Pricing . . . . .	63
2.5.1	Quantifying Selection Effects . . . . .	64
2.5.1.1	Specification . . . . .	64
2.5.1.2	Intuition . . . . .	64
2.5.2	Results . . . . .	65
2.5.3	CalvoPlus Model Extension . . . . .	66
2.6	Conclusion . . . . .	67
2.7	Appendix . . . . .	68
2.7.1	Empirics . . . . .	68
2.7.1.1	Variance Decomposition . . . . .	69
2.7.2	Model . . . . .	71
<b>3</b>	<b>Temporary Sales in Response to Demand Shocks . . . . .</b>	<b>72</b>
3.1	Introduction . . . . .	72
3.2	Data . . . . .	74
3.2.1	Filtering Process . . . . .	74
3.2.2	Temporary Sales . . . . .	76
3.3	Stylized Facts . . . . .	77
3.3.1	Do stores use a mixed strategy? . . . . .	78
3.3.2	Temporary sales and demand uncertainty . . . . .	79
3.3.2.1	Empirical Strategy . . . . .	80
3.3.2.1.1	Variable Definitions . . . . .	80
3.3.2.1.2	Specification . . . . .	81
3.3.2.2	Results . . . . .	81
3.3.3	Temporary sales and demand shocks . . . . .	83
3.3.3.1	Empirical Strategy . . . . .	83
3.3.3.1.1	Specification . . . . .	83
3.3.3.1.2	Interpretation . . . . .	83
3.3.3.2	Results . . . . .	84
3.4	Shocks to the Aggregated Demand for a Product . . . . .	86
3.4.1	Specification . . . . .	86

3.4.2	Sales Response . . . . .	87
3.4.3	Contribution to Price Dynamics . . . . .	87
3.5	A Model . . . . .	89
3.5.1	Uncertain and Sequential Trade Model . . . . .	89
3.5.1.1	Uncertain . . . . .	89
3.5.1.2	Sequential . . . . .	89
3.5.1.3	Equilibrium . . . . .	90
3.5.2	Extended Model with Temporary Sales . . . . .	91
3.5.2.1	Equilibrium . . . . .	92
3.5.3	Accounting for the stylized facts . . . . .	94
3.6	Conclusion . . . . .	95
3.7	Appendix . . . . .	96
3.7.1	Empirics . . . . .	96
<b>References</b>	. . . . .	<b>98</b>

## LIST OF TABLES

Table	Page	
1.1	Timeline of Stay-at-Home Orders . . . . .	7
1.2	Effect of Stay-at-Home Orders on the Labor Market and Expenditures . . . . .	11
1.3	Cumulative and Relative Effects of Stay-at-Home Orders . . . . .	20
1.4	Baseline Model Parameters . . . . .	25
1.5	Initial Conditions for Census Region Application (May 6) . . . . .	28
1.6	US Census Region Data vs Model Prediction (August 26) . . . . .	31
1.7	Balance across states . . . . .	36
1.8	Difference-in-difference Estimates . . . . .	38
1.9	Effect of Stay-at-Home Orders on the Economy (Composition Effects) . . . . .	40
1.10	Opportunity Insights Employment Measure (State vs County) . . . . .	41
1.11	Effect of Non-essential Business Closures on the Economy . . . . .	43
2.1	Summary Statistics of IRI Data . . . . .	54
2.2	Synchronization of Price Changes . . . . .	56
2.3	Benchmark Parameters . . . . .	61
2.4	Selection Effects in Retail Chain Pricing . . . . .	66
2.5	CalvoPlus Parameters . . . . .	71
3.1	Summary Statistics of IRI Data . . . . .	75
3.2	Summary Statistics for Sales . . . . .	76
3.3	Behavior of Temporary Sales . . . . .	79
3.4	Demand Uncertainty and Temporary Sales . . . . .	82
3.5	Impulse Response of Temporary Sales to Demand Shocks . . . . .	85
3.6	Impulse Response of Average Prices to Demand Shocks . . . . .	88
3.7	Demand Uncertainty and Temporary Sales (Regressions that address Potential Endogeneity Issues) . . . . .	96
3.8	Aggregate Impulse Response of Temporary Sales to Demand Shocks . . . . .	97

## LIST OF FIGURES

Figure	Page	
1.1	Coronavirus outbreak by US state . . . . .	6
1.2	Economic Response to Covid-19 . . . . .	8
1.3	Effect of Stay-at-Home Orders on the Economy . . . . .	10
1.4	Opportunity Insights Employment Measure . . . . .	13
1.5	Heterogeneity across Consumption Expenditures . . . . .	15
1.6	Effect of Treatment Timing on Hours Worked DD Estimate . . . . .	16
1.7	Effect of Treatment Timing on Event Study Estimates . . . . .	18
1.8	Effect of Non-essential Business Closures on the Economy . . . . .	19
1.9	Multiple-location Model vs Single-location Model . . . . .	26
1.10	US Census Region Data vs Model (Susceptibles) . . . . .	30
1.11	Optimal Mitigation Policy . . . . .	33
1.12	Heterogeneous Response to Covid-19 by Income and Sector . . . . .	34
1.13	Timeline of Stay-at-Home Orders . . . . .	35
1.14	Counties by Treatment Status . . . . .	37
1.15	Goodman-Bacon (2018) Decomposition . . . . .	39
1.16	Opportunity Insights Employment Measure by NAICS Sector . . . . .	42
1.17	Heterogeneity across Consumption Expenditures (Extra) . . . . .	42
1.18	Effect of Stay-at-Home Orders on the Economy (Abraham and Sun) . . . . .	43
1.19	Extended US Census Region Application . . . . .	46
1.20	Comparative Statics (Susceptibles) . . . . .	47
1.21	Optimal National Policy . . . . .	49
2.1	Example of Retail Chain Price Synchronization . . . . .	55
2.2	Variance Decomposition . . . . .	58
2.3	Example Pricing Decision . . . . .	63
2.4	Selection Effects with Retail Chain Pricing . . . . .	65
2.5	Retail Chain Price Synchronization . . . . .	68
2.6	Variance Decomposition (Robustness Checks) . . . . .	70
3.1	Variation in Sales Frequency . . . . .	78
3.2	Demand Uncertainty and Temporary Sales . . . . .	81
3.3	Impulse Response of Temporary Sales to Demand Shocks . . . . .	84
3.4	Impulse Response of Aggregate Sales Frequency to Demand Shocks . . . . .	86
3.5	Impulse Response of the Average Price to Demand Shocks . . . . .	87
3.6	Prices and Quantities in the Baseline UST Model . . . . .	90
3.7	Possible Equilibria in the UST Model with Storage . . . . .	93



## CHAPTER 1

### Stay-at-Home Orders in a Fiscal Union

#### 1.1 Introduction

State and local governments throughout the United States attempted to mitigate the outbreak of Covid-19 using stay-at-home orders to limit social interactions and mobility. Many praised these policies for slowing the spread of the virus while others protested them due to their perceived negative impact on the economy. The implementation of stay-at-home orders also created conflict between policy makers. Some called for a nationwide order while others believed discretion should be left to state and local governments.

This paper studies the economic and political controversy associated with stay-at-home orders and their implementation, both empirically and theoretically. Empirically, we use an event-study framework to estimate the marginal impact of stay-at-home orders on the economy. We build a daily, state-level panel of labor market and consumer spending variables using data from the time-clock company Homebase and non-profit organization Opportunity Insights. Our main sample spans several weeks before the issuance of the first stay-at-home order in March up to the first reopening in April. Both hours worked and consumer spending were already declining before the first statewide stay-at-home order was issued on March 19th. Thus, the economic decline was likely attributed to a combination of government-issued policies and behavioral responses by consumers and firms to the health risks posed by the pandemic. The event-study framework allows us to separate these economic effects by exploiting differences across states that issued a stay-at-home order and those that did not. Controls for the daily, case load of infections in each state and non-essential business closures help isolate the impact of stay-at-home orders. The control for non-essential closures also helps approximate the cumulative effect of mitigation policy.

Stay-at-home orders negatively impacted the labor market. Hours worked initially declines by 2 percentage points and continues downward to a trough of 6 percentage points (pp) two weeks after the order is enacted. High-income workers experience larger declines in employment compared to low-income workers. Early-treated states experience about twice the decline in hours worked (4pp) relative to late-treated states (2pp). Non-essential business closures did not exacerbate or mitigate the labor market response.

Stay-at-home orders also decreased consumer spending by 4 percentage points for several weeks. In contrast to employment, spending declines are larger in magnitude among low-income consumers compared to high-income consumers. Consumer anticipation effects vary across sectors. For example, we find evidence of stockpiling in grocery and food stores, but not in accommodation and food services. Early and late-treated

states experience similar post-treatment declines in spending. However, stockpiling in the pre-treatment period is primarily driven by early-treated states. Unlike employment, non-essential business closures did exacerbate the consumer spending response.

The baseline estimates suggest that stay-at-home orders decreased consumer spending by \$10 billion and total earnings by \$15 billion. They also suggest that stay-at-home orders account for 10%-25% of the total employment decline observed during our sample. The non-essential business closure estimates suggest that the cumulative effect of mitigation policy is at least as large in magnitude as implied by the stay-at-home order estimates. The results are not driven by outliers. Lastly, in response to the virus, some counties implemented their own stay-at-home orders before their respective states. We also conduct our analysis at the county-level to account for these orders and find similar impacts compared to our baseline estimates.

In order to study the optimal implementation of stay-at-home orders, we extend the Susceptible-Infected-Recovered (SIR) model of Eichenbaum et al. (2020) to multiple locations. Each location initially experiences an idiosyncratic virus shock. Following this, the virus can spread through both consumption and employment *within* a given location. This interaction captures the natural societal responses to the virus as people reduce their economic activity in order to limit their risk of infection. Although locations are closed economies with respect to their economic transactions, individuals may also spread the virus *across* locations by travelling from place to place. Thus, the economic activity of one location is dependent upon other locations due to virus spillovers. The model incorporates travel cost across locations in order to specify the extent of virus spillovers. Lastly, in response to the transmission of the initial virus shocks, an aggregate social planner can implement mitigation policy—proxied by a consumption tax—to maximize the welfare of the country. Mitigation policy is able, but not required, to vary across locations. This allows for economic and welfare comparisons between a common national policy and location-specific policy.

The main contribution of our model is the cross-sectional distribution of consumption, hours worked, and infection rates that it generates. The single-location model of Eichenbaum et al. (2020) is equivalent to our multiple-location model if and only if the initial infection rates of locations are symmetric and travel costs across locations are zero. If either of these restrictions fails to hold, then aggregate dynamics will depend on the behavior of the cross-section.

We explore the behavior of the cross-sectional distribution and dynamics in the model through an application using US Census region data. The model suggests that travel cost across locations plays a significant role in the spread of the virus in the US. Allowing for complete pass-through of virus transmission overestimates the infection rates of every Census region. In contrast, allowing for no virus transmission across regions tends to underestimate the infection rates seen in the data. We propose a method to determine the rate of transmission across locations through the associated travel costs. This method provides a good fit of both the

cross-sectional distribution and dynamics in the US Census data, and the procedure is generalized for use in other applications.

Lastly, we turn our analysis to optimal policy with multiple locations. From a national welfare perspective, it is optimal for locations with higher infection rates to set stricter mitigation policies. A common, national policy is too restrictive for mildly infected locations and causes greater consumption and employment declines than are optimal in these areas. Stricter policy in more severely affected locations causes infection rates, and subsequently mitigation policy, to converge over time. Thus, our finding that policy should be set at a local level is contingent on sufficient asymmetries in the infection rates of locations.

Other research has examined the empirical relationship between the economy, the spread of Covid-19, and the ensuing policy responses (Baker et al., 2020; Binder, 2020; Chetty et al., 2020; Rojas et al., 2020). Our empirical contribution is closely related to other difference-in-difference estimates that find negative effects of stay-at-home orders on the labor market (Allcott et al., 2020; Baek et al., 2020; Bartik et al., 2020; Beland et al., 2020; Gupta et al., 2020; Kong and Prinz, 2020; Lin and Meissner, 2020) and consumer spending (Allcott et al., 2020; Coibion et al., 2020; Goolsbee and Syverson, 2020). Importantly, our estimates for both the employment and spending effects are found using an event study specification using data at a daily frequency which builds on many of these papers. The high frequency data allows us to provide a clear characterization of both when and where stay-at-home orders were issued. The event study specification allows us to assess whether our results are driven by differences in observable pre-treatment trends across treated and control states, for which we do not find evidence. Of the papers mentioned above, our paper is most closely related to Allcott et al. (2020) who find similar results for both consumer spending and hours worked. The Opportunity Insights data allows us to analyze these effects by consumer and worker income quartiles as well as by industry, which are both new results.

Our paper is also related to the quantitative theoretical literature of infectious disease modelling (Kermack and McKendrick, 1927; Angulo et al., 2013; Zakary et al., 2017; Bisin and Moro, 2020) and its application to economic activity and mitigation policy (Acemoglu et al., 2021; Atkeson, 2020; Jones et al., 2020). We extend the model of Eichenbaum et al. (2020) to multiple locations. This allows us to illustrate how the economy in one location can be affected by the spread of the virus in other locations. Antràs et al. (2020), Bognanni et al. (2020), and Fajgelbaum et al. (2020) also study the economic spillovers of virus transmission across locations. Antràs et al. (2020) focus primarily on endogenous trade adjustments while Fajgelbaum et al. (2020) assess cross-location mitigation policy such as travel restrictions. Similar to our paper, Bognanni et al. (2020) analyze domestic mitigation policies such as stay-at-home orders. However, we focus specifically on the optimality of location-specific versus nationwide mitigation policy. Overall, these models provide complementary, rather than contradictory, results which help to more fully characterize the economic effects

of virus transmission across locations and the subsequent policy responses. Lastly, Baek et al. (2020) develop a model with stay-at-home orders characterized as an exogenous productivity or demand shock. By explicitly modelling the spread of the virus, we are able to endogenously model mitigation policy such as stay-at-home orders.<sup>1</sup> This allows us to evaluate their economic impact as well as their optimal implementation in response to the spread of the virus.

The rest of the paper proceeds as follows. Section 1.2 presents a timeline of the spread of Covid-19 throughout the U.S., the subsequent economic and policy responses, and formalizes the event study specification. Section 1.3 presents the baseline event study results and robustness checks. Section 1.4 describes the economic SIR model. Section 1.5 compares the multiple-location model with the single-location model, applies the multiple-location model to Census region data, and presents the results for optimal mitigation policy. Section 1.6 concludes.

## **1.2 Empirical Framework: Isolating the economic impact of stay-at-home orders**

We begin by creating a daily panel of several labor market and consumer spending variables. Our panel spans from March up to April 24, the first day a stay-at-home order was lifted. This allows our analysis to focus specifically on the initial virus outbreak and asymmetric policy responses across states. The second subsection provides a timeline of the outbreak of the virus, the enactment of stay-at-home orders, and the economic response during this period. The third subsection formalizes our event-study specification.

### **1.2.1 Variable Definitions**

#### **1.2.1.1 Baseline Labor Variables**

We obtain our primary labor market measures from the time-clock company Homebase.<sup>2</sup> Homebase helps over 60,000 businesses track the hours worked of 1 million hourly employees. Their data provides three labor-market measures: businesses open, employees working, and hours worked. A business is defined as open if at least one employee clocked in for a given day. Employees working is counted as the number of unique employees who clocked in at least once. Hours worked is then determined from the time cards of all employees. Total hours worked is the primary outcome variable as it corresponds to the model developed in Section 1.4. All variables are measured at the state level. Each variable is normalized relative to the median for that day of week using data from January 4, 2020 through January 31, 2020.

---

<sup>1</sup>This is one of the novel contributions of Eichenbaum et al. (2020) which remains in our extension.

<sup>2</sup>An interactive database from Homebase can be found here: <https://joinhomebase.com/data/>

### 1.2.1.2 Supplemental Labor Variables

We supplement our Homebase data with an employment measure constructed by the non-profit organization Opportunity Insights (OI) in Chetty et al. (2020). Employment is normalized relative to the average level from Jan. 4-31. This employment data provides several advantages and disadvantages compared to the Homebase data. The first advantage is that the OI data is available at both the state and county level. This allows us to account for stay-at-home orders issued by individual counties. Second, employment is measured by low (<\$27,000 per year), medium (\$27,000-\$60,000), and high income (>\$60,000) earners at both the state and county level.<sup>3</sup> Lastly, the state-level data is disaggregated into four NAICS sectors: Professional and Business Services, Education and Health Services, Retail and Transportation, Leisure and Hospitality.

However, the OI data has several disadvantages. The OI data is constructed using worker-level data from *Earnin* and firm-level payroll data from *Paychex* and *Intuit*.<sup>4</sup> This creates the conflict that observations are recorded when workers are paid rather than the specific date they are employed as in the time-clock data from Homebase. To circumvent this issue, Chetty et al. (2020) make the assumption that individuals are employed for all days in their pay period. If a stay-at-home order occurred in the middle of an individual's pay period, employment effects would not be seen until their next pay period. Thus, the implications of this assumption are too strong for an event study analysis which hinges upon variation at the time of treatment. To account for the difference between the end of the pay period and the date of employment, we lag the data by an additional week.<sup>5</sup> Lastly, their employment series is constructed as a seven-day look-back moving average. This attenuates the overall effect of a one-day change in employment which is of primary interest in the context of stay-at-home orders. For these reasons, we use the Homebase data for our baseline estimates on the labor market and the OI data for comparisons and robustness checks.<sup>6</sup>

### 1.2.1.3 Expenditure Variables

We use credit and debit card spending data from Affinity Solutions to estimate changes in consumer expenditures. This expenditure data is also constructed by Chetty et al. (2020) and is available at more granular levels than the Homebase data. At the state level, expenditure data is provided by income quartile. Consumers are considered high-income if they live in a ZIP code with median income in the top quartile of ZIP codes. Similarly, low-income consumers live in the bottom quartile of ZIP codes. Middle-income consumers

---


<sup>3</sup>These are approximate measures. Quartiles are based upon the hourly wage distribution in 2019. Income levels are more accurately < \$13.00/Hour, \$13.00 – \$29.17, and > \$29.17/Hour.

<sup>4</sup>Please see Chetty et al. (2020) for a more detailed description of their employment series.

<sup>5</sup>Chetty et al. (2020) use *Earnin* data only for workers that are paid weekly or bi-weekly. They make the assumption that there is a one-week lag from the end of the pay period to the date an individual is paid. By lagging their employment series an additional week, we essentially shift the data to the beginning of the two-week pay cycle and assume there is no lag between the end of the pay period and payment date. Chetty et al. (2020) also convert weekly employment data from *Paychex* to daily by assuming that employment is constant within each week. Lagging an additional week also helps control for the effects of this assumption.

<sup>6</sup>Appendix Figure 1.12 plots these two measures for comparison.

Figure 1.1: Coronavirus outbreak by US state



Note: This figure plots the natural log of cases per 1,000 residents for each state. To account for periods of zero infection in the beginning of our sample, we add one to the number of cases for all observations before calculating the infection rate. The black line represents the national total. Data is obtained from USAFacts.

are then defined as living in the second and third quartiles. Expenditure data is also separated into several merchant groups defined by Affinity Solutions: Apparel and General Merchandise, Entertainment and Recreation, Grocery, Health Care, Restaurants and Hotels, Transportation. Total expenditure data is also available at the county level, but is not available by income quartile or merchant group. A seven-day look-back moving average is taken for all consumption variables, and then measured relative to the median for that day of week using data from January 4, 2020 through January 31, 2020.<sup>7</sup>

#### 1.2.1.4 Infection variables

Our daily panel spanning this period uses cases reported by the non-profit organization *USA Facts*.<sup>8</sup> Cases are reported at the county level which we then aggregate to the state level. Primary source information for infections and tests comes from the CDC and state and local health agencies. Population data are obtained from the US Census to calculate case load per capita.

### 1.2.2 Context and timeline

The United States recorded its first case of the novel coronavirus disease of 2019 (Covid-19) on January 20th, 2020 when a man from the state of Washington became ill after a visit to Wuhan, China. The timing and spread of the virus differed across states after this initial infection. To see this, Figure 1.1 plots the natural log of cases per capita for each state from March 9th through April 24th. West Virginia was the last state

<sup>7</sup>See Chetty et al. (2020) for a full description of the construction of the expenditure series.

<sup>8</sup>Specifically, cases and deaths are obtained from the website: <https://usafacts.org/visualizations/coronavirus-covid-19-spread-map/>.

Table 1.1: Timeline of Stay-at-Home Orders

March 15	March 16	March 17	March 18	March 19	March 20	March 21
				California		
March 22	March 23	March 24	March 25	March 26	March 27	March 28
Illinois New Jersey	New York Oregon Washington	Connecticut Delaware Louisiana Massachusetts Michigan New Mexico Ohio	Hawaii Idaho Indiana West Virginia Wisconsin	Colorado Vermont	Kentucky	Minnesota Montana New Hampshire Rhode Island
March 29	March 30	March 31	April 1	April 2	April 3	April 4
Alaska	Kansas Virginia	North Carolina Maryland	Arizona District of Columbia Nevada Tennessee	Maine Pennsylvania Texas	Florida Georgia	Mississippi
April 5	April 6	April 7	April 8	April 9	April 10	April 11
Alabama	Missouri		South Carolina			

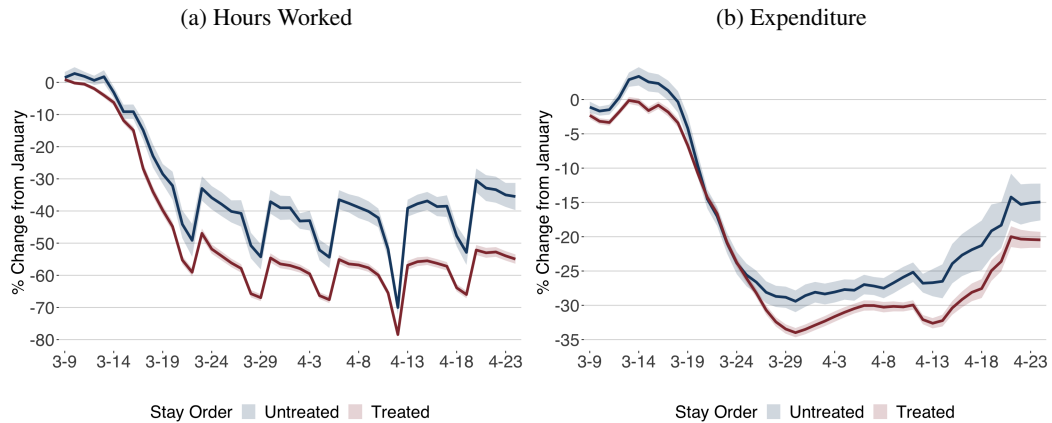
Note: This table presents the timeline of the enforcement of stay-at-home orders. Some policies began at a specific time of the day such as 6:00am or 6:00pm. Orders that occur after noon are recorded on the following day. Dates of stay-at-home orders are obtained from the *New York Times*.

to document a confirmed case on March 17th, almost two months after the first case in Washington. At this point, Montana had a case rate similar to the national total. However, low case growth led them to have the least cases per capita on April 24th. In contrast, persistently high case growth in New York led them to maintain the highest case rate per capita from mid-March through the end of April.

As infection rates grew, many local and state governments issued various social distancing policies to help limit personal contact and combat the spread of the coronavirus. Our analysis centers on statewide stay-at-home orders. These orders generally closed all non-essential businesses, and encouraged people to stay at home except for essential activities such as shopping for food or seeking medical care. Over 90% of the US population was placed under a stay-at-home order in the span of three weeks across forty-two states and the District of Columbia. Table 1.1 documents the timeline of these stay-at-home orders.<sup>9</sup> California was the first state to announce an order on March 19th. Twenty-five additional states issued orders the following week. South Carolina was the last state to issue an order on April 8. Some orders were issued to start at a specific time of the day (e.g. 6:00am or 6:00pm). As our panel is at a daily frequency and cannot capture

<sup>9</sup>Appendix Figure 1.13 presents a timeline of the percentage of the US population under a stay-at-home order.

Figure 1.2: Economic Response to Covid-19



Note: This figure plots the time series of spending and hours worked for states which issued a stay-at-home order (red) and states that did not (blue) from March 9th to April 24th. Estimates are averages over states in their respective treatment groups. Shaded regions represent the standard error for that date. Hours and expenditure data are provided by Homebase and Affinity Solutions, respectively.

this, we consider states with orders enacted after Noon to start treatment the following day.<sup>10</sup>

Despite widespread recognition of their need from a public health perspective, stay-at-home orders were met with opposition. Protests over the economic impact of these orders occurred in many states (NBC News, 2020; Fox News, 2020). One protester in New Jersey said, “Businesses are suffering, unemployment checks are not being sent, landlords are not getting rent. We feel like these directives are causing more suffering than is necessary.” (NJ.com, 2020) A poll conducted in May estimated that 35% of people “strongly or somewhat agreed” that restrictions and closures had been too severe (USA Today, 2020).

Figure 1.2 plots the the response of hours worked and consumer expenditures during the outbreak of Covid-19 from March 9th to April 24th. The red line represents the average over states which issued a stay-at-home order, and the blue line depicts those that did not. Standard errors are shaded in their respective colors. Both panels show that employment and expenditure were declining before the first statewide stay-at-home order was issued on March 19th. We also see that these declines differed across treatment groups throughout most of the sample.

Panel (a) shows that total hours worked in states which issued a stay-at-home order declined by 34% on March 18th. On this date, states that did not issue an order only had a decline of 23%. As a stay-at-home order had not yet been issued, they cannot account for this economic decline or the 11 percentage point (pp) difference across groups. However, these differences did grow over time. On April 23rd, the day before the first lifting of a stay-at-home order, declines in hours worked were 54% and 35% for the treated and untreated groups, respectively. Thus, the differences in hours worked across groups had grown

<sup>10</sup>Primary source information for state announcements was compiled by the *New York Times*: <https://www.nytimes.com/interactive/2020/us/coronavirus-stay-at-home-order.html>



by 8 percentage points from March 18th (11pp) to April 23rd (19pp). Panel (b) shows a similar pattern in expenditures. The percentage point difference in expenditure across groups grew by 2 percentage points from on March 18th (3pp) to on April 23rd (5pp).<sup>11</sup>

### 1.2.3 Event Study Specification

The simple analysis above helps provide the intuition behind a difference-in-differences estimation. Although it provides suggestive evidence that stay-at-home orders may have contributed to the economic decline of states, it does not account for differences in characteristics of states as in Figure 1.1 or the variation in treatment timing in Table 1.1. To account for these complications, the following event study specification is estimated:

$$y_{it} = \alpha_i + \gamma_t + \Gamma X_{it} + \sum_{\substack{k \neq -1 \\ -10 \leq k \leq 15}} \beta_k I(\text{StayHome}_{it} = k) + \varepsilon_{it} \quad (1.1)$$

This empirical strategy uses variation in both *when* and *where* stay-at-home orders were issued. Our primary outcome variable,  $y_{it}$ , is either hours worked or consumption expenditures in state  $i$  at time  $t$ . State fixed effects,  $\alpha_i$ , control for both observable and unobservable differences across states that do not change over time. Time fixed effects,  $\gamma_t$ , control for variables that affect all states in a given period such as the weekly seasonality of hours worked seen in Panel (a) of Figure 1.2 or national news about the coronavirus. Our sample begins on March 9th, ten days before the first statewide stay-at-home order in California, and ends on April 23rd, the day before the first statewide reopening in Alaska.  $\text{StayHome}_{it}$  represents the current date relative to the date a stay-at-home order was issued in a given state. This variable is always equal to  $-1$  for non-treated states which serve as the control group.

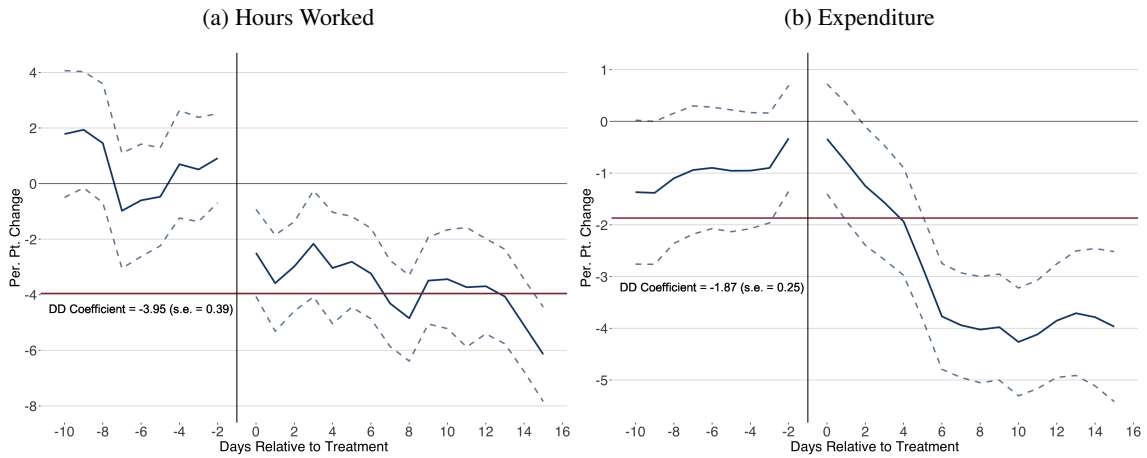
The coefficients of interest, the  $\beta_k$ 's, represent the average difference between the control and treated group at period  $k$  relative to the difference at  $k = -1$ , the day before the order is enacted for treated states. We present estimates for event times  $k \in [-10, 15]$ . Event times 0 to 15 estimate the causal effect of a stay-at-home order on hours worked and consumption expenditures. These estimates represent the average treatment effect for the treated (ATT). Event times -10 to -2 serve as a falsification test that our results are not driven by pre-existing differences in trends between the treatment and control groups before a stay-at-home order is issued. All of our presented estimates use a balanced set of states. This means our coefficients do not reflect changes in the composition of the treatment group.<sup>12</sup>

Goodman-Bacon and Marcus (2020) raise substantive issues and provide recommendations when using a difference-in-differences design in the context of Covid-19. Our baseline specification contains several

<sup>11</sup>Appendix Table 1.7 performs a similar analysis with the remaining labor, expenditure, and infection variables.

<sup>12</sup>For example, if we extended the event times to  $k \in [-10, 16]$ , South Carolina would not be included in the estimate of  $k = 16$  due to the timing of its stay order (recall Table 1). Thus,  $\beta_{16}$  would be partially driven by this change in the composition of the treated group.

Figure 1.3: Effect of Stay-at-Home Orders on the Economy



Note: This figure plots the stay-at-home order event study coefficients estimated from Equation 1.1. Subcaptions represent the dependent variable. Dashed lines represent 95% confidence intervals using robust standard errors. The red line plots the difference-in-difference estimate. Observations are at the state-by-day level. Hours and expenditure data are obtained from Homebase and Affinity Solutions, respectively.

controls to help us address these issues. Namely, we include a control for log cases per capita and a similar event study variable for non-essential business closures.<sup>13</sup> The inclusion of a measure of case incidence helps us to separate the endogenous responses of individuals and firms to the risk of infection from the economic impact of stay-at-home orders. Controlling for non-essential business closures helps separate the effect of these complementary policy tools and avoid possible confounding effects of the two.<sup>14</sup>

### 1.3 Results: Effect of Stay-at-Home Orders on the Economy

Figure 1.3 analyzes the effect of stay-at-home orders on the economy. Panel (a) shows that pre-trends in hours worked are slightly elevated eight to ten days before a stay-at-home order, but are flat 1-7 days before a stay-at-home order is enacted. Differences in trends are statistically insignificant at the 95% level using robust standard errors for the entire pre-treatment period. Hours worked declined by 4 percentage points (s.e. = 0.39) on average after a stay-at-home order. These effects vary over time. Immediately after the order, treated states experience over a 2 percentage point decrease in hours. This effect is persistent with a slight downward trend. Fifteen days after the order, treated states experience a trough decline of  $-6$  percentage points. Effects are statistically significant for the entire post-treatment period.

Panel (b) plots the coefficients for statewide expenditure. Expenditures experience a statistically significant 4pp decline five days after the order. This effect remains stable and significant up to fifteen days after the

<sup>13</sup>Non-essential business closures at the state level are obtained from Raifman et al. (2020).

<sup>14</sup>Our main results analyze the effect of stay-at-home orders on the economy. Non-essential business closure estimates can be found in Section 1.3.4 and Appendix 1.7.1.

Table 1.2: Effect of Stay-at-Home Orders on the Labor Market and Expenditures

Days	Labor Variables (Per. Pt.)			Expenditure Variables (Per. Pt.)				
	Hours Worked (1)	Employment (2)	Businesses Open (3)	State Level (4)	County Level (5)	Low Income (6)	Middle Income (7)	High Income (8)
-7	-0.98 (-0.93)	-0.16 (-0.16)	0.37 (0.38)	-0.94 (-1.49)	-0.91 (-1.35)	0.19 (0.2)	-1.05 (-1.42)	-0.39 (-0.55)
-2	0.91 (1.11)	0.88 (1.19)	0.7 (0.9)	-0.33 (-0.63)	-0.44* (-2.48)	-0.29 (-0.33)	-0.47 (-0.81)	0.13 (0.19)
0	-2.50* (-3.12)	-2.03* (-2.78)	-2.46* (-3.45)	-0.34 (-0.63)	-0.33* (-2.07)	-0.4 (-0.45)	-0.35 (-0.61)	-0.21 (-0.32)
5	-2.81* (-3.38)	-2.59* (-3.53)	-2.62* (-3.55)	-2.83* (-5.53)	-2.72* (-3.89)	-2.41* (-2.9)	-3.12* (-5.52)	-2.52* (-3.92)
10	-3.44* (-3.79)	-3.27* (-4.02)	-4.14* (-5.22)	-4.26* (-8.02)	-3.66* (-3.8)	-4.24* (-4.59)	-4.62* (-8.21)	-3.21* (-4.85)
15	-6.14* (-7.08)	-5.68* (-7.08)	-5.59* (-6.91)	-3.97* (-5.36)	-2.88* (-2.54)	-3.03* (-2.59)	-4.2* (-5.84)	-3.62* (-5.25)
Observations	2,346	2,346	2,346	2,346	73,451	2,254	2,346	2,300

Note: This table presents the stay-at-home order event study coefficients estimated from Equation 1.1. Columns represent the dependent variable. Observations are at the state-by-day level except column (5) which uses county-by-day observations. Counties begin treatment on the date they issued a stay-at-home order or the date the state they reside in issued an order. T-values are presented using robust standard errors for state estimates and standard errors clustered at the state-level for county estimates. Labor and consumption variables come from Homebase and Affinity Solutions, respectively.  
\* denotes estimates significant at the 95% level.

order. We do not find statistically significant differences in pre-trends before a stay-at-home order. However, they do exhibit a 1 percentage point upward trend. It is known that consumers would often stockpile items, such as toilet paper, in the days preceding a stay-at-home order. This upward trend may be representative of the consumer anticipation effect. In this case, our event study estimates for expenditure may be biased towards finding a negative effect and overstate the effect of the order.

Table 1.2 provides several alternative specifications and further explores the extent of the anticipation effect in the baseline results of Figure 1.3.<sup>15</sup> T-values are presented to ease comparison across specifications. The baseline results are provided in columns (1) and (4) for reference. Results for employment and businesses open in columns (2) and (3) test the robustness of the benchmark labor market results. Immediately after the order, we see slightly weaker effects of -2.03pp and -2.46pp in employment and businesses open compared to the baseline estimate of -2.50pp in hours worked. Both variables also suggest a slight downward trend with an estimated decrease of about 5.6pp fifteen days after the order is enacted.

The baseline results represent the average effect of statewide stay-at-home orders on treated states' eco-

<sup>15</sup>Additional robustness checks for population-weighted regressions, outliers, and composition effects can be found in Appendix Table 1.9. These checks have no significant effects on the results.

conomic activity. However, some counties issued local orders before their state did. These policies could bias the estimates towards not finding a result. To account for this, we also estimate Equation 1.1 using county-level orders as the basis for treatment. In this case, a county is counted as treated if it issued a stay-at-home order, or if the state it resides in did.<sup>16</sup> As the Homebase dataset does not contain county-level information, this is only estimated for the consumption data from Affinity Solutions in column (5). Furthermore, we expect spending and saving patterns to be related to one's income level. This heterogeneity may lead to differences in consumption smoothing. We test this hypothesis in columns (6)-(8).<sup>17</sup>

Columns (4) and (5) of Table 1.2 are similar in terms of their coefficient estimates. However, we see that the county estimates are statistically significant at the 95% level two days before and the day the order is enacted.<sup>18</sup> This implies that the anticipation effect mentioned earlier is statistically significant. Middle-income earners in column (7) and our baseline estimates have similar upward pre-trends as evidenced by days -7 and -2. However, these upward trends are not as apparent for low (6) and high (8) income consumers. This suggests that the anticipation effect is mostly driven by middle-income consumers. This is further exemplified as all of the significant point estimates are larger in magnitude for middle-income consumers compared to low and high-income estimates. Despite this, we still find statistically significant estimates for both high and low-income consumers 5-15 days after an order is enacted. Lastly, comparing columns (6) and (8), we see that low-income consumers see a trough marginal decline in expenditure of 4.2pp which occurs ten days after the order. High-income consumers see a smaller trough decline of 3.6pp fifteen days after the order.

### 1.3.1 Opportunity Insights Employment Measure

Chetty et al. (2020) mention that the Homebase employment composition skews towards restaurant and retail workers. In contrast, they show that their employment measure more accurately tracks nationally representative statistics such as the Quarterly Census of Employment and Wages (QCEW) and the Occupational Employment Statistics (OES). To explore the possibility that our results are driven by the Homebase composition, we re-estimate Equation 1.1 using the OI employment measure at the county-level. This serves the dual purpose of accounting for county-level stay-at-home orders.

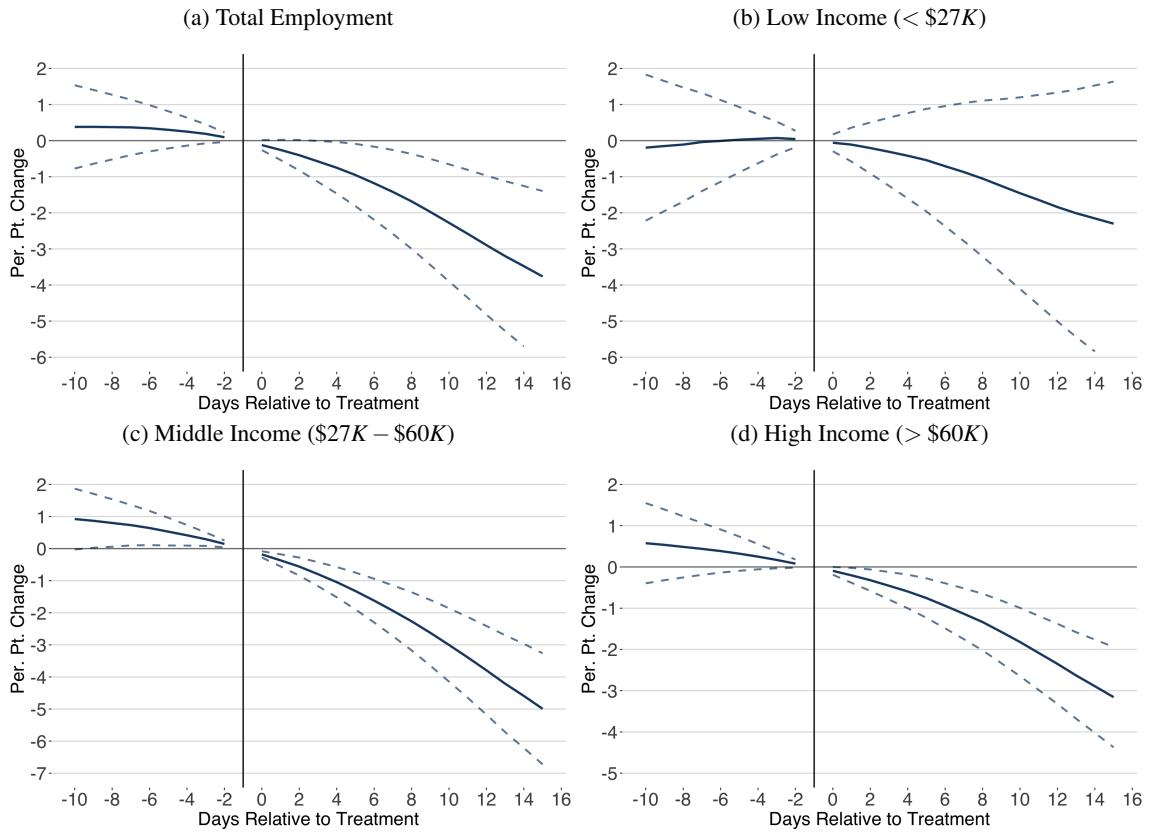
Figure 1.4 plots the results using the Opportunity Insights employment measure as the dependent variable. Observations are at the county level. Dashed lines represent 95% confidence intervals using standard

<sup>16</sup>Data for county-level stay-at-home orders and non-essential business closures and their timing is obtained from <https://ce.naco.org/?dset=COVID-19&ind=Emergency%20Declaration%20Types>. Over 100 counties issued a stay-at-home order before their state did. Appendix Figure 1.14 plots counties by their treatment status to gain an understanding of the cross-sectional variation induced by the county orders.

<sup>17</sup>Observations may differ across specifications due to missing values. Specifically, high-income expenditure data is missing for Montana, and low-income data is missing for Alaska and Vermont. Columns (7) and (8) in Appendix Table 1.9 re-estimate our baseline results when removing these states from the sample to provide evidence that these differences across income quartiles are not driven by differences in the sample composition.

<sup>18</sup>Standard errors for the county-level estimates are clustered at the state level.

Figure 1.4: Opportunity Insights Employment Measure



Note: This figure plots the stay-at-home order event study coefficients estimated from Equation 1.1. Subcaptions represent the dependent variable. Employment data is obtained from Opportunity Insights. Observations are at the county-by-day level. Counties begin treatment on the date they issued a stay-at-home order or the date the state they reside in issued an order. Dashed lines represent 95% confidence intervals using standard errors clustered at the state level.

errors clustered at the state level. Panel (a) plots the effect of issuing a stay-at-home order on total employment. There is no evidence of differential trends up to ten days before the issuance of an order. There is a statistically significant 1 percentage point decline in employment five days after the order. Similar to the baseline estimates, the effect on employment is amplified over time and remains statistically significant for the remainder of the fifteen day post-period. After fifteen days, the effect peaks at -3.8 percentage points. Overall, the Opportunity Insights data provides similar estimates as the Homebase data, but the seven-day moving average has created a significant lag in the employment effects compared with Column two of Table 1.2. The Day 15 estimate of -3.76pp from the OI data is more similar to the Day 10 estimate of -3.44pp from the Homebase data. The Day 10 and Day 5 estimates are respectively (-2.28pp, -2.81pp).

Panels (b)-(d) plot the employment response by worker’s yearly income level. We see heterogeneous effects across the income distribution. Low-income workers (< \$27K/Year) experience slightly more than a 2 percentage point decline in employment fifteen days after a stay-at-home order. This estimate is statistically

insignificant, as are all other point estimates for low-income workers. We see the strongest effect of -5pp in middle-income workers (\$27K – \$60K) fifteen days after an order. All post-period estimates are statistically significant for middle-income workers. However, we do find statistically significant estimates of a downward trend in employment in the ten days leading up to the order. High-income workers (> \$60K) experience a significant 3pp decline in employment fifteen days after the order. High-income workers also have a slight downward trend in employment leading up to the issuance of an order, but these estimates are statistically insignificant.<sup>19,20</sup>

### 1.3.2 Heterogeneity across Sectors and the Anticipation Effect

Our baseline expenditure estimates showed some evidence of an anticipation effect by consumers. This is likely the result of stockpiling goods in the days preceding a stay-at-home order. Many consumers shifted their spending from food services, such as restaurants, to grocery and food stores during the pandemic (Cav-allo, 2020). These markets and their goods vary significantly in their ability to be stockpiled. An expenditure anticipation effect would be expected for groceries, but not restaurant meals, for example. We now turn to an exploration of the heterogeneous effects in expenditures and the anticipation effect across several consumption categories.

Figure 1.5 plots the results of Equation 1.1 for (a) Grocery and Food Stores, (b) Apparel and General Merchandise Stores, (c) Accommodation and Food Services, and (d) Arts & Recreation as the dependent variable. Observations are at the state level. Dashed lines represent confidence intervals at the 95% confidence level using robust standard errors. All panels suggest similar dynamics in the post-treatment period. There is an initial decline in expenditure which experiences a statistically significant trough decline 6-8 days after a stay-at-home order is issued. After 6-8 days, expenditures begin to recover. However, these effects differ significantly in their magnitudes across sectors. Grocery and food stores exhibit the largest decrease in expenditures of 10 percentage points. In contrast, accommodation and food services only experience a 2pp decline. Trough declines are about 3pp for both apparel & general merchandise stores and arts & recreation expenditures.

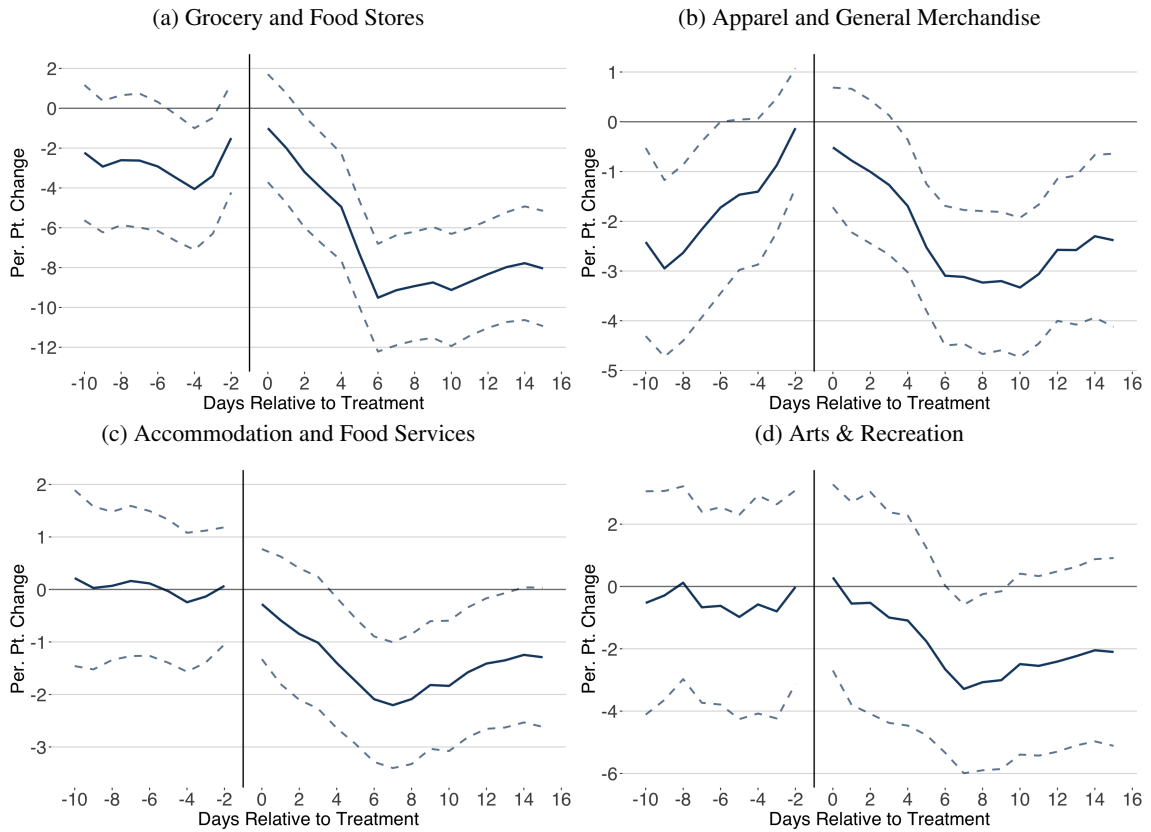
Pre-trends also vary significantly across specifications. The anticipation effect seen in our baseline expenditure results appears to be driven primarily by grocery stores and general merchandise stores. Panel (a) of Figure 1.5 shows that grocery and food stores experienced a statistically significant upward trend of 4pp in

---

<sup>19</sup>Workers may shift in the income distribution over time due to wage cuts. As an example, low-income workers may see an increase in employment after a stay-at-home order if middle-income workers on the threshold see a decrease in wages. Chetty et al. (2020) provide evidence that wages for individuals who remained employed during this period of the pandemic were largely unchanged. This suggests, given the significance levels in each panel and the large differences in magnitude across panels, that it is unlikely that shifts in the wage distribution will affect the overall qualitative story presented in Panels (b)-(d). Although, these shifts may affect the quantitative interpretation.

<sup>20</sup>Appendix Table 1.10 compares the county-level results in Figure 1.4 with the event study results at the state level. Additional heterogeneity checks at the state level for employment by NAICS sector are in Appendix Figure 1.16.

Figure 1.5: Heterogeneity across Consumption Expenditures



Note: This figure plots the stay-at-home order event study coefficients estimated from Equation 1.1. Subcaptions represent the dependent variable. Expenditure data come from Affinity Solutions. Merchant groups are defined by Affinity Solutions. Observations are at the state-by-day level. Dashed lines represent 95% confidence intervals using robust standard errors.


the days preceding a stay-at-home order. Apparel and general merchandise stores also see an upward trend that starts at -3pp nine days before the issuance of an order. These upward trends are completely absent in the ten days preceding an order for accommodation and food services as well as arts & recreation expenditures. Panels (c) and (d) provide further evidence that our baseline estimates are not significantly driven by consumer anticipation effects, and that stay-at-home orders had large impacts on the economy.<sup>21</sup>

### 1.3.3 Identifying Variation induced by Treatment Timing

Recent work has emphasized the impact that treatment timing can have on difference-in-difference and event study estimates (Goodman-Bacon, 2018; Callaway and Sant’Anna, 2020; Abraham and Sun, 2020). Goodman-Bacon (2018) shows that the baseline DD estimates (red lines in Figure 1.3) can be expressed as

<sup>21</sup>Appendix Figure 1.17 contains additional heterogeneity checks for Transportation & Warehousing and Health Care Services.

Figure 1.6: Effect of Treatment Timing on Hours Worked DD Estimate



Note: This figure plots the results of a Goodman-Bacon (2018) decomposition on our difference-in-difference estimate. The dependent variable is hours worked. The y-axis is the treatment-timing cohort DD estimate. The x-axis is the weight for the respective estimates. Cohorts are further divided into two groups: Treatment occurred before March 29th (Early, blue circles), Treatment occurred on or after March 29th (Late, red triangles). Hours worked data is obtained from Homebase. Observations are at the state-by-day level.

the weighted average of the 2x2 DD estimates for each cohort of treatment timing.<sup>22,23</sup> There are 19 cohorts in our framework ( $C_0 = \{ControlStates\}, C_1 = \{CA\}, C_2 = \{IL, NJ\}, \dots, C_{18} = \{SC\}$ ).

Each treated cohort has a DD estimate with the control group for a total of 18 estimates.<sup>24</sup> Figure 1.6 plots these estimates for hours worked. Cohorts are divided into two groups, those that received treatment before March 29th (Early, blue circles) and those that received treatment on or after March 29th (Late, red triangles).<sup>25</sup> Each treatment-cohort/control DD estimate is negative, so hours worked decreased after treatment for all cohorts relative to the control group. However, there is a significant difference in the DD estimates across the early/late treatment groups. Early treated groups, on average, experience about twice the relative decline in hours worked (14.8pp) compared to late treated groups (7.5pp).

<sup>22</sup>2x2 stands for two-period/two-group. Intuitively, the two periods are pre/post treatment. Groups are defined their treatment-timing. DD coefficients are then estimated for each two-group combination, hence 2x2.

<sup>23</sup>Weights rely on three main factors: (1) when units are treated, (2) the size of the treatment cohort, and (3) the length of the sample. Please see Goodman-Bacon (2018) for a full discussion on the estimation of these weights.

<sup>24</sup>The Goodman-Bacon (2018) decomposition also contains DD estimates across treated cohorts. In our case, this results in a total of 18 (treated/control) + 153 (treated/treated) = 171 observations. Appendix Figure 1.15 provides the full decomposition for both hours worked and consumption expenditures.

<sup>25</sup>Stay-at-home orders occurred over four weeks (recall Table 1.1). March 29th is used as the determinant for late treatment as it groups the first two weeks (Early) and last two weeks (Late).



### 1.3.3.1 Refined Event Study Specification

Figure 1.6 suggests that there may be heterogeneity in the estimated treatment effects across treatment cohorts. We refine the event study specification in Equation 1.1 to test this hypothesis:

$$y_{it} = \alpha_i + \gamma_t + \Gamma X_{it} + \sum_{g=1}^G \sum_{k \neq -1, -10 \leq k \leq 15} \beta_k^g I(i \in g) I(StayHome_{it} = k) + \varepsilon_{it} \quad (1.2)$$

where  $g$  specifies the group that state  $i$  belongs to. Notice that, compared to Equation 1.1,  $\beta_k$  is now allowed to vary across groups,  $\beta_k^g$ . We again consider the two groups of Early (Before March 29th) and Late (On or After March 29th) treatment. Thus, formally, our specification is

$$y_{it} = \alpha_i + \gamma_t + \Gamma X_{it} + \sum_{-10 \leq k \leq 15}^{k \neq -1} \beta_k^{Early} I(i \in Early) I(StayHome_{it} = k) + \sum_{-10 \leq k \leq 15}^{k \neq -1} \beta_k^{Late} I(i \in Late) I(StayHome_{it} = k) + \varepsilon_{it}$$

This specification is similar but not equivalent to the specification in Abraham and Sun (2020) as we relax their assumption that groups must belong to the same treatment-timing cohort.<sup>26</sup>

### 1.3.3.2 Results


Figure 1.7 plots the results of Equation 1.2 with the early treatment estimates in blue and late treatment in red. Similar to Figure 1.6, Panel (a) suggests that early treated states experience about twice the relative decline in hours worked than later treated states. Early treated states experience about a 4pp decline in hours worked for the entire post-period while late treated states experience a 2pp decline. Panel (b) estimates similar expenditure declines after the issuance of a stay-at-home order in early and late treated states, but significantly different pre-trends. The consumer anticipation effect appears to be driven primarily by early treated states which see a 3pp increase in expenditures before an order is enacted. Given the heterogeneous effects induced by treatment timing, we implement the estimator of Abraham and Sun (2020) in Appendix Figure 1.18. Overall, the estimates remain similar to our baseline results and suggest significant declines in both hours worked and expenditure.

### 1.3.4 Non-essential Business Closures

Section 1.2.3 mentioned that the specification in Equation 1.1 also includes an event-study covariate for non-essential business closures. The cumulative effect of implemented policies may be of interest in addition to the individual effect of stay-at-home orders. Figure 1.8 plots our baseline estimates for hours worked and total

<sup>26</sup>This specification is also similar to stratifying the event study by early and late treatment groups.

Figure 1.7: Effect of Treatment Timing on Event Study Estimates




Note: This figure plots the stay-at-home order event study coefficients estimated from Equation 1.2. Subcaptions represent the dependent variable. Treated states are divided into two groups: Treatment occurred before March 29th (Early, blue), Treatment occurred on or after March 29th (Late, red). Observations are at the state-by-day level. Hours and expenditure data are obtained from Homebase and Affinity Solutions, respectively.

expenditure side-by-side with non-essential business closures in blue and stay-at-home orders in red. Dashed lines represent confidence intervals at the 95% confidence level using robust standard errors. Standard errors are omitted for stay-at-home orders since they are shown in Figure 1.3. Panel (a) shows that non-essential business closures did not have a statistically significant effect on hours worked. This may be the result of a break in pre-trends before their issuance.

Panel (b) shows that non-essential closures did have a statistically significant effect on expenditures 4-10 days after their implementation with a trough on Day 6. The significant point estimates have a range from a 1-2 percentage point decline in expenditures. We find no evidence of a break in pre-trends up to ten days before the order. The pre-trends also show that the consumer anticipation effect seen for stay-at-home orders is not evident for non-essential business closures. This may be due to differences in the supply and demand implications between the two policies as non-essential business closures directly targeted the supply side (firms) of the market while stay-at-home orders directly targeted both supply (firms) and demand (consumers). An analog of Table 1.2 with robustness/heterogeneity checks for non-essential business closures can be found in the appendix. The results are similar to Figure 1.8 in that we find null effects of non-essential business closures as in Panel (a) or mitigated effects with limited statistical significance compared to stay-at-home orders as in Panel (b). Overall, the inclusion of non-essential business closures suggest that the cumulative effect of mitigation policy in response to Covid-19 is at least as great as our estimates from stay-at-home orders.

Figure 1.8: Effect of Non-essential Business Closures on the Economy



Note: This figure plots the non-essential business closure event study coefficients (blue) and stay-at-home order event study coefficients (red) estimated from Equation 1.1. Subcaptions represent the dependent variable. Dashed lines represent 95% confidence intervals using robust standard errors. Observations are at the state-by-day level. Hours and expenditure data are obtained from Homebase and Affinity Solutions, respectively.

### 1.3.5 Interpretation of Event Study Results

The event study estimates represent the average effect of issuing a stay-at-home order on hours worked and consumer expenditure for treated states. The previous sections presented these results as percentage point effects (e.g. 4pp decrease in hours worked). These percentage point effects can be used to estimate the absolute and cumulative effects of stay-at-home orders for treated states. Recall that all variables are measured relative to January 4-31. The average amount of consumer spending per day from Jan. 4-31 was \$21.82 billion (Chetty et al., 2020). To estimate the absolute effect of a stay-at-home order, we multiply the event study estimates by \$21.82 billion. The baseline estimate for consumer spending fifteen days after a stay-at-home order was -3.97%. The absolute effect is then  $-3.97\% * \$21.82 = \$0.87$  billion. The cumulative effect after fifteen days can be found by summing the absolute effect for all of the post-treatment days (0-15). For consumer expenditures, this equates to \$10.5 billion. However, this estimate assumes that the entire US was treated. We then multiply \$10.5 billion by the percentage of the population that was placed under a stay-at-home order to account for this.<sup>27</sup> The final estimate of the cumulative effect of stay-at-home orders on consumer spending is then  $95\% * \$10.5 \text{ billion} = \$9.94 \text{ billion}$ .

Panel A of Table 1.3 plots the cumulative spending loss for treated states calculated above in Column (1). The bracketed terms are calculated by summing the upper 95% confidence band and the lower band from our baseline expenditure estimates (Figure 1.3) over the entire post-period. This results in a range from \$6.28 billion to \$13.61 billion. A similar analysis is conducted for earnings per worker and total earnings for states

<sup>27</sup>This uses the simplifying assumption that consumer spending per capita is equal across states as we do not have access to the actual percentage of consumer spending by treated states.

Table 1.3: Cumulative and Relative Effects of Stay-at-Home Orders

<b>Panel A: Spending and Hours Worked</b>			
	Total Spending (1)	Earnings Per Worker (2)	Total Earnings (3)
January Benchmark	\$21.92 billion/day	\$139.31/day	\$26.23 billion/day
Cumulative Effect (Days 0-15)	\$9.94 billion [6.28, 13.61]	\$82.40 [44.19, 120.61]	\$15.51 billion [8.32, 22.71]
<b>Panel B: Employment</b>			
	Stay-at-Home Orders (1)	BLS Employment (2)	CPS Supplement (3)
Employment Loss	5,797,349	21,909,000	49,839,000
Relative Effect	100%	26.5%	11.6%

Note: This table presents the cumulative effect of stay-at-home orders on consumer spending and hours worked in Panel A. The cumulative effects are obtained by summing the stay-at-home order event-study coefficients of Equation 1.1 and multiplying this sum by the reported January benchmarks. We then multiply this result by the percent of the population that experienced a stay-at-home order during our sample period. Numbers in brackets present the confidence interval for the calculated cumulative effect by summing the 95% lower and upper confidence bands for the event-study coefficients. The January benchmark for consumer spending is obtained from Chetty et al. (2020). To obtain January benchmarks for earnings per worked and total earnings, we use BLS data on average private hourly earnings (CES0500000003), average private weekly hours (CES0500000002), and total nonfarm employment (CES0000000001). Column (1) of Panel B uses the difference-in-difference estimate for employment to compute the employment loss from stay-at-home orders. The BLS employment loss is computed by subtracting the BLS estimate of total nonfarm employment in April 2020 by the estimate in January 2020. The employment loss in Column (3) is obtained from the supplemental Covid-19 questions used in the May 2020 CPS. The relative effect is then calculated by dividing the employment loss from stay-at-home orders by the total employment loss in the respective columns.

that issued a stay-at-home order.<sup>28</sup> This suggests a cumulative loss of \$80 per worker shown in Column (2). We aggregate this effect over treated states by multiplying the January BLS non-farm employment numbers (152 million) by the percent of the treated population in the US. This suggests about a \$15.5 billion loss in earnings in the first fifteen days after a stay-at-home order. Bartik et al. (2020) show that stay-at-home orders had significant effects for almost 30 days after their issuance. Given this, our results likely serve as a lower bound for the cumulative effect of stay-at-home orders.

Panel B of Table 1.3 plots the relative employment loss from stay-at-home orders. Unlike consumer spending and hours worked, the effect of stay-at-home orders on employment should not be measured cumulatively. For example, an individual who becomes unemployed one day after a stay-at-home order is likely to be unemployed two days after the order as well. The cumulative effect would then count this individual more than once in the total employment loss. To circumvent this issue, we express the average effect of stay-

<sup>28</sup>We first calculate the average earnings per worker per day in January. To do this, we obtain average hourly earnings (\$28.43) and average weekly hours (34.3) from the BLS for all employees on private payrolls in January. Average daily earnings can then be calculated as the hourly earnings multiplied by average daily hours,  $\$28.43 \times (34.3/7) = \$139.31$ . The cumulative effect for earnings per worker is then calculated by summing our baseline event study estimates for hours worked multiplied by the daily wage.

at-home orders on employment relative to the total employment loss during our sample. The average effect of stay-at-home orders can be found by multiplying the difference-in-difference estimate for employment (-4%) by total employment in January (152 million).<sup>29</sup> As before, this estimate is scaled by the percent of the population under a stay-at-home order during this period which results in  $-4\% * 95\% * 152 \text{ million} = -5.8 \text{ million}$ . We normalize this 5.8 million employment loss relative to the total employment loss to isolate the relative impact of stay-at-home orders. The BLS employment estimate for April (130 million) suggests that 22 million jobs were lost in total. The relative effect is then computed by dividing the employment loss from stay-at-home orders by the total employment loss,  $\frac{5.8M}{22M} = 26\%$ .

In May, the BLS added several questions to the the Current Population Survey (CPS) to help gauge the effects of the coronavirus on the labor market (BLS, 2020). One of the questions was “At any time in the last 4 weeks, were you unable to work because your employer closed or lost business due to the coronavirus pandemic?”<sup>30</sup> The survey results from this supplemental question suggest 50 million people were unable to work. Column (3) of Panel B uses this estimate as the total employment loss which results in a 11.6% contribution of stay-at-home orders to the total employment loss.

#### **1.4 Model Framework: Optimal mitigation policy with multiple locations**

The empirical results of the previous section showed that stay-at-home orders had a significant effect on the economy. Beyond general concerns about the economic impact of the orders, policy makers often argued about how to best implement them. On April 1st, in regards to issuing a nationwide stay-at-home order, President Trump responded “There are some states that are different. There are some states that don’t have much of a problem . . . You have to give a little bit of flexibility.” (Press Briefing, 2020) Dr. Anthony Fauci expressed a different opinion the next day saying, “I just don’t understand why we’re not doing that [issuing a federally-mandated stay-at-home order]. We really should be.” (Politico, 2020)

This section builds on our empirical results by developing an economic framework to analyze virus transmission and optimal mitigation policy in a fiscal union such as the United States. The model will allow us to address virus transmission across locations. Thus, the economic activity of one location depends upon the spread of the virus in other locations.

##### **1.4.1 Initial Conditions**

We model the United States as a multi-location, closed economy version of the SIR macro model in Eichenbaum et al. (2020). We interpret these locations as states which we denote by  $i \in [1, N]$ , where  $N$  is the

<sup>29</sup>See Appendix Table 1.8 for difference-in-difference estimates. As before, non-farm employment in January is obtained from the BLS.

<sup>30</sup>The CPS is conducted during the week of the 19th. Thus, the four weeks before May 19th slightly abstracts from the sample period of our analysis.

total number of states. The initial population in a state as a fraction of the national population is written as  $Pop_{i,0} \in (0, 1]$ .<sup>31</sup> A fraction of the state population becomes infected at  $t = 0$ ,  $T_{i,0} = \varepsilon_i$ . All agents are initially susceptible,  $S_{i,0} = 1 - T_{i,0}$ . Thus, active infections in state  $i$  at date 0 is  $I_{i,0} = T_{i,0}$ . This allows states to vary in both initial population and initial infections.

#### 1.4.2 Virus Transmission

As in Eichenbaum et al. (2020), we assume the virus is transmitted within a state through consumption and employment. We further allow the virus to transfer within and across states through an exogenous variable. For a given Home state, which we denote by  $H$ , this transmission mechanism at time  $t$  is given by:

$$T_{H,t} = \pi_{s1}(S_{H,t}C_{H,t}^S)(I_{H,t}C_{H,t}^I) + \pi_{s2}(S_{H,t}N_{H,t}^S)(I_{H,t}N_{H,t}^I) + \pi_{s3}S_{H,t} \sum_{i=1}^N (1 - \tau_{H,i}^d)(I_{i,t}Pop_{i,0}) \quad (1.3)$$

where  $C_{H,t}^S$  is consumption at time  $t$  by susceptible agents in the Home state, and  $C_{H,t}^I$  is consumption by infected agents. Similarly,  $N_{H,t}^S$  and  $N_{H,t}^I$  are hours worked by susceptible and infected agents in the Home state, respectively. Note that  $\tau_{H,i}^d$  represents the travel cost between locations  $H$  and  $i$ . This yields an interpretation of  $1 - \tau_{H,i}^d$  as the percentage of susceptible agents in state  $H$  that interact with infected agents in state  $i$ . The parameters  $\pi_{s1}$ ,  $\pi_{s2}$ , and  $\pi_{s3}$  are the probability of transmitting the virus through each of the three associated mechanisms.

The first two terms of our transmission mechanism, apart from the location indices, are identical to those in Eichenbaum et al. (2020). The substantive difference between the two models is the presence of the third term. This is evident by the fact that our model replicates Eichenbaum et al. (2020) if we assume that the travel cost between states,  $\tau_{H,i}^d$ , equals zero and infection rates are symmetric, or in the trivial case of modelling one location. The third and novel mechanism is similar in spirit to the empirical model of Adda (2016), but we loosen the restriction that the travel cost is equal between all pairs of states. Our only assumptions are that (1) travel cost within a state is equal to zero,  $\tau_{H,H}^d = 0$ , and (2)  $\tau_{H,i}^d \in [0, 1]$  to prohibit negative transmission across locations.

The spread of the virus in each state can then be modeled by iterating

$$S_{t+1} = S_t - T_t \quad (1.4)$$

$$I_{t+1} = (1 - \pi_r - \pi_d)I_t + T_t \quad (1.5)$$

$$R_{t+1} = R_t + \pi_r I_t \quad (1.6)$$

$$D_{t+1} = D_t + \pi_d I_t \quad (1.7)$$

<sup>31</sup>The initial population of a state equals one if and only if  $N = 1$ .

where the state index  $i$  is suppressed for notational simplicity.  $T_t$  follows the transmission mechanism in Equation 1.3 in each state.  $R_t$  and  $D_t$  represent the total fraction of recovered and deceased. The parameters  $\pi_r$  and  $\pi_d$  represent the probability an infected person recovers or dies.

### 1.4.3 The Economy

We assume that states are closed economies with the exception of virus transmission.<sup>32</sup> This allows us to isolate the economic spillovers of virus transmission across locations while replicating the Eichenbaum et al. (2020) results as a special case. This allows us to focus specifically on the economic effects of the cross-state virus transmission mechanism, and simplifies the solution to the state-level decision problems. In the text that follows, the location index  $i$  is suppressed except where needed for clarity of exposition. Individual state economies can be written as follows:

#### 1.4.3.1 Households

We assume agents' discounted lifetime utility depends on their infection status  $j \in (s, i, r)$ . As infection status follows the transition matrix presented in Equations 1.3-1.7, lifetime utility of the various agents can be written as

$$U_t^r = u(c_t^r, n_t^r) + \beta U_{t+1}^r \quad (1.8)$$

$$U_t^i = u(c_t^i, n_t^i) + \beta [(1 - \pi_r - \pi_d)U_{t+1}^i + \pi_r U_{t+1}^r + \pi_d \times 0] \quad (1.9)$$

$$U_t^s = u(c_t^s, n_t^s) + \beta [(1 - \pi_{I,t})U_{t+1}^s + \pi_{I,t}U_{t+1}^i] \quad (1.10)$$

where  $c_t$  is consumption,  $n_t$  is hours worked, and  $\pi_{I,t}$  represents the probability a susceptible agent becomes infected. In the model, this is found by dividing Equation 1.3 by  $S_t$  which yields

$$\pi_{I,t} = \pi_{s1}(C_t^S)(I_t C_t^I) + \pi_{s2}(N_t^S)(I_t N_t^I) + \pi_{s3} \sum_{i=1}^N (1 - \tau_{H,i}^d)(I_{i,t} Pop_{i,0}) \quad (1.11)$$

The remainder of the economic environment follows Eichenbaum et al. (2020). Preferences of agents in each state take the functional form,

$$u(c_t^j, n_t^j) = \ln c_t^j - \frac{\theta}{2} (n_t^j)^2 \quad (1.12)$$

<sup>32</sup>The closed economy assumption abstracts from endogenous trade fluctuations and international policy such as travel restrictions in response to the virus. Antràs et al. (2020) and Fajgelbaum et al. (2020) develop open-economy models to analyze these avenues. Importantly, these papers provide support that the qualitative results in Section 1.5 are likely to hold if we relax the closed economy assumption.

Consumers maximize utility over an infinite horizon subject to the budget constraint

$$(1 + \mu_{ct})c_t^j = w_t \phi^j n_t^j + \Gamma_t \quad (1.13)$$

where  $w_t$  is the real wage,  $\phi$  is labor productivity, and  $\Gamma_t$  are lump-sum transfers from the government financed by the tax on consumption,  $\mu_{ct}$ . Intuitively, the consumption tax can be viewed as a proxy for mitigation policies such as stay-at-home orders as it creates a tradeoff between mitigating the spread of the virus and negatively impacting the economy.<sup>33</sup>

#### 1.4.3.2 Firms

In each state, a continuum of representative firms maximize profits

$$\Pi_t = Y_t - w_t N_t \quad (1.14)$$

where  $N_t$  is hours worked and  $Y_t = AN_t$  is output.

#### 1.4.3.3 Government

Each state government is subject to a balanced-budget requirement:

$$\mu_{ct} (S_t c_t^s + I_t c_t^i + R_t c_t^r) = \Gamma_t (S_t + I_t + R_t) \quad (1.15)$$

for all  $t$ .<sup>34</sup>

### 1.5 Model Implications: Multiple locations and optimal policy

We illustrate the dynamics of our model through several numerical examples. First, we compare the SIR macro model with one location to our multiple location model. We use the baseline parameters in Eichenbaum et al. (2020) for this exercise (see Table 1.4). Eichenbaum et al. (2020) study how various parameterizations affect the dynamics of the single-location model, so we do not replicate that sensitivity analysis here. Various changes to their parameters are unlikely to significantly affect the main contribution of our model which is the cross-sectional distribution we obtain. However, our model does contain several new parameters: (1) state population, (2) initial state infection rates, and (3) travel cost across locations. The second subsection develops a procedure to parameterize travel costs across locations which we apply to US Census region data.

<sup>33</sup>The interpretation of the consumption tax as a proxy for mitigation is discussed more in Section 1.5.3.

<sup>34</sup>Requiring the government budget constraint to hold at the state level implies that fiscal transfers across states are equal to zero. Loosening this restriction to hold at the national level—net fiscal transfers are zero—would create the opportunity to study optimal transfers across locations. This is a possible avenue for future research to study the effect of large fiscal programs such as the Paycheck Protection Program or the Cares Act, but it is not as relevant for our analysis of mitigation policies such as stay-at-home orders.



Table 1.4: Baseline Model Parameters

	Parameter (1)	Value (2)
<b>Panel A: New Parameters</b>		
State Population (Baseline)	$Pop_{i0}$	$\frac{1}{N}$
Initial State Infection Rate (Baseline)	$\varepsilon_{i0}$	0.001
Travel Cost (Baseline)	$\tau_{H,i}^d$	$\frac{ H-i }{N}$
<b>Panel B: Eichenbaum et al. (2020)</b>		
Initial Country Population	$POP_0^{Agg}$	1
Initial Country Infected	$\varepsilon^{Agg}$	0.001
Productivity	$A$	10
Labor Productivity (Susceptible)	$\phi^s$	1
Labor Productivity (Infected)	$\phi^i$	0.8
Labor Productivity (Recovered)	$\phi^r$	1
Disutility of labor	$\theta$	9
Discount Rate	$\beta$	$0.98^{1/52}$
Probability of consumption infection	$\pi_{s1}$	0.0185
Probability of labor infection	$\pi_{s2}$	1.8496
Probability of exogenous infection	$\pi_{s3}$	0.2055
Probability of recovery	$\pi_r$	$7 \times \frac{0.99}{18}$
Probability of death	$\pi_d$	$7 \times \frac{0.01}{18}$
Consumption Tax	$\mu_{ct}$	0

Note: This table presents the parameters used in our baseline model specification in Section 1.5. Panel A presents variables specific to our multiple-location model. Panel B presents the baseline parameters used in Eichenbaum et al. (2020) which remain in our model.

This application allows us to perform a sensitivity analysis to study the effect that travel cost across locations has on the cross-sectional distribution of infection rates.<sup>35</sup> In turn, we study how the cross-section affects the aggregate dynamics of the model. Lastly, we study optimal mitigation policy when there are multiple locations with asymmetric infection rates.

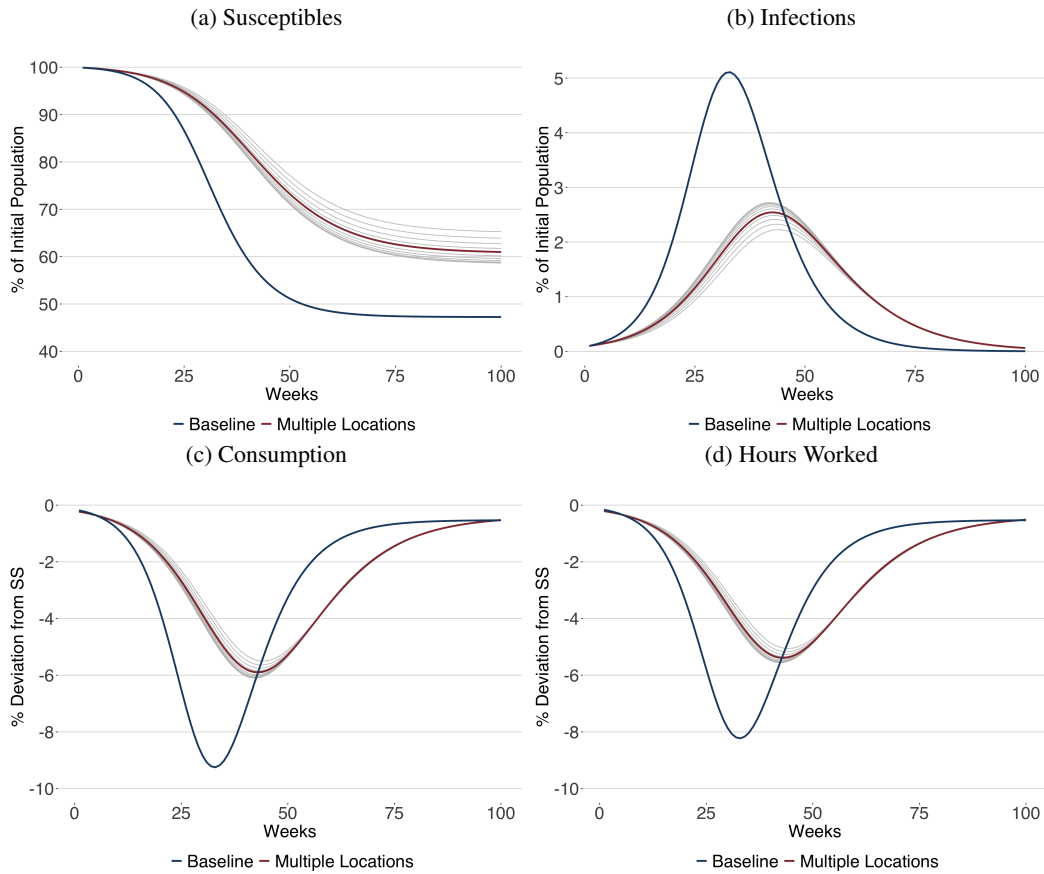
### 1.5.1 Multiple-location model vs single-location model

Figure 1.9 plots results from our simulated model with one location (blue) and twenty-one locations (red). The grey lines represent the individual states. The red line is the national total computed as the population-weighted average of the grey lines. All states in the country are symmetric in terms of their initial populations and initial infected. Thus, these results are entirely driven by our travel cost which is  $\tau_{H,i}^d = \frac{|H-i|}{N}$  where  $N$  is the total number of locations in the country.<sup>36</sup> As mentioned earlier, setting travel cost to 0 when initial

<sup>35</sup>The appendix contains a full sensitivity analysis of state population, state infection rates, and travel cost across locations.

<sup>36</sup>This travel cost is selected as it meets the assumptions that (1) travel cost within a state is equal to 0, and (2)  $\tau_{H,i} \in [0, 1]$  for all state pairs. The dynamics of the cross-section will depend on the travel cost selected. Figures 1.10 and 1.20 provide a more in-depth analysis of the effect of different travel costs on the cross-sectional distribution.

Figure 1.9: Multiple-location Model vs Single-location Model



Note: This figure plots the results of simulations using the baseline SIR macro model with one location (blue) and our extended multiple-location version (red). Grey lines are produced from the multiple-location model and represent individual states. The red line is the national total from the multiple-location model, or equivalently the population-weighted average of the grey lines. Twenty-one locations were simulated.

infection rates are symmetric would result in all states following the blue line.

Panel (a) of Figure 1.9 plots the simulated fraction of susceptibles. With one location, we see that about 45% of the country’s population becomes infected by the end of the sample. The multiple-location results suggest that a little less than 30% of the population contracts the virus. Furthermore, the spread of the virus is not symmetric across locations even though the initial conditions are. This creates a cross-section in which 25% of the state population contracts the virus in the least affected state while over 30% contract it in the most severely impacted state.

Panel (b) of Figure 1.9 plots the simulated infection rates for both models. We see that the baseline SIR macro model suggests a peak infection rate of 3.5% at about 30 weeks. With multiple locations, the country-wide peak occurs almost 20 weeks later and is much lower at about 1.5%. There are two interesting results seen in the cross-section of peak infection rates. Not only do peak infection rates differ across states,

but they also differ in their timing. The intuition for this can be seen in Panels (c) and (d) which plot the percent deviation of consumption and hours worked from their steady state values. As states experience more infections, susceptible agents endogenously choose to decrease their consumption and hours to lower their probability of being infected. In turn, this contributes to earlier turning points in the infection rate.

## 1.5.2 Application to US Census Region Data

The previous subsection showed the capability that our model has to generate a cross-sectional distribution of economic variables and infection rates. We demonstrated the important role played by the travel cost across locations,  $\tau_{H,i}^d$ . Although our chosen travel cost function is convenient to illustrate the differences between the single and multiple-location models, it is impractical for applied settings. In this section, we develop a procedure to parameterize trade costs across locations. We then apply the procedure to US Census region data to illustrate the advantages of this approach.

### 1.5.2.1 Calibration Procedure

Panel A of Figure 1.10 plots the spread of the virus in the US by Census Region (Midwest, Northeast, South, West) and the national total (dashed line) from May 6th through August 26th. National and Census region data are obtained by aggregating the county-level infection and population data from Section 1.2. The daily data are transformed to a weekly frequency in accordance with the model. The following steps are followed to generate simulations from the model:

1. Set Number of Locations ( $N$ )
2. Set Initial Population of Locations ( $Pop_0$ )
3. Set Susceptible Percent of Population ( $S_t$ )
4. Set New Infections ( $T_t$ )
5. Set Active Infections ( $I_t$ )
6. Set Parameters ( $\tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3}$ )<sup>37</sup>

In this application, the number of locations is set equal to the number of Census regions,  $N = 4$ . The corresponding vector of populations by region is  $Pop_0 = (20.9\%, 17.2\%, 38.1\%, 23.8\%)$  for the Midwest, Northeast, South, and West respectively. The percent of susceptibles at each date is 1 less the cumulative fraction of the population that has been infected up to that date,  $S_t = 1 - \sum_{k=0}^t T_k$ . On May 6th, the susceptible vector is  $S_t = (99.7\%, 98.9\%, 99.8\%, 99.8\%)$ . The number of new infections per capita for the week ending on May 6th were  $T_t = (0.07\%, 0.11\%, 0.04\%, 0.03\%)$ . We approximate current active infections as  $I_t = T_t + (1 - \pi_r - \pi_d)T_{t-1}$ . This approximation produces  $I_t = (0.11\%, 0.22\%, 0.06\%, 0.05\%)$ . We then jointly set

<sup>37</sup>Appendix 1.7.2.2 provides details on why these parameters are chosen to fit the data.

Table 1.5: Initial Conditions for Census Region Application (May 6)

<b>Panel A: Initial Conditions</b>				
	Midwest	Northeast	South	West
Percentage of US Population ( $Pop_0$ )	20.9%	17.2%	38.1%	23.8%
Susceptible ( $S_t$ )	99.7%	98.9%	99.8%	99.8%
New Infections ( $T_t$ )	0.07%	0.11%	0.04%	0.03%
Active Infections ( $I_t$ )	0.11%	0.22%	0.06%	0.05%

<b>Panel B: Transmission Parameters</b>		
	Parameter	Value
Probability of consumption infection	$\pi_{s1}$	0.0009
Probability of labor infection	$\pi_{s2}$	1.53
Probability of exogenous infection	$\pi_{s3}$	0.732

<b>Panel C: Transmission Matrix (<math>1 - \tau_{H,i}^d</math>)</b>				
Virus Destination ( $H$ )	Virus Origin ( $i$ )			
	Midwest	Northeast	South	West
Midwest	1.00	0.216	0.030	0.326
Northeast	1.1e-4	1.00	5.6e-4	0.001
South	0.009	0.472	1.00	0.247
West	0.220	0.391	0.264	1.00

Note: This table presents the initial conditions and parameters used in the Census region application. Panel A presents the initial conditions for each Census region on May 6th. Panels B and C present the estimated parameters from solving Equation 1.16.

the travel cost across locations as well as the other parameters to minimize the sum of squared errors between susceptible population percentages in the model and the data,<sup>38</sup>

$$\min_{\tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3}} \sum_{i=1}^4 \sum_{t=May6}^{Aug26} (S_{i,t}^{Data} - S_{i,t}^{Model})^2 \quad (1.16)$$

As in Section 1.4, we impose the restrictions on travel cost:  $\tau_{H,i}^d \in [0, 1]$  and  $\tau_{H,i}^d = 0$  for  $H = i$ . We also impose the restrictions that the virus transmission parameters are positive  $\pi_{s1}, \pi_{s2}, \pi_{s3} \geq 0$ . This process yields estimates of ( $\pi_{s1} = 0.0009, \pi_{s2} = 1.53, \pi_{s3} = 0.732$ ).

### 1.5.2.2 Interpretation of Transmission Matrix

Table 1.5 provides a summary of the initial conditions and the calibrated parameters from this procedure. Among the more interesting estimation results are those pertaining to the travel costs,  $\tau_{H,i}^d$ . We refer to the set of  $1 - \tau_{H,i}^d$  as the transmission matrix as it relates to the interaction between susceptible agents in region  $H$

<sup>38</sup>The appendix provides a more detailed procedure to match both the time series of susceptibles and consumer spending in the Census region data.

with infected individuals in region  $i$ . Specifically,  $1 - \tau_{H,i}^d$  represents the probability that a susceptible individual in region  $H$  interacts with an infected individual from region  $i$ .<sup>39</sup> Rows represent the virus destination (H), and columns the virus origin (i). Thus, the first row represents the vector  $1 - \tau_{Midwest,i}^d$ . As we imposed the restriction that susceptible and infected agents fully interact within regions, we see that  $1 - \tau_{Midwest,Midwest}^d = 1$ . The second column produces an estimate of  $1 - \tau_{Midwest,Northeast}^d = 0.216$ . This suggests that, over the sample period of May 6th through August 26th, susceptible agents in the Midwest interacted with 21% of infected agents from the Northeast on average. The third and fourth columns suggest that susceptibles in the Midwest interacted at lower rates with the South (3%) but higher rates with the West (32.6%).

The transmission matrix suggests that there is substantial and highly asymmetric interaction across regions as most elements are non-zero and unequal. One exception is the row for the Northeast region,  $1 - \tau_{Northeast,i}^d$ , in which the cross-region estimates are all close to zero. This suggests the susceptible agents in the Northeast have not interacted with infected agents of other regions. In contrast, the Northeast contributed significantly to the spread of the virus to other regions. The two largest elements of the matrix are the transmission to the South from the Northeast (47.2%) and to the West from the Northeast (39.1%). It is reassuring that the transmission matrix backed out of this SIR economic model accords with the statement by Yale epidemiologist, Nathan Grubaugh, “We now have enough data to feel pretty confident that New York was the primary gateway for the rest of the country” (New York Times, 2020).

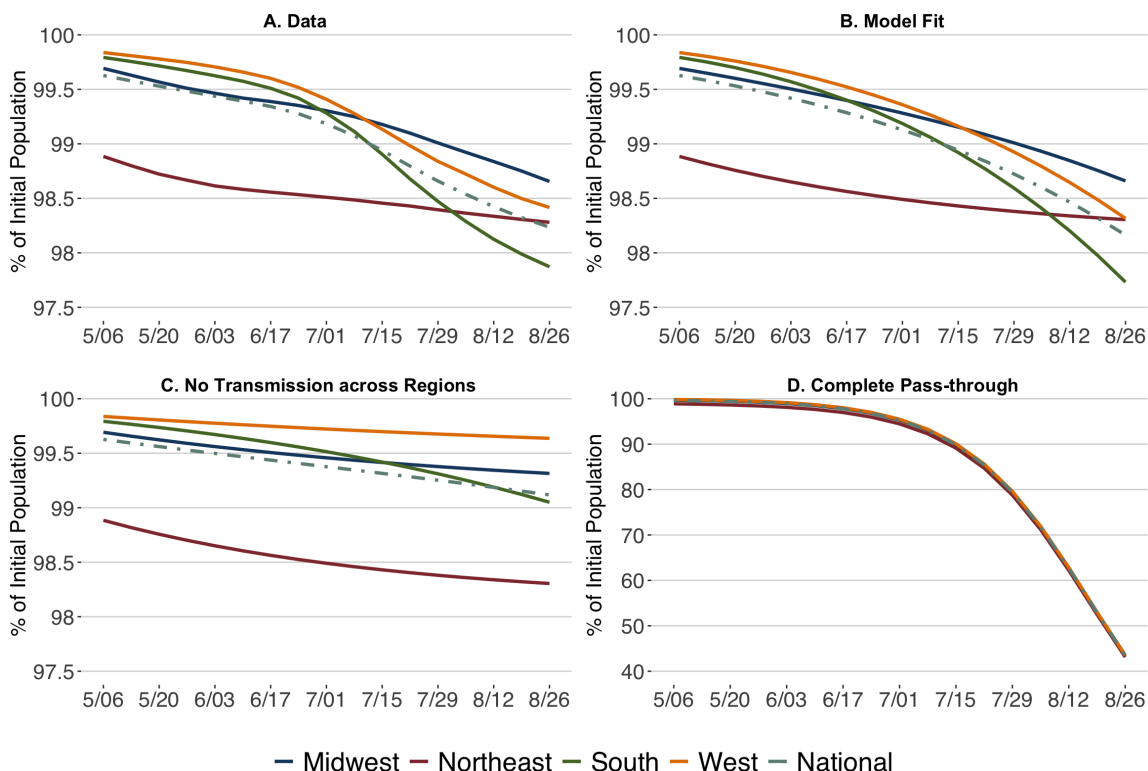
### 1.5.2.3 Results

Panel B of Figure 1.10 plots the results using the above procedure and then simulating our model for 17 weeks in accordance with the length of the sample period. We see that the model performs well in replicating the cross-sectional distribution. The susceptible population in the data for the week starting on August 26 is  $S_t^{Data} = (98.7\%, 98.3\%, 97.9\%, 98.4\%)$  for the Midwest, Northeast, South and West. The model produces susceptible estimates of  $S_t^{Model} = (98.7\%, 98.3\%, 97.7\%, 98.3\%)$ . We also see that the model performs well in replicating the individual time series for each region as well as the crossing patterns across regions. In the data, the South first surpasses the Midwest in their cumulative cases per capita on July 1st and subsequently surpasses the Northeast on August 5th. In the model, these dates are June 24th and August 12th. Lastly, the West experiences higher cumulative cases per capita than the Midwest starting on July 15th, and the model predicts a crossing on July 22nd.

---

<sup>39</sup>In a reduced-form manner, our calibration procedure for  $1 - \tau_{H,i}^d$  captures the average real-life interactions across regions over the sample period (e.g. travel to and from college, vacations, visits to care for dependents).

Figure 1.10: US Census Region Data vs Model (Susceptibles)



Note: This figure plots the susceptible percentage of the population in the United States and each Census region. Panel A plots the observed time series in the data. Panel B plots the model fit after jointly optimizing the travel distance across locations and the transmission parameters ( $\tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3}$ ). Panels C and D present comparative static exercises using  $\pi_{s1}, \pi_{s2}$ , and  $\pi_{s3}$  from Panel B. Panel C allows no virus transmission across Census regions. Panel D allows complete pass-through of the virus across regions. Infection data is obtained from *USA Facts*.

### 1.5.2.4 Comparative Statics

Panels C and D emphasize the importance of modelling the interaction of virus transmission across locations. We again present the two extreme cases of no transmission across regions ( $\tau_{H,i}^d = 1$  for all  $H \neq i$ ) and complete pass-through of infections across locations ( $\tau_{H,i}^d = 0$ ). The remaining parameters are the same as in Panel B. Thus, differences across panels are due entirely to differences in the transmission matrix,  $1 - \tau_{H,i}^d$ . Panel C plots the results for no transmission across regions. This underestimates the infection rates of all regions compared to the data. The percentage of susceptible individuals predicted by the model on August 26 are now  $S_i^{Model} = (99.3\%, 98.3\%, 99.1\%, 99.6\%)$ . The model now predicts only one cumulative infection rate crossing. Cumulative infections in the South exceeds the Midwest on July 22nd, several weeks past the equivalent data moment.

Panel D plots the results when there is complete pass-through of infections across locations. This provides the opposite result and significantly overestimates the spread of the virus. The model now predicts that all regions have converged with  $\sim 57\%$  of their population contracting the virus by August 26th. Table 1.6

Table 1.6: US Census Region Data vs Model Prediction (August 26)

<b>Panel A: Cumulative Infections</b>					
	Aggregate	Midwest	Northeast	South	West
Data	1.8%	1.3%	1.7%	2.1%	1.6%
Calibrated Model	1.8%	1.3%	1.7%	2.3%	1.7%
No Transmission	0.9%	0.7%	1.7%	0.9%	0.4%
Complete Pass-through	56.5%	56.4%	56.8%	56.4%	56.4%

<b>Panel B: Crossing Patterns</b>			
	South > Midwest	South > Northeast	West > Midwest
Data	July 1	August 5	July 15
Calibrated Model	June 24	August 12	July 22
No Transmission	July 22	NA	NA
Complete Pass-through	NA	NA	NA

Note: This table presents the results of our Census region application. Panel A plots the cumulative infections on August 26th. Panel B plots the dates in which a given Census region passed another region in cumulative infections. Both panels present the results seen in the data and three model simulations: (1) the calibrated model using the parameters in Table 1.5, (2) allowing no virus transmission across locations, and (3) allowing complete pass-through of the virus across locations.

summarizes the results of our Census region application for ease of comparability across model specifications. Overall, the model simulations suggest that the virus was transmitted partially, but not fully, across locations in the US.

### 1.5.3 Optimal Mitigation Policy

The previous subsection showed that there is substantial interaction across regions in the US. Importantly, the Census region application also suggests that there is not complete pass-through of the virus across locations. This section builds on this finding by analyzing optimal mitigation policy when initial infection rates are asymmetric and the virus is partially, but not fully, transmitted across locations.

#### 1.5.3.1 Planner's Problem

The equilibrium conditions in the appendix show that infected agents do not internalize the negative externalities they impose on other agents in the economy. This creates the role for government intervention to help mitigate virus transmission through the consumption tax of the model,  $\mu_{ct}$ . This tax can be viewed as a proxy for containment measures as the government can reduce the consumption and hours worked of infected agents by increasing  $\mu_{ct}$ .<sup>40</sup> However, like stay-at-home orders, the consumption tax does not specifically target infected agents. This creates a tradeoff between decreasing transmitted infections,  $T_t$ , and negatively affecting

<sup>40</sup>Importantly, the event study results in Section 1.3 showed that stay-at-home orders did indeed reduce spending and hours worked which provides empirical support for the use of the consumption tax as a proxy for mitigation policy.

the economy.

To formalize this in a multiple-location model, assume that a social planner sets  $\mu_{ct}$  to maximize the population-weighted average of the discounted lifetime utility of all agents:

$$\max_{\mu_{ct}} U_0 = \sum_{i=1}^N Pop_{i0} (S_{i0}U_{i0}^S + I_{i0}U_{i0}^I + R_{i0}U_{i0}^R) \quad (1.17)$$

where lifetime utility of an agent of type  $j \in (s, i, r)$  follows Equations 1.8-1.10. This problem best describes a national government which maximizes the above welfare function with the ability, but not the requirement, to restrict state governments to follow the same policy rule.<sup>41</sup> This allows us to evaluate the tradeoffs that exist between national mandates as suggested by Dr. Fauci and allowing state-level discretion as implied by President Trump.

### 1.5.3.2 Results

Figure 1.11 plots the results of a two-location simulation using the optimal policy rule derived from the welfare function in Equation 1.17. Both states have an initial population of  $\frac{1}{2}$ . We consider one state to be severely infected (blue) with an initial infection of 1%, and the other to have a mild infection (red) of 0.1%. As before, we set  $\tau_{H,i}^d = \frac{|H-i|}{N}$ . The dashed lines illustrate the baseline consumption tax rate of zero in all periods. The optimal policy of the severely and mildly infected states differs for almost 20 weeks. During this period, the severely infected state sets a higher consumption tax to mitigate their infection rate. After twenty weeks, the infection rates of the two states converge which leads to similar mitigation policies.

These results suggest that it is not optimal to set a nationwide mitigation policy. Appendix Figure 1.21 plots an analogous simulation when restricting both states to the same policy rule to understand this result. A common, national policy leads to similar virus transmission dynamics in both states compared to Figure 1.11. However, the mildly infected state experiences significantly larger declines in consumption at the onset of the pandemic. Simply put, a nationwide mandate is too restrictive for the economy of the mildly infected state and does not help contribute to the mitigation of the virus relative to local mandates. These results suggest that policy should be set at a local level contingent on sufficient asymmetries in infection rates.

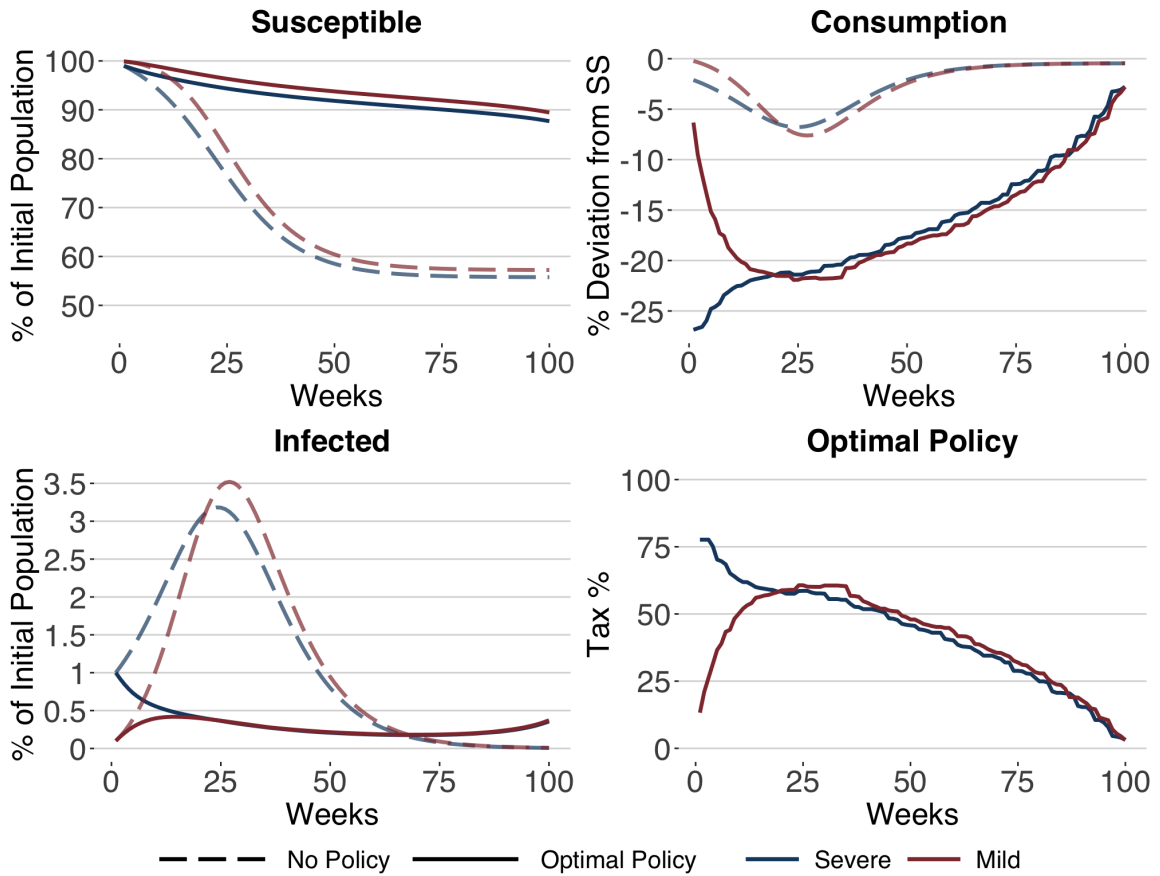
## 1.6 Conclusion

We show that stay-at-home orders reduced consumption and employment by 4 percentage points using an event study framework. Our estimates suggest that stay-at-home orders contributed to the total employment loss by almost 6 million employees. This accounts for 10%-25% of the total employment decline that oc-

<sup>41</sup>State governments are likely to place different weights on the welfare of their constituents and the welfare of the rest of the country. Our planning problem allows us to abstract from this complication. It is within this context that our model is that of a fiscal union.



Figure 1.11: Optimal Mitigation Policy



Note: This figure plots the results of a 2-location SIR macro model with (solid lines) and without (dashed lines) policy intervention. Both locations have an initial population of  $\frac{1}{2}$ . Initial infections to 1% and 0.1% for the severe (blue) and mild (red) infected locations, respectively.

curred during our sample period. Cumulatively, stay-at-home orders contributed a loss of about \$10 billion in consumer spending and \$15 billion in total wages.

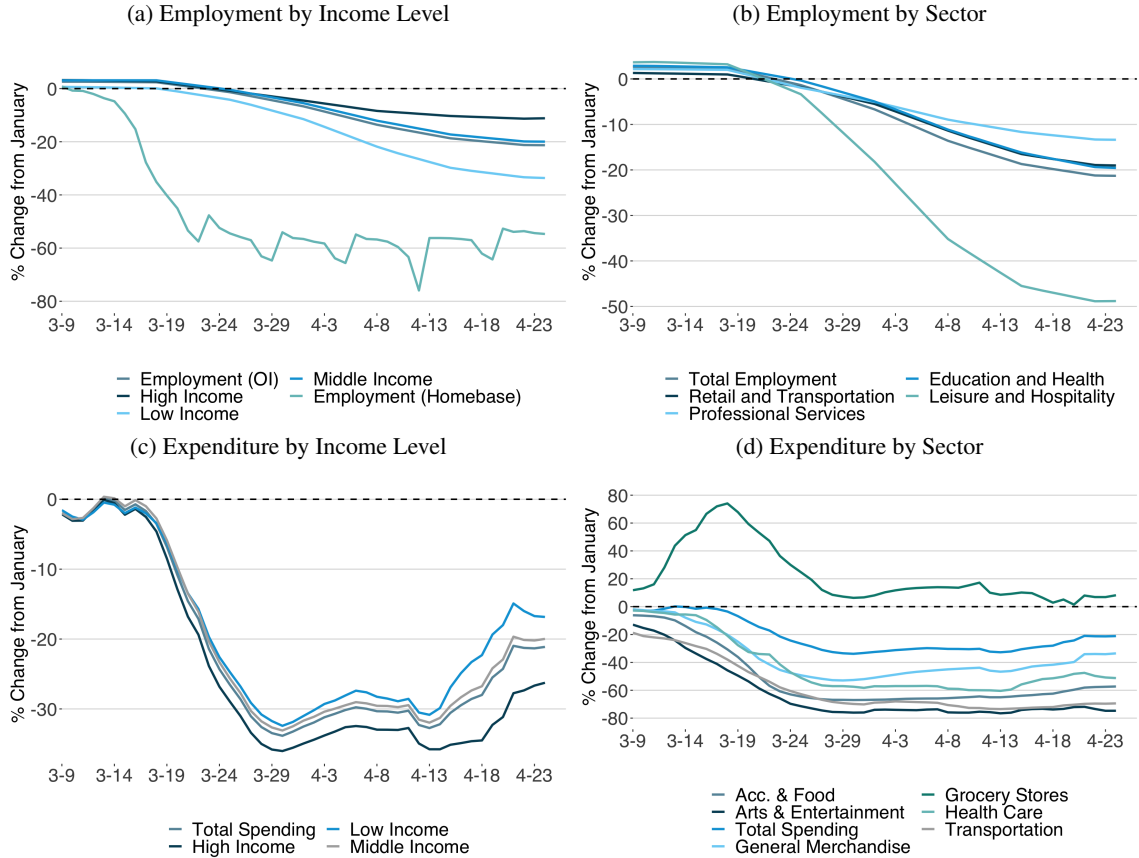
We pair our empirical findings with an economic SIR model to evaluate virus transmission and optimal mitigation policy in a fiscal union. We show that travel costs are important for both the aggregate and cross-sectional estimation of the virus spread and economic equilibrium. The extended model is shown to fit both the infection transmission and consumption time series at the Census region level. Lastly, we find that it is optimal for states to set their own mitigation policies when infection rates are asymmetric. Our model, combined with our simulation approach, demonstrates the value of data-driven public policy decisions that factor in the economic and health consequences of Covid-19 as well as the value of local and state policy discretion.

## 1.7 Appendix

### 1.7.1 Data

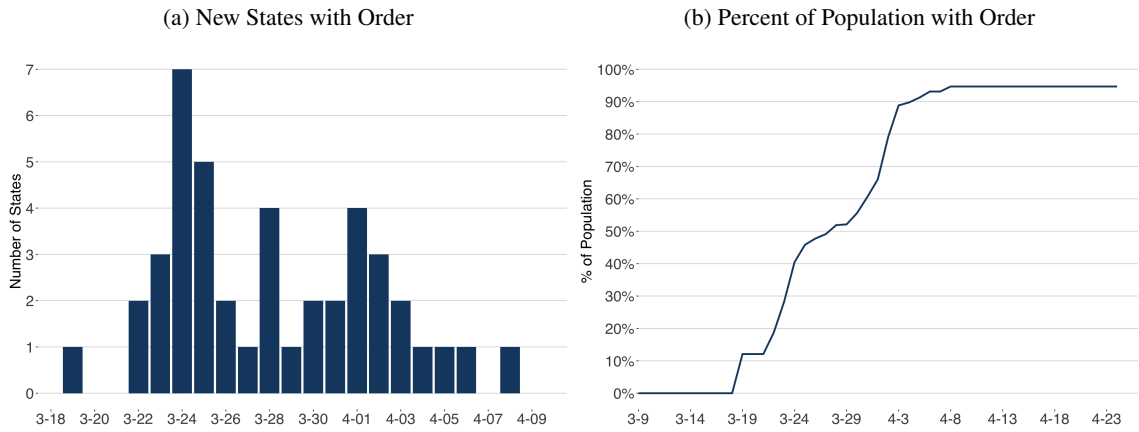
#### 1.7.1.1 Summary Statistics

Figure 1.12: Heterogeneous Response to Covid-19 by Income and Sector



Note: This figure plots the aggregate time series of employment and consumption expenditures from March 9th to April 24th. Employment data is obtained from Homebase and Opportunity Insights. Expenditure data is sourced from Affinity Solutions. Panel (a) plots the employment response by income level. Panel (b) plots the employment response by sector. Panel (c) plots the expenditure response by income level. Panel (d) plots the expenditure response by sector. All variables are measured relative to their respective counterparts from January 4th-31st.

Figure 1.13: Timeline of Stay-at-Home Orders



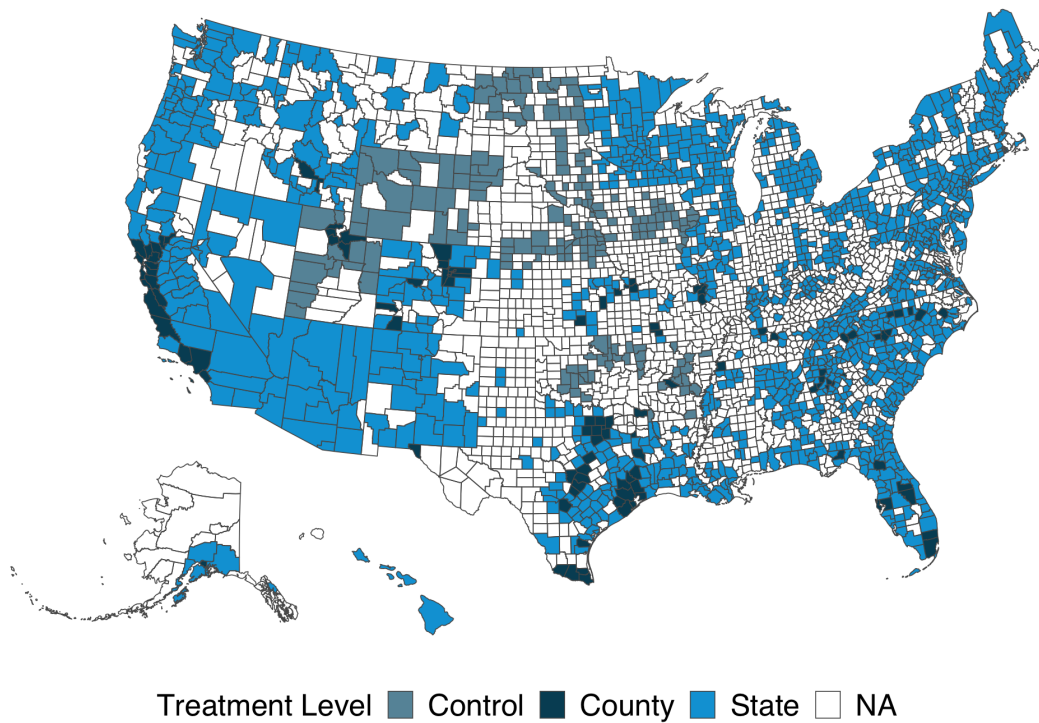
Note: This figure plots the number of new states that issued a stay-at-home order from March 18th through April 9th in Panel (a). Panel(b) plots the percent of the total US population under a stay-at-home order from March 9th through April 24th. Dates of stay-at-home orders are obtained from the *New York Times*.

Table 1.7: Balance across states

	March 18 (Pre-treatment)			April 23 (End of Sample)		
	(1) Treatment Mean/SD	(2) Control Mean/SD	(3) Difference (1) - (2)	(1) Treatment Mean/SD	(2) Control Mean/SD	(3) Difference (1) - (2)
Cases	208.95 (506.52)	29.38 (18.4)	179.58	19,644 (42,255)	2,256 (1,290)	17,388
Population (in millions)	7.20 (7.74)	2.18 (1.31)	5.03	7.20 (7.74)	2.18 (1.31)	5.03
Cases per 10,000 Residents	2.46 (3.25)	1.51 (0.76)	0.95	238 (279)	109 (51)	129
Businesses Open	-0.22 (0.08)	-0.14 (0.06)	-0.08	-0.46 (0.08)	-0.29 (0.07)	-0.17
Employment	-0.34 (0.1)	-0.23 (0.1)	-0.12	-0.55 (0.09)	-0.37 (0.1)	-0.18
Hours Worked	-0.34 (0.1)	-0.23 (0.1)	-0.11	-0.54 (0.1)	-0.35 (0.11)	-0.19
Expenditures (Total)	-0.03 (0.04)	0.00 (0.04)	-0.03	-0.20 (0.08)	-0.15 (0.08)	-0.05
Expenditures (Low Income)	-0.04 (0.09)	-0.01 (0.08)	-0.03	-0.17 (0.12)	-0.12 (0.12)	-0.04
Expenditures (Middle Income)	-0.03 (0.04)	0.01 (0.04)	-0.03	-0.20 (0.07)	-0.14 (0.08)	-0.06
Expenditures (High Income)	-0.04 (0.05)	-0.03 (0.06)	-0.01	-0.26 (0.06)	-0.19 (0.06)	-0.07
Number of observations	43	8	51	43	8	51

Note: This table presents means and standard deviations for outcomes of interest and other observable characteristics at the state level by treatment status for March 18th and April 23rd. Treatment is defined by the issuance of a statewide stay-at-home order. March 18th is the day before the issuance of the first statewide stay-at-home order. April 23rd is the day before the first lifting of a stay-at-home order. The third column reports the difference in means between the two groups. Low income expenditure is missing for Alaska and Vermont. High income expenditure is missing for Montana. This affects the number of observations for these variables.

Figure 1.14: Counties by Treatment Status



Note: This figure plots counties that remain in our final sample by their treatment status. Stay-at-home data at the county level is obtained from the National Association of Counties. Statewide orders are obtained from the *New York Times*. Counties that issued a stay-at-home order before their state are shaded in dark blue. Counties that did not issue a stay-at-home order but reside in a state that did are shaded in light blue. Counties that did not experience a county-level or state-level stay-at-home order are shaded in grey-blue. Counties missing from our combined dataset are not shaded and labelled as NA.

### 1.7.1.2 Difference-in-difference Estimates

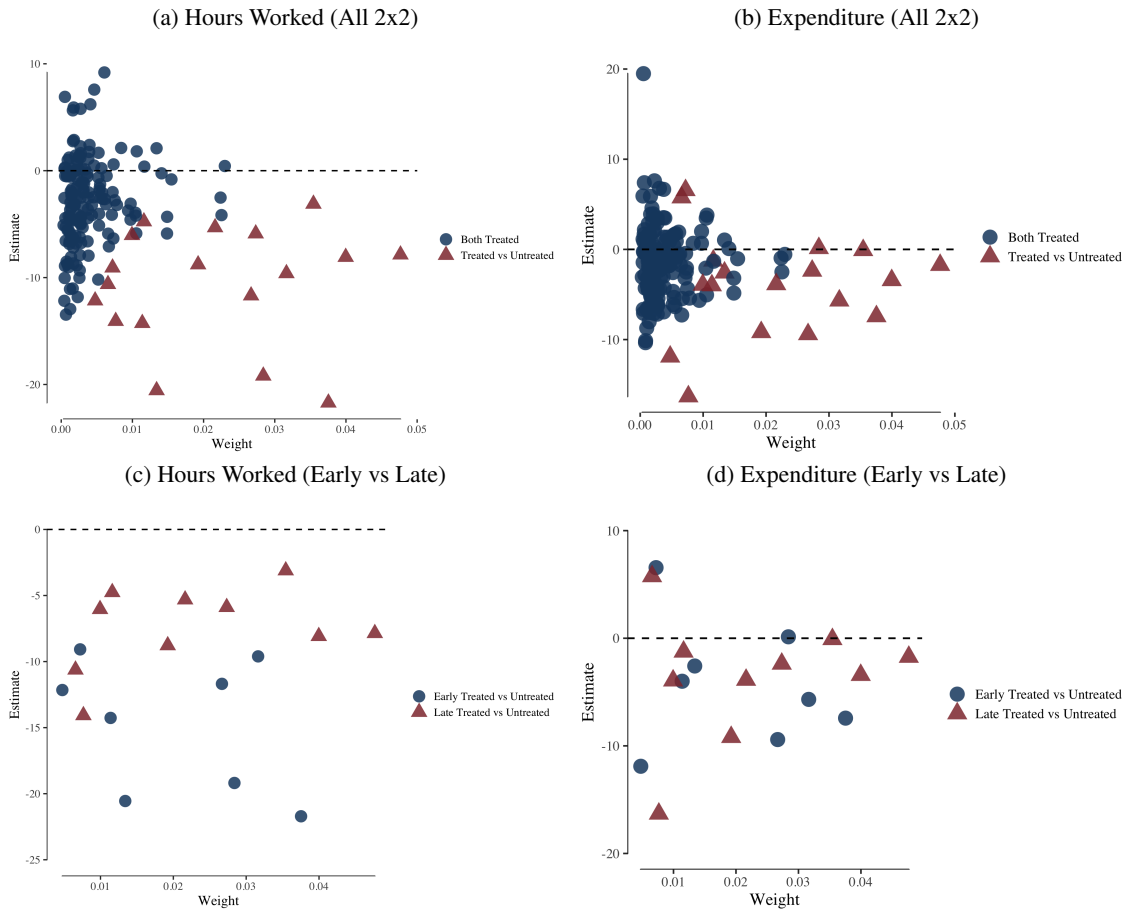
Table 1.8: Difference-in-difference Estimates

	Labor Variables (Per. Pt.)			Expenditure Variables (Per. Pt.)			
	Hours Worked (1)	Employment (2)	Businesses Open (3)	State Level (4)	Low Income (5)	Middle Income (6)	High Income (7)
Stay-Home Order	-3.954*** (0.394)	-4.009*** (0.386)	-4.719*** (0.359)	-1.869*** (0.253)	-1.964*** (0.370)	-1.959*** (0.277)	-1.913*** (0.266)
Business Closure	-3.103*** (0.411)	-2.983*** (0.406)	-3.690*** (0.383)	-0.216 (0.267)	-0.771* (0.411)	-0.460 (0.283)	-0.384 (0.271)
Observations	2,346	2,346	2,346	2,346	2,254	2,346	2,300

Note: This table presents difference-in-difference estimates for stay-at-home orders and non-essential business closures. Coefficients are jointly estimated in the same regression. Column names represent the dependent variable. Robust standard errors are reported. All observations are at the state-by-day level. Labor market and expenditure variables come from Homebase and Affinity Solutions, respectively.

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01

Figure 1.15: Goodman-Bacon (2018) Decomposition



Note: This figure plots the results of a Goodman-Bacon (2018) decomposition on our difference-in-difference estimate for hours worked in Panel A and expenditure in Panel B. The dependent variable is hours worked. The y-axis is the treatment-timing cohort DD estimate. The x-axis is the weight for the respective estimates. The top panel plots all 2x2 difference-in-difference coefficients. Estimates are divided into two groups: Both Treated (blue circles) and Treated vs Untreated (red triangles). In the bottom panel, only the treated vs untreated DD estimates are plotted. Estimates are further divided into two groups: Treatment occurred before March 29th (Early, blue circles), Treatment occurred on or after March 29th (Late, red triangles). Hours worked and expenditure data is obtained from Homebase and Affinity Solutions, respectively. Observations are at the state-by-day level.

### 1.7.1.3 Additional Event Study Estimates

Table 1.9: Effect of Stay-at-Home Orders on the Economy (Composition Effects)

Days	Hours Worked (Homebase, State, Per. Pt.)			Total Expenditure (Affinity Solutions, State, Per. Pt.)				
	Drop California (1)	Drop New York (2)	Pop. Weighted (3)	Drop California (4)	Drop New York (5)	Pop. Weighted (6)	Drop Alaska/Vermont (7)	Drop Montana (8)
-7 (t-value)	-1.17 (-1.1)	-0.82 (-0.77)	-0.57 (-0.64)	-0.92 (-1.43)	-0.92 (-1.43)	-0.63 (-1.09)	-0.91 (-1.53)	-0.9 (-1.44)
-2	0.89 (1.09)	0.9 (1.07)	0.91 (1.03)	-0.34 (-0.63)	-0.34 (-0.63)	-0.45 (-0.83)	-0.26 (-0.48)	-0.27 (-0.52)
0	-2.61* (-3.25)	-2.29* (-2.84)	-2.06* (-2.55)	-0.37 (-0.66)	-0.37 (-0.66)	-0.23 (-0.41)	-0.3 (-0.53)	-0.34 (-0.63)
5	-2.88* (-3.43)	-2.84* (-3.37)	-2.59* (-3.28)	-2.98* (-5.86)	-2.98* (-5.86)	-2.25* (-4.14)	-2.78* (-5.29)	-2.8* (-5.53)
10	-3.59* (-3.93)	-3.39* (-3.69)	-3.23* (-3.83)	-4.38* (-7.99)	-4.38* (-7.99)	-3.25* (-5.75)	-4.17* (-7.94)	-4.18* (-7.95)
15	-6.31* (-7.25)	-6.09* (-6.97)	-5.36* (-5.79)	-4.03* (-5.28)	-4.03* (-5.28)	-2.83* (-3.65)	-4* (-5.28)	-3.87* (-5.21)
Observations	2,300	2,300	2,346	2,300	2,300	2,346	2,254	2,300

Note: This table presents estimates of equation 1.1 for several specifications. All estimates are conducted at the state level. T-values are presented using robust standard errors for state estimates. Hours worked and total expenditure variables come from Homebase and Affinity Solutions, respectively. Specifications (1) and (4) drop California from the sample. Specifications (2) and (5) drop New York from the sample. Specifications (3) and (6) weight regressions by state population. Specification (7) shows that composition effects are not driving our results in Column (6) of Table 1.2. Similarly, specification (8) shows that Column (8) of Table 1.2 is not driven by composition effects.

\* denotes estimates significant at the 95% level.



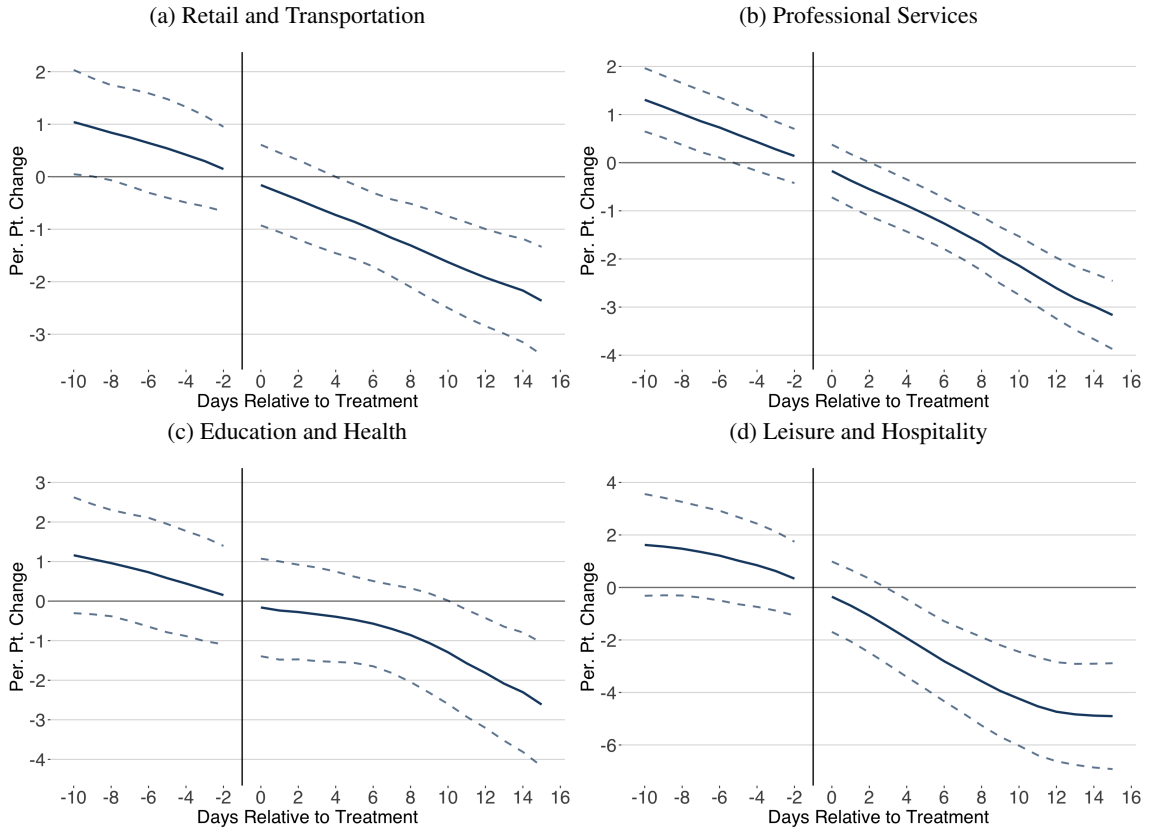
Table 1.10: Opportunity Insights Employment Measure (State vs County)

Days	State Employment (Per. Pt.)				County Employment (Per. Pt.)			
	Total (1)	Low (2)	Middle (3)	High (4)	Total (5)	Low (6)	Middle (7)	High (8)
-7 (t-value)	0.71 (1.55)	0.23 (0.33)	0.92 (1.83)	0.68 (1.8)	0.37 (0.94)	-0.04 (-0.06)	0.73* (2.27)	0.44 (1.36)
-2	0.15 (0.39)	0.09 (0.15)	0.19 (0.44)	0.12 (0.37)	0.1 (1.39)	0.04 (0.38)	0.15* (2.69)	0.08 (1.59)
0	-0.17 (-0.46)	-0.13 (-0.22)	-0.21 (-0.49)	-0.13 (-0.44)	-0.12 (-1.73)	-0.06 (-0.49)	-0.18* (-3.61)	-0.09 (-1.95)
5	-1.02* (-2.82)	-0.74 (-1.35)	-1.29* (-3.05)	-0.76* (-2.61)	-0.95* (-2.17)	-0.54 (-0.75)	-1.32* (-4.52)	-0.76* (-3.09)
10	-2.01* (-4.38)	-1.41* (-2.00)	-2.62* (-5.18)	-1.69* (-4.89)	-2.28* (-2.75)	-1.46 (-1.07)	-3.00* (-5.15)	-1.82* (-4.31)
15	-2.67* (-5.05)	-1.56 (-1.88)	-3.67* (-6.46)	-2.75* (-6.29)	-3.76* (-3.12)	-2.3 (-1.15)	-4.99* (-5.65)	-3.16* (-5.08)
Observations	2,346	2,300	2,300	2,346	34,937	23,912	26,558	16,415

Note: This table presents estimates of equation 1.1 for the Opportunity Insights employment data. Specifications (1)-(4) are estimated at the state level, and columns (5)-(8) at the county level. T-values are presented using robust standard errors for state estimates and standard errors clustered at the state-level for county estimates. Low-income workers make < \$27K per year. Middle-income workers make \$27K – \$60K. High-income workers make > \$60K per year.

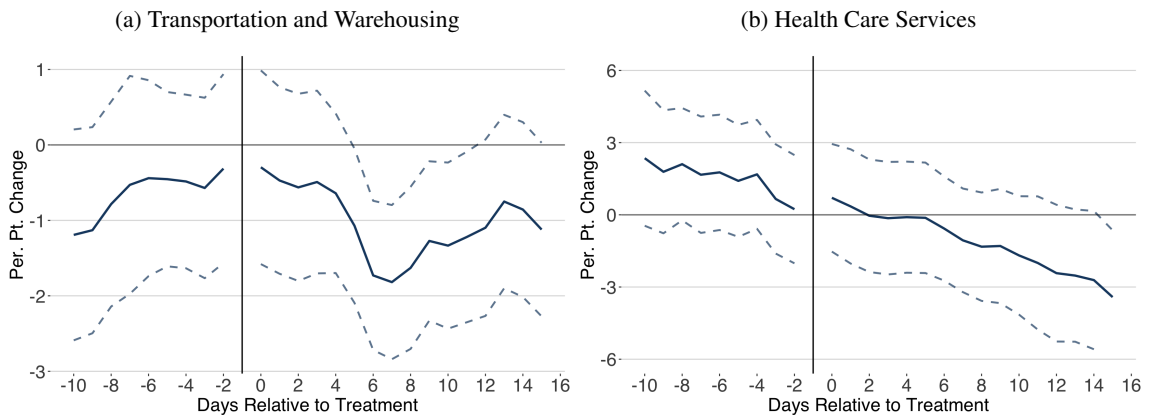
\* denotes estimates significant at the 95% level.

Figure 1.16: Opportunity Insights Employment Measure by NAICS Sector



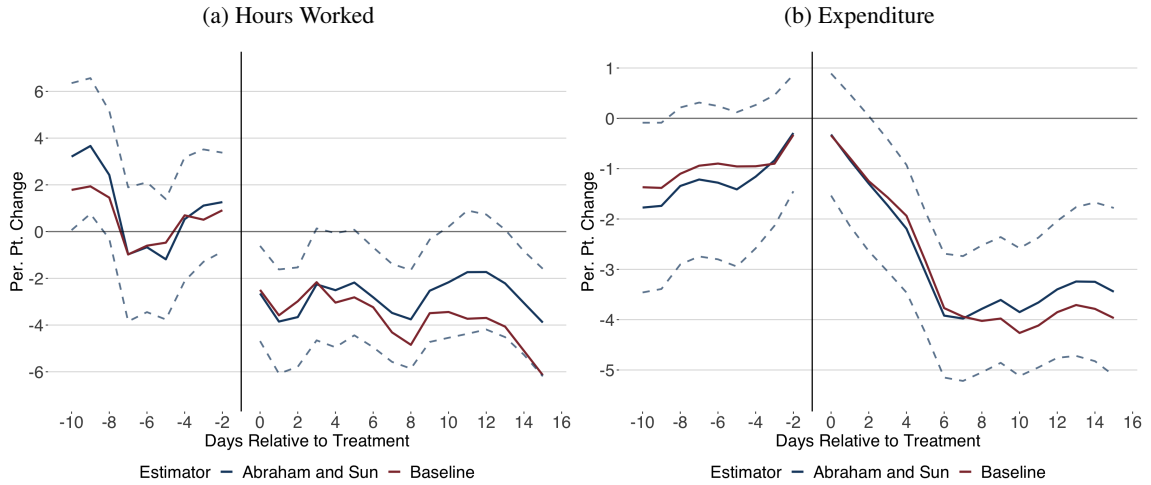
Note: This figure plots the results of our event study presented in Equation 1.1. Employment data is obtained from Opportunity Insights. Observations are at the state level. Dashed lines represent 95% confidence intervals using robust standard errors. Retail and Transportation is NAICS supersector 40, Professional Services (60), Education and Health (65), Leisure and Hospitality (70).

Figure 1.17: Heterogeneity across Consumption Expenditures (Extra)



Note: This figure plots the results of our event study presented in Equation 1.1. Expenditure data come from Affinity Solutions. Merchant groups are defined by Affinity Solutions. Observations are at the state level. Dashed lines represent 95% confidence intervals using robust standard errors.

Figure 1.18: Effect of Stay-at-Home Orders on the Economy (Abraham and Sun)



Note: This figure plots the stay-at-home order event study coefficients estimated using the method in Abraham and Sun (2020). Sub-captions represent the dependent variable. Dashed lines represent 95% confidence intervals using robust standard errors. The red line plots the baseline estimates for comparison. Observations are at the state-by-day level. Hours and expenditure data are obtained from Homebase and Affinity Solutions, respectively.

Table 1.11: Effect of Non-essential Business Closures on the Economy

Days	Labor Variables (Per. Pt.)			Expenditure Variables (Per. Pt.)				
	Hours Worked (1)	Employment (2)	Businesses Open (3)	State Level (4)	County Level (5)	Low Income (6)	Middle Income (7)	High Income (8)
-7 (t-value)	4.12* (3.23)	3.45* (2.94)	3.74* (3.28)	-0.25 (-0.34)	-0.26 (-0.37)	-1.71 (-1.56)	-0.22 (-0.27)	-0.31 (-0.41)
-2	0.96 (0.94)	0.71 (0.79)	1.11 (1.22)	-0.02 (-0.04)	-0.06 (-0.29)	-0.31 (-0.3)	-0.03 (-0.04)	-0.2 (-0.26)
0	-0.44 (-0.48)	-0.27 (-0.33)	-0.85 (-0.96)	-0.12 (-0.2)	0.06 (0.4)	-0.31 (-0.3)	-0.08 (-0.11)	-0.42 (-0.63)
5	-0.82 (-0.84)	-0.76 (-0.89)	-1.6 (-1.76)	-1.73* (-3)	-1.42* (-2.08)	-2.54* (-2.75)	-1.93* (-2.92)	-1.93* (-2.72)
10	-0.88 (-0.89)	-0.58 (-0.65)	-1.74 (-1.92)	-1.1 (-1.83)	-1.51 (-1.49)	-1.95 (-1.91)	-1.16 (-1.84)	-1.99* (-2.59)
15	0.33 (0.35)	0.14 (0.16)	-1.78 (-1.97)	0.06 (0.08)	0.01 (0.01)	-1.37 (-1.13)	-0.11 (-0.16)	-0.64 (-0.88)
Observations	2,346	2,346	2,346	2,346	73,451	2,254	2,346	2,300

Note: This table presents estimates of equation 1.1 for several specifications with non-essential business closures as the covariate of interest. All estimates are conducted at the state level except column (5) which provides county level estimates. T-values are presented using robust standard errors for state estimates and standard errors clustered at the state-level for county estimates. Labor and consumption variables come from Homebase and Affinity Solutions, respectively.

\* denotes estimates significant at the 95% level.

## 1.7.2 Model

### 1.7.2.1 Equilibrium Conditions

Our model can be summarized by the following first-order conditions:

### Susceptible

$$\frac{1}{C_t^s} - (1 + \mu_{ct})\lambda_{bt}^s + \lambda_{\pi_t} \pi_{s1}(I_t C_t) = 0 \quad (1.18)$$

$$-\theta \phi^s n_t^s + A \lambda_{bt}^s + \lambda_{\pi_t} \pi_{s2}(I_t N_t^I) = 0 \quad (1.19)$$

$$\beta(U_{t+1}^i - U_{t+1}^s) - \lambda_{\pi_t} = 0 \quad (1.20)$$

### Infected

$$\frac{1}{C_t^i} - (1 + \mu_{ct})\lambda_{bt}^i = 0 \quad (1.21)$$

$$-\theta \phi^i n_t^i + A \lambda_{bt}^i = 0 \quad (1.22)$$

### Recovered

$$\frac{1}{C_t^r} - (1 + \mu_{ct})\lambda_{bt}^r = 0 \quad (1.23)$$

$$-\theta \phi^r n_t^r + A \lambda_{bt}^r = 0 \quad (1.24)$$

### Market Clearing

$$S_t c_t^s + I_t c_t^i + R_t c_t^r = C_t \quad (1.25)$$

$$S_t n_t^s + I_t n_t^i + R_t n_t^r = N_t \quad (1.26)$$

#### 1.7.2.2 Parameters chosen to Optimize Model

In Section 1.5.2, we optimize  $(\tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3})$  to best fit the data. We chose these parameters as they summarize the endogenous labor responses by susceptible agents in the model without policy intervention. Recall that the probability of being infected is

$$\pi_{I,t} = \pi_{s1}(C_t^S)(I_t C_t^I) + \pi_{s2}(N_t^S)(I_t N_t^I) + \pi_{s3} \sum_{i=1}^N (1 - \tau_{H,i}^d)(I_{i,t} Pop_{i,0}) \quad (1.27)$$

Optimizing the parameters  $(\tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3})$  directly controls the probability of being infected. Furthermore, the equilibrium conditions show that the lagrange multiplier for this constraint is

$$\lambda_{\pi_t} = \beta(U_{t+1}^i - U_{t+1}^s) \quad (1.28)$$

In the steady-state, this can be re-written as

$$\lambda_{\pi_t} = \beta(U_{ss}^i - U_{ss}^s) \quad (1.29)$$

where

$$U_{ss}^s = \frac{1}{1 - \beta} u_{ss}^s \quad (1.30)$$

$$U_{ss}^i = \frac{1}{1 - \beta(1 - \pi_r - \pi_d)} (u_{ss}^i + \beta \pi_r u_{ss}^r) \quad (1.31)$$

However, in the steady-state,  $u_{ss}^s = u_{ss}^r$ . Substituting and simplifying, Equation 1.29 can be re-written as

$$\lambda_{\pi_t} = \left( \frac{\beta}{1 - \beta(1 - \pi_r - \pi_d)} u_{ss}^i + \frac{\beta^2 \pi_r}{1 - \beta(1 - \pi_r - \pi_d)} u_{ss}^s \right) - u_{ss}^s \quad (1.32)$$

$$\lambda_{\pi_t} = (u_{ss}^i - u_{ss}^s) + c \quad (1.33)$$

where  $c$  is a constant that depends on the discount factor and the probabilities of recovering and dying. Substituting for steady-state utility, we obtain

$$\lambda_{\pi_t} = [(\log(c_{ss}^i) - \frac{\theta}{2}(n_{ss}^i)^2) - (\log(c_{ss}^s) - \frac{\theta}{2}(n_{ss}^s)^2)] + c \quad (1.34)$$

$$\lambda_{\pi_t} = \log(c_{ss}^i) - \log(c_{ss}^s) + c \quad (1.35)$$

$$\lambda_{\pi_t} = \log\left(A\phi^i\left(\frac{1}{\theta}\right)^{1/2}\right) - \log\left(A\left(\frac{1}{\theta}\right)^{1/2}\right) + c \quad (1.36)$$

$$\lambda_{\pi_t} = \log(\phi^i) + c \quad (1.37)$$

Thus, the endogenous response by susceptible agents to decrease their probability of infection can be summarized by the loss in labor productivity and  $c$ . One could potentially treat  $\phi^i$  and  $\beta$ , which affects  $c$ , as free parameters. However, the other components of  $c$ ,  $\pi_r$  and  $\pi_d$  are set to match their empirical counterparts and should not be treated as free parameters.

### 1.7.2.3 Extended Model Calibration Procedure

Section 1.5.2.1 provided a procedure to calibrate our model in applied settings. This procedure focused solely on matching the time series of susceptibles in the model to that in the data. We extend this procedure to match both the time series of susceptibles and consumption in the data. The following steps are followed to generate simulations from the model:

1. Set Number of Locations ( $N$ )
2. Set Initial Population of Locations ( $Pop_0$ )
3. Set Susceptible Percent of Population ( $S_t$ )
4. Set New Infections ( $T_t$ )
5. Set Active Infections ( $I_t$ )
6. Set Parameters ( $\mu_{ct}, \tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3}$ )

The first five steps are identical to the procedure in Section 1.5.2.1. Step six now includes the consumption tax,  $\mu_{ct}$ , in the calibration procedure. This increases the number of calibrated parameters by  $N * T$ , where  $N$  is the number of locations and  $T$  is the number of time periods in the application. Due to the increase of parameters, we use the following four step procedure for calibration:

**Step One** Jointly set  $(\tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3})$  to match the empirical time series of susceptibles.

$$\min_{\tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3}} \sum_{i=1}^4 \sum_{t=May6}^{Aug26} (S_{i,t}^{Data} - S_{i,t}^{Model})^2 \quad (1.38)$$

**Step Two** Jointly set  $\mu_{ct}$  to match the empirical time series of consumption holding the calibrated transmission parameters fixed  $(\tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3})$  fixed.

$$\min_{\mu_{ct}} \sum_{i=1}^4 \sum_{t=May6}^{Aug26} (C_{i,t}^{Data} - C_{i,t}^{Model})^2 \quad (1.39)$$

**Step Three** Repeat Step One holding the calibrated consumption taxes,  $\mu_{ct}$ , fixed.

**Step Four** Jointly set  $(\mu_{ct}, \tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3})$  to match both the inverse-variance weighted time series of susceptibles and consumption.

$$\min_{\mu_{ct}, \tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3}} \frac{\sum_{i=1}^4 \sum_{t=May6}^{Aug26} (S_{i,t}^{Data} - S_{i,t}^{Model})^2}{var(S_{i,t}^{Data})} + \frac{\sum_{i=1}^4 \sum_{t=May6}^{Aug26} (C_{i,t}^{Data} - C_{i,t}^{Model})^2}{var(C_{i,t}^{Data})} \quad (1.40)$$

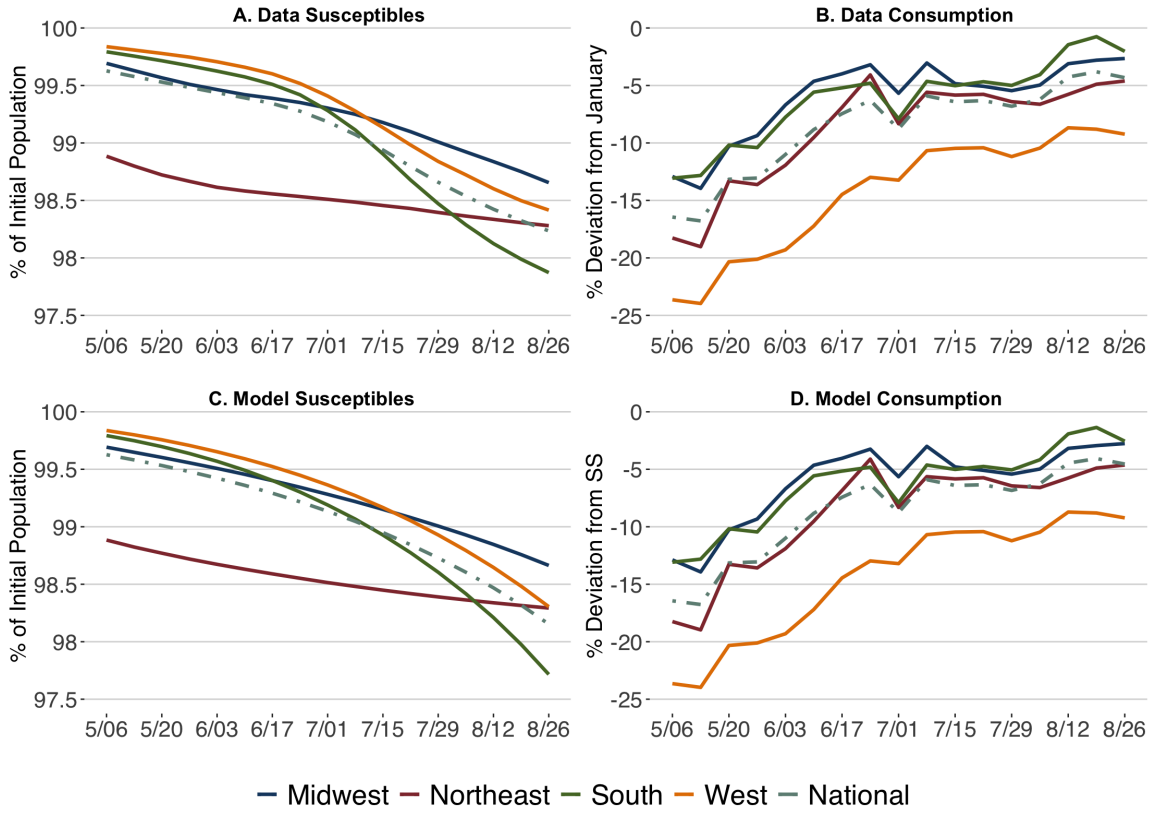
subject to

$$(\mu_{ct}, \tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3}) \geq 0.9 * (\mu_{ct}, \tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3}) \quad (1.41)$$

$$(\mu_{ct}, \tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3}) \leq 1.1 * (\mu_{ct}, \tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3}) \quad (1.42)$$

As in Section 1.5.2.1, we impose the restrictions on travel cost:  $\tau_{H,i}^d \in [0, 1]$  and  $\tau_{H,i}^d = 0$  for  $H = i$  and that the virus transmission parameters are positive  $\pi_{s1}, \pi_{s2}, \pi_{s3} \geq 0$  for all steps in the calibration procedure. Figure 1.19 plots both the time series for susceptibles and consumption in the data and the model after following this extended procedure. Overall, the results are similar to those in Figure 1.10 with the addition of matching the time series of consumption in the data.

Figure 1.19: Extended US Census Region Application



Note: Panel A plots the observed time series of susceptibles in the United States and each Census region from May 6 through August 26. Panel B the observed consumption paths. Panels C and D plot the equivalent model fits for susceptibles and consumption, respectively, after jointly optimizing the consumption taxes, the travel distance across locations, and the transmission parameters  $(\mu_{ct}, \tau_{H,i}^d, \pi_{s1}, \pi_{s2}, \pi_{s3})$ . Infection data is obtained from *USA Facts*. Consumption data is obtained from *Affinity Solutions*.

#### 1.7.2.4 Additional Comparative Statics

Figure 1.20 plots the results of a three-location SIR model under different parameterizations. Columns represent different initial conditions. The first column uses initial populations of  $\frac{1}{3}$  and initial infections of 0.1% for all states. The second column uses the initial population across states,  $Pop_0 = (36.6\%, 56.5\%, 6.9\%)$ , while maintaining the same initial infections. The third column randomizes the initial infections of locations,  $T_0 = (0.16\%, 0.05\%, 0.16\%)$ , using the randomized populations from the second column.<sup>42</sup> Rows represent the travel cost used. We require that aggregate initial population and infections are equal 1 and 0.1%, re-


<sup>42</sup>These population and infection rates were determined using a random draw. The random draws are then scaled to maintain the same initial aggregate conditions as our baseline simulation. The actual dynamics will depend on the specific random draw, but will not change the qualitative comparative static results.

spectively, for all simulations. Thus, differences in the aggregate must be attributed to the cross-sectional interactions rather than initial aggregate conditions. All other parameters remain the same as in Table 1.4.


Figure 1.20: Comparative Statics (Susceptibles)

**1. No Transmission,  $\tau_{H,i}^d = 1$  for  $H \neq i$**


(a) Pop = (33.3%, 33.3%, 33.3%)    (b) Pop = (36.6%, 56.5%, 6.9%)    (c) Pop = (36.6%, 56.5%, 6.9%)  
 Infections = (0.1%, 0.1%, 0.1%)    Infections = (0.1%, 0.1%, 0.1%)    Infections = (0.16%, 0.05%, 0.16%)



**2. Partial Transmission,  $\tau_{H,i}^d = \frac{|H-i|}{N}$**



**3. Full Transmission,  $\tau_{H,i}^d = 0$**



Note: This figure plots the path of susceptibles using a 3-state SIR macro model. In all specifications, aggregate initial infections are equal to 0.001. Columns represent different initial conditions. The first column presents results when locations are symmetric. The second column randomizes the population of locations. The third column randomizes the initial infections of locations using the populations from the second column. Rows represent the travel cost used. The first row plots the results using no interaction across locations. The second row uses  $\tau_{H,i}^d = \frac{|H-i|}{N}$ . The third row plots the results using no travel cost across locations.

**No Transmission across Locations.** The first row plots the results using no interaction across locations,  $\tau_{H,i}^d = 1$  for  $H \neq i$ . Panel (a) shows that, with symmetric locations and no transmission across locations, all locations have the same time series for the fraction of susceptible agents. Panel (b) shows that virus

transmission within a state is increasing in its population. The intuition is that individuals are more likely to catch the virus from a larger number of people in more populous states. Panel (c) shows that different initial infection conditions can create crossing patterns in the cross-section. For these particular randomized population and initial infection draws, we see that State 2 takes about twenty and forty weeks to surpass States 3 and 1, respectively, in cumulative infections. This differs from Panel (b) where states with higher population always experience higher cumulative infections than lower populated states.

**Partial Transmission.** The second row simulates the same scenarios using  $\tau_{H,i}^d = \frac{|H-i|}{N}$ . Comparing Panels (1a) and (2a) show that asymmetric travel costs across locations can lead to cross-sectional heterogeneity of outcomes even when initial population and infections are symmetric. Panel (2b) shows that, as in Panel (1b), initial population affects both the cross-sectional and aggregate time series. Panel (2c) shows that allowing for virus transmission across locations mitigates, but does not eliminate, the effect that initial infections has on the cross-sectional crossing patterns.

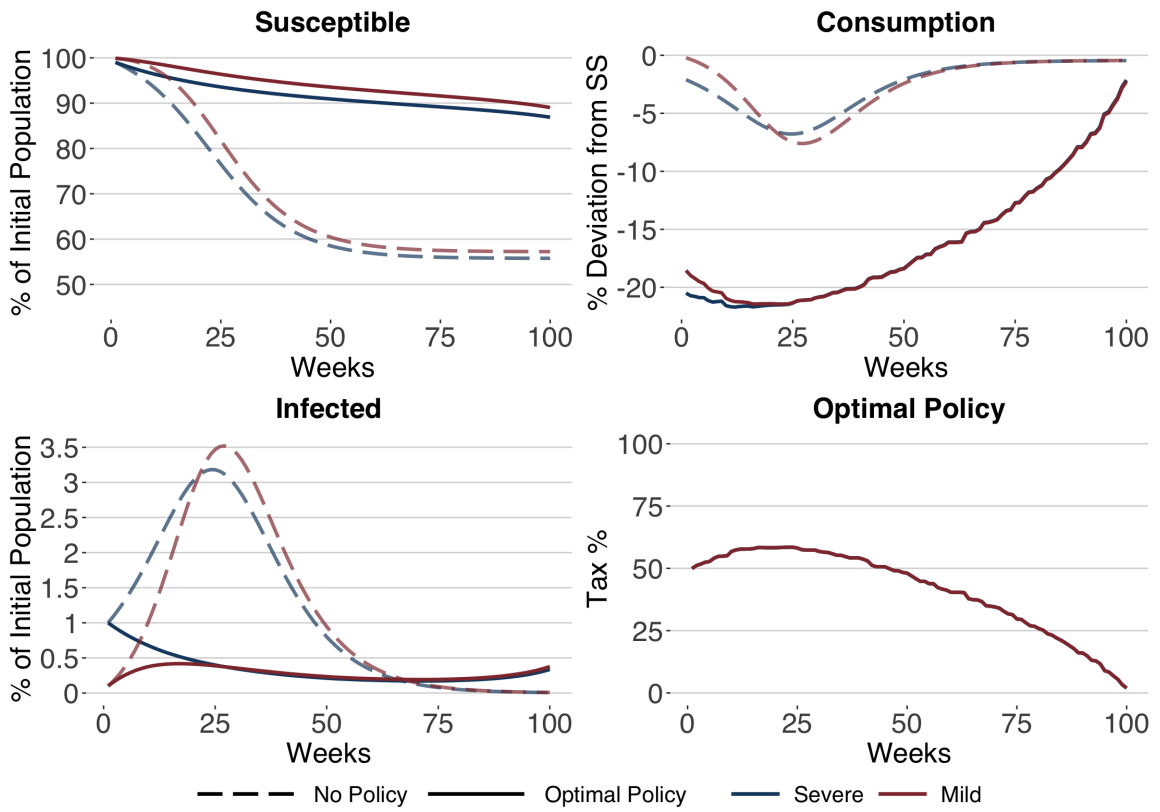
**Full Transmission.** The third row plots the results using no travel cost across locations,  $\tau_{H,i}^d = 0$ . Panel (3a) shows that when there is full transmission of the virus across locations that our model replicates Eichenbaum et al. (2020) when locations are symmetric. This is also the case when locations have different populations in Panel (3b). Our transmission mechanism in Equation 1.3 shows that idiosyncratic initial infection rates are needed to generate a departure from Eichenbaum et al. (2020) when travel cost is zero. However, Panel (3c) shows that the consequences of idiosyncratic infection rates are small when there is complete virus pass-through across locations even with asymmetric population sizes.

Overall, Panels (1a)-(1c) show that both initial population and initial infection rates can have significant consequences for the dynamics of our model. Comparing these panels to their counterparts in Panels (2) and (3) show that the extent of these consequences depends heavily on the modelling of virus transmission across locations. In the limiting case of no travel cost across locations, our model converges to the single-location model of Eichenbaum et al. (2020) regardless of initial population and infection symmetry. Figure 1.20 also shows that cross-sectional assumptions have important implications for aggregate dynamics. In the most severe cases, over 50% of the total population becomes infected after 100 weeks as seen in Panel 3 with full transmission across locations. However, Panel (1a) suggests that less than 20% of the nation becomes infected when there is no transmission across locations. Our results suggest that our model is not only important for modelling cross-sectional behavior, but also aggregate dynamics even when holding aggregate initial conditions fixed.



### 1.7.2.5 National Mitigation Policy

Figure 1.21: Optimal National Policy



Note: This figure plots the results of a 2-location SIR macro model restricted to a common, national policy. Both locations have an initial population of  $\frac{1}{2}$ . Initial infections to 1% and 0.1% for the severe (blue) and mild (red) infected locations, respectively.

## CHAPTER 2

### Selection Effects in Retail Chain Pricing

#### 2.1 Introduction

Selection effects—the concept that firms optimally choose the timing of their price changes—are key determinants of how the economy responds to aggregate nominal shocks. As selection effects increase, firms respond more to aggregate nominal shocks which mitigates their impact. Consequently, the degree of monetary non-neutrality is not directly linked to micro-level price stickiness as non-neutrality also depends on *how* firms determine to change their price.

This paper examines selection effects in the context of retail chain pricing using a menu cost model. It is commonly assumed in menu cost models that firms set their prices autonomously.<sup>1</sup> However, stores belonging to the same retail chain (e.g. Kroger, Publix) set nearly identical prices in practice. Additionally, retail chains synchronize the timing and magnitude of their price changes across stores violating the independence assumption.

I find that accounting for retail chains more than doubles the degree of selection effects relative to the standard menu cost model. This relationship suggests that the standard menu cost model overestimates the degree of monetary non-neutrality by ignoring synchronization in retail chain pricing. The intuition behind this result lies in the store’s pricing decision. In the standard model, stores decide whether or not to change their price contingent on their idiosyncratic productivity and their current price relative to the aggregate price level. Firms then change their price if the expected profit gains outweigh the menu cost.

Introducing retail chains to the model adds another component to the store pricing decision. First, store-level productivity shocks include a common retail chain component. This common component helps account for the price synchronization seen in the data for stores belonging to the same chain. I also assume that the retail chain occasionally sets the store’s price. However, when determining the store’s price, the retailer cannot observe the store’s productivity and thus chooses the store’s price based on the retailer-level productivity. The combination of these two assumptions leads to the amplified selection effect in the retail chain model.

Pricing in this environment can be interpreted as a two-step process. First, the retailer sets a target price for all stores belonging to its chain using chain-level state variables. Second, stores choose to keep the retail chain price or set their own price. The assumptions above serve as a reduced form modelling approach for the cost that a store pays to deviate from the chain price. Thus, firms in the retail chain model face a “constrained” optimization problem whereas firms in the standard model are “unconstrained”. This constraint causes stores

---

<sup>1</sup>See Sheshinski and Weiss (1977); Golosov and Lucas (2007); Nakamura and Steinsson (2008) for examples of pricing in standard menu cost models.

to be less responsive to their idiosyncratic productivity shocks (as it is more costly) and more responsive to aggregate shocks.

I estimate selection effects by simulating both the standard model and retail chain model calibrated to scanner-level data for over 600 goods sold in the United States from 2001-2007. Data come from Information Resources Inc. (IRI) which records the weekly revenue and quantity sold for each product at the store level. Importantly, the dataset records the retail chain to which each store belongs.

Before calibrating the model, I document descriptive evidence of retail chain price synchronization. I begin by documenting synchronization in the timing of price changes. I find that if one store in a retail chain changes their price for a good in a given week, there is a 70% probability that at least half of stores within that chain change their price during the same week (averaged over goods and chains). The conditional probability remains over 40% with the increased restriction that all stores within the chain change prices. Requiring that all price changes be in the same direction yields an estimate of about 30% for both price increases and decreases. I also find that these price changes are similar in magnitude. A variance decomposition shows that 62% of price dispersion from a store's average price can be explained by chain-week fixed effects on average over goods while the idiosyncratic store component accounts for only 25% of price variation. This result is prevalent across most goods with the chain-week component explaining at least half of the price variation for 625 of the 655 goods in the sample.

The retail chain model is then calibrated to match several price-setting statistics including the chain-week component of the variance decomposition for each good. Thus, the empirical variance decomposition serves as a key factor in determining the probability that a retail chain sets an individual store's price in the model (i.e. the degree of the retail chain "constraint"). I then simulate store-level price paths for each of the 655 goods using the retail chain model. After simulating the retail chain model, I perform a counterfactual analysis where stores do not adhere to the retail chain constraint. Each store faces the same inflation process and analogous productivity process as in the retail-chain simulation. Thus, differences between the models are not driven by different productivity draws and are solely driven by each store selecting its unconstrained optimal price in each period.

I measure selection by regressing the change in a store's log price on the change in the aggregate price level controlling for changes in store-level productivity. The partial equilibrium nature of the model facilitates this regression as store-level price changes do not feedback into the aggregate price level which follows an exogenous process. The coefficient on the aggregate price level thus serves as an estimate of the responsiveness of store-level prices to exogenous changes in the aggregate price level (i.e. the selection effect). The weighted mean estimates for the retail chain model and standard model are 0.25 and 0.09, respectively. These estimates suggest that a 1% increase in the aggregate price level leads to selection effects that are more than

twice as large in the presence of retail-chains compared to the standard model (0.25% versus 0.09%).

These results build on our knowledge of selection effects in sticky-price models. The Calvo (1983) model of price adjustments represents one extreme in sticky price models. In the Calvo model, a subset of firms are selected at random each period to change their price. Thus, as noted in Nakamura and Steinsson (2013), firms cannot optimally time their price changes, and aggregate shocks have no effect on how many and which firms change their price. Caballero and Engel (2007) illustrate that aggregate shocks only affect the intensive margin of price adjustment in the Calvo model (i.e. only the magnitude of price changes is affected for firms that were already going to adjust their price) which leads to large real effects of monetary shocks.

Menu cost models such as the one in this paper introduce an extensive margin of price adjustment (Caballero and Engel, 2007). Thus, aggregate shocks also affect how many and which firms change their price referred to as the selection effect in Golosov and Lucas (2007). In general, the extensive margin of adjustment leads nominal shocks to have smaller real effects on the economy compared to the Calvo model. For example, nominal shocks only produce 20% of the real effects in Golosov and Lucas (2007) compared to the Calvo model. However, the extent of the real effects varies significantly depending on the modelling assumptions used. Golosov and Lucas (2007) commonly serves as the lower bound where additional assumptions typically predict effects closer to those in the Calvo model (Leptokurtic shocks and scale economies: Midrigan (2011), Alvarez and Lippi (2014), Bonomo et al. (2020); Random menu costs: Dotsey et al. (1999); Sectoral heterogeneity: Nakamura and Steinsson (2010)).

In contrast to the examples above, my results suggest that retail chain pricing decreases the real effects of nominal shocks. Nakamura and Steinsson (2010) help illustrate the intuition for this result in menu cost models. They show, conditional on the same frequency of price change, reducing the variance of stores' idiosyncratic productivity shocks leads to less real effects of nominal shocks. This is a result of the average inflation rate becoming a more important determinant in stores' pricing decisions. Golosov and Lucas (2007) also illustrate this result and show that their setup converges to Caplin and Spulber (1987) in the absence of idiosyncratic shocks. The Caplin and Spulber (1987) model represents the opposite extreme of the Calvo model in which the aggregate price level is completely flexible even in the presence of micro price stickiness. Thus, nominal rigidities have no real effect on the economy. Although the partial equilibrium nature of the model does not allow me to directly estimate the real effects of nominal shocks, I show that my regression specification captures the relationship described by Nakamura and Steinsson (2010) and Golosov and Lucas (2007). This relationship suggests that the standard menu cost model overestimates the degree of monetary non-neutrality by ignoring synchronization in retail chain pricing as illustrated by my regression results.

This paper also builds on the literature of retail chain pricing. DellaVigna and Gentzkow (2019) and Adams and Williams (2019) show that chains often set prices uniformly across stores or according to a small

set of retail pricing zones. Similarly, online and offline prices for a given retailer tend to be synchronized (Cavallo, 2017, 2019). Previous studies have found that chains predominantly account for price-level differences and variation in the frequency of price changes (Daruich and Kozlowski, 2021; Nakamura et al., 2011). The variance decomposition in this paper is most closely related to that in Nakamura (2008) which focuses on time series variation in prices and finds similar results.

Lastly, this paper helps provide a rationale for the stickiness of local prices to local conditions (Daruich and Kozlowski, 2021; Gagnon and López-Salido, 2020; Della Vigna and Gentzkow, 2019). This paper is most closely related to Daruich and Kozlowski (2021) who develop a model of multi-region firms with uniform pricing. Their paper finds that local elasticities are likely biased estimates of aggregate elasticities when accounting for uniform pricing in retail chains. My paper does not explicitly examine the impact of retail chain pricing on regional versus aggregate shocks. Instead, my paper highlights the degree of selection effects with versus without synchronization in retail chain pricing. As retail chain synchronization increases, stores respond less to their idiosyncratic shocks and respond more to aggregate shocks which helps rationalize the effects found in Gagnon and López-Salido (2020).

The rest of the paper proceeds as follows. Section 2.2 describes the IRI scanner dataset. Section 2.3 presents descriptive statistics of price synchronization within retail chains. Section 2.4 introduces the standard menu cost model and the retail chain extension. Section 2.5 uses model simulations to estimate selection effects. Section 2.6 concludes.

## 2.2 Data

My primary analysis uses the Information Resources Inc. (IRI) retail scanner dataset from 2001-2007. The IRI records the total weekly revenue and quantity sold for over 100,000 products and 3,000 stores.<sup>2</sup> Products are defined by their Universal Product Code (UPC), and stores are defined as a key provided by the IRI. The average price of product  $i$  in store  $j$  for the week  $t$  is computed as the total revenue ( $Rev_{ijt}$ ) divided by total quantity sold ( $Q_{ijt}$ ),  $P_{ijt} = \frac{Rev_{ijt}}{Q_{ijt}}$ .

### 2.2.1 Sample Formation

#### 2.2.1.1 Stores

The IRI provides data for both grocery and drug stores. My primary sample consists only of grocery stores which are a majority of the dataset. Each grocery store may belong to a specific retail chain (e.g. Kroger, Publix). In order to analyze the effect of retail chain pricing, I require that stores (1) belong to a retail chain, (2) do not switch chains over time, (3) and are open for more than one year.

---

<sup>2</sup>A complete description of the dataset can be found in Bronnenberg et al. (2008).

Table 2.1: Summary Statistics of IRI Data

<b>Panel A: Store Requirements</b>					
	All Stores	Grocery Stores	Belong to Chain	Do not Switch Chain	In Sample > 1 Year
Number of Stores	3,150	2,378	2,367	1,534	1,234
Number of Chains	147	128	127	119	101
<b>Panel B: Sample Formation</b>					
	Stores	Chains	Products	Categories	Observations
Initial Sample	3,150	147	105,929	31	1,367,985,544
Store Requirements	1,234	101	91,102	31	656,570,883
Product Sold by at least half of Chains	1,184	99	7,342	31	360,875,433
Product Sold for at least half of store-weeks	1,123	99	655	29	60,698,877

Note: This table presents summary statistics of the IRI data. Panel A presents the effects of the store-level requirements on the total number of stores and chains in the sample. Panel B presents total counts for both stores and products throughout the complete sample formation.

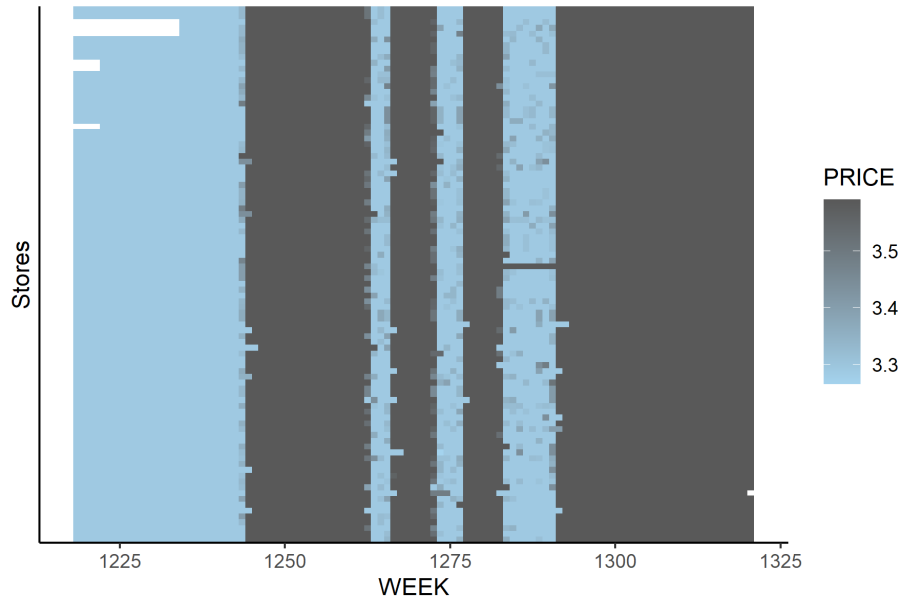
Panel A of Table 2.1 highlights the effects of these requirements. The first restriction facilitates the variance decomposition presented in the next section. Only 11 of the over 2,000 grocery stores in the dataset do not belong to a retail chain. The second restriction helps account for pricing behavior related to retail chain switching and reduces the number of stores in the sample by one-third. DellaVigna and Gentzkow (2019) conduct an event-study analysis and show that pricing behavior shifts substantially when stores switch chains which could bias the results of the variance decomposition. The final requirement that stores remain open for more than one year helps to remove pricing decisions that may be related to the opening and closing of a particular store. Overall, the store requirements have a strong effect on the total number of stores which reduce from 3,150 to 1,234. This effect is more modest for the number of retail chains which reduces from 147 to 101.

### 2.2.1.2 Products

I further refine the sample through several product restrictions. I require that products (1) are carried by at least half of the chains and (2) are sold for at least half of store-weeks in a given year. These requirements help focus the sample on a set of widely available and commonly sold products. They also help avoid retail chain specific products.

Panel B of Table 2.1 documents the total effect of both the store and product requirements on the final sample of observations. The initial sample contained information for over 100,000 products. The store requirements had a relatively small effect on the total number of products (91,102) compared to the product requirements which reduced the final sample to 655 goods. The product requirements also removed about 100 stores. The final sample contains 1,123 stores belonging to 99 retail chains which sell 655 products

Figure 2.1: Example of Retail Chain Price Synchronization



Note: This figure plots an example of price synchronization within a retail chain. Each point on the y-axis represents an individual store belonging to the retail chain. Darker (lighter) shades represent higher (lower) prices. Missing values are represented by white space.

across 29 categories.

### 2.2.2 Retail Chain Pricing

Importantly, the IRI dataset records the retail chain  $r$  that each store  $j$  belongs to. This allows me to document several stylized facts about retail chain pricing. The first fact is that retail chains implement uniform pricing—the phenomenon that stores belonging to the same retail chain set nearly identical prices regardless of their respective market characteristics. Furthermore, the timing and magnitude of their price changes are often identical.

Figure 2.1 graphs an example of uniform pricing for a product belonging to the sugar/sugar substitute category within one retail chain. Each point on the y-axis represents an individual store belonging to the retail chain. Darker (lighter) shades represent higher (lower) prices. Missing values are represented by white space.

We see that there is small (or zero) price variation across stores for most weeks. Stores seldom make idiosyncratic price changes with most price changes occurring across all stores within the chain. Furthermore, prices change by the similar magnitudes across all stores. Appendix Figure 2.5 illustrates that this pricing pattern is not limited to this specific good, nor the sugar/sugar substitute category.<sup>3</sup>

<sup>3</sup>Alternatively, see DellaVigna and Gentzkow (2019) for similar pricing patterns found in the Nielsen scanner dataset across various product categories.

Table 2.2: Synchronization of Price Changes

	Any Change	Price Increases	Price Decreases
More than One Store	69.2%	60.8%	62.4%
At least Half of Stores	68.7%	54.1%	55.6%
All but One Store	56.4%	43.9%	44.9%
All Stores	42.9%	30.8%	32.1%

Note: This table presents statistics for the synchronization of price changes within a retail chain. All statistics are conditional on at least one price change in a store. The frequency represents the percent of weeks in which the specific row occurred (averaged over goods and chains). The interpretation of the first column in the last row is conditional on at least one store within a retailer changing its price, all stores belonging to that chain changed their price 43% of the time. The second and third columns document whether these price changes are in the same direction. Thus, conditional on a price increase/decrease within a chain, all stores change their price 30.8%/32.1% of the time.

### 2.3 Price Synchronization within Retail Chains

Figure 2.1 and Appendix Figure 2.5 helped provide illustrative examples of price synchronization within retail chains across various products and product categories. This section attempts to quantify these illustrative examples. The first subsection documents stylized facts on the timing of price changes. The second subsection conducts a variance decomposition to help analyze both the direction and magnitude of price changes.

#### 2.3.1 Price Synchronization

Table 2.2 presents statistics for the synchronization in the timing of price changes within a retail chain. Statistics in the first column are conditional on at least one price change in a store.<sup>4</sup> The frequency represents the percent of weeks in which the specific row occurred (averaged over goods and chains). The interpretation of the first row in the first column is conditional on at least one store within a retailer changing its price, more than one store belonging to that chain changed their price 69.2% of the time. Increasing the restriction that at least half of stores in the chain change their price has a negligible effect on the frequency. We see that, conditional on a price change, all stores in that chain change their price 42.9% of weeks. This is about a 25 percentage point reduction compared to the frequency for at least half of stores. However, much of this reduction can be accounted for by one store not changing its price.

The second and third columns of Table 2.2 serve as a check that these price changes are in the same direction. We see that requiring price changes to be in the same direction reduces the frequency of synchronized changes by about 10 percentage points across all specifications. Conditional on a price increase in one store, more than one store in the same chain increases its price 60.8% of the time. The frequencies for all stores in a chain increasing/decreasing their price in the same period are 30.8% and 32.1%, respectively. Differences

<sup>4</sup>I follow a similar procedure as Hee Hong et al. (2021) and Alvarez et al. (2016) in limiting a price change to be at least one cent and less the infinity. Recall, posted prices are not provided by the IRI, and instead I calculate price as  $P_{ijt} = \frac{Rev_{ijt}}{Q_{ijt}}$ . Thus, fractional price changes may occur due to this method. Eliminating infinite price changes further helps account for measurement error.



in the conditional price increase and price decrease frequencies are less than 2 percentage points across all specifications.

### 2.3.2 Variance Decomposition

Table 2.2 suggests that retailers highly synchronize the timing of their price changes across stores. These price changes are also typically in the same direction. However, the statistics presented do not provide information on the magnitude of these price changes. Although all stores in a chain change their price in the same direction over 30% of the time, the size of price changes may vary significantly across stores. To account for this, I conduct a variance decomposition of store's relative prices. This variance decomposition also helps account for the feature shown in Figure 2.1 where individual stores are often a week early or late to update their price to the retail chain price which can bias the previous price change statistics downward.

I begin by modelling the log price ( $p_{j,r,t}$ ) for store  $j$  belonging to retail chain  $r$  in week  $t$  as

$$p_{j,r,t} = \alpha_t + \delta_j + \gamma_{r,t} + \epsilon_{j,r,t} \quad (2.1)$$

where  $\alpha_t$  is a week fixed effect,  $\delta_j$  is a store fixed effect,  $\gamma_{r,t}$  is a chain-by-week fixed effect, and  $\epsilon_{j,r,t}$  is the residual. This equation is estimated separately for each good which eliminates the need to include a product component.<sup>5</sup> Using these estimated parameters, I perform the following variance decomposition for each good:

$$\text{Var}(p_{j,r,t} - \hat{\delta}_j) = \text{Var}(\hat{\alpha}_t) + \text{Var}(\hat{\gamma}_{r,t}) + \text{Var}(\hat{\epsilon}_{j,r,t}) \quad (2.2)$$

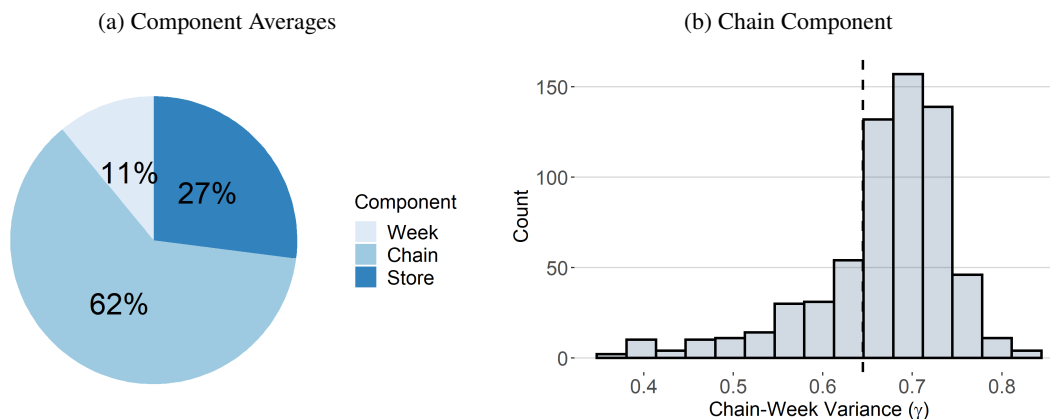
I normalize  $p_{j,r,t}$  by  $\hat{\delta}_j$ , the average price in a store, in order to analyze price variation over time rather than constant differences in the average price across chains.<sup>6</sup> Thus, the total variance of the relative price is decomposed into three components: price changes that occur across all stores in the same week regardless of the retail chain ( $\hat{\alpha}_t$ ), price changes that occur across all stores in the same week within a retail chain ( $\hat{\gamma}_{r,t}$ ), and individual store-level price changes ( $\hat{\epsilon}_{j,r,t}$ ).

Panel (a) of Figure 2.2 presents the weighted averages from the variance decomposition. The week component  $\alpha_t$  estimate of 11% suggests that price changes are not highly correlated across chains. However, the chain-week component  $\gamma_{r,t}$  explains 62% of price variation on average. This suggests that price changes and their magnitude are highly correlated for stores within the same chain. Although the results suggest that individual store managers do maintain some flexibility when changing their price with the store component

<sup>5</sup>Estimation follows a similar procedure used in Daruich and Kozlowski (2021) and Kaplan et al. (2016). See Appendix 2.7.1.1 for more details.

<sup>6</sup>The variance decomposition for relative price variation over time and constant differences in the average price across chains need not necessarily be the same. As detailed in Crucini and Telmer (2020), the decomposition for relative price variation over time is the relevant statistic for business cycle models such as the one in this paper.

Figure 2.2: Variance Decomposition



Note: This figure presents the results of the variance decomposition in Equation (2.2). The variance decomposition is conducted separately for each good. Panel (a) presents the mean estimate for each component of the decomposition. Panel (b) plots the distribution of the chain-week variance ( $\gamma$ ) over all goods. The vertical dashed line represents the cutoff for the first quartile.

$\varepsilon_{j,r,t}$  explaining about one quarter of relative price dispersion.

Panel (b) plots the distribution of the estimated chain component for all goods. The vertical dashed line represents the cutoff for the first quartile at 64.5%. The minimum variation explained by the chain-week component is 35.7%. Overall, the chain-week component can explain at least half of the price variation for about 625 out of the 655 goods that the decomposition was computed for.

### 2.3.2.1 Robustness Checks

I perform several robustness checks of the above variance decomposition. As the behavior of temporary sale prices often behave differently than regular prices (Eichenbaum et al., 2011; Anderson et al., 2017), I replace temporary sale prices with the previous regular/non-sale price. I follow the method in Eden et al. (2021) of defining a temporary sale price as a 10% drop in price followed by a price equal to or above the pre-sale price within four weeks.<sup>7,8</sup> Second, I estimate the variance decomposition using monthly rather than weekly prices.<sup>9</sup> Lastly, I perform the variance decomposition with the combination of these two restrictions.

The results of these robustness checks are presented in Appendix Figure 2.6. Replacing temporary sale prices with the last observed regular/non-sale observation reduces the weighted average chain-week component by about 15 percentage points to 47%. Nearly all of this decline is a result of an increase in the idiosyncratic store component which is now 41.5%. Sampling monthly observations rather than weekly observations has negligible effects on the results for the analogous sale/non-sale price decompositions. In the

<sup>7</sup>Limitations of the IRI sales indicator are provided in Eden et al. (2021).

<sup>8</sup>This definition is similar to others in the literature (Coibion et al., 2015; Nakamura and Steinsson, 2008).

<sup>9</sup>I follow Gagnon and López-Salido (2020) in selecting the weekly price observation that spans the 15th of the month to reflect BLS price sampling methods.

weakest specification (monthly/non-sale prices), the chain-week component still accounts for at least half of relative price variation for 380 of the 655 goods. Loosening this restriction slightly to 40% of price variation returns the number of goods to a similar level as in the baseline specification, 624 out of 655.

## 2.4 Menu Cost Model with Retail Chains

This section analyzes a menu cost model extended to account for price synchronization within retail chains. I begin by documenting store-level pricing decisions in a standard menu cost model without retail chains as in Nakamura and Steinsson (2008).

### 2.4.1 Standard Model

Consider a firm ( $z$ ) with real profits given by:

$$\Pi_t(z) = \frac{p_t(z)}{P_t} c_t(z) - \frac{W_t}{P_t} L_t(z) - K \frac{W_t}{P_t} I_t(z) \quad (2.3)$$

where  $P_t$  represents the aggregate price level. The first term  $\frac{p_t(z)}{P_t} c_t(z)$  is the firm's revenue where  $\frac{p_t(z)}{P_t}$  is the firm's relative price and  $c_t(z)$  the firm's demand. The firm's total cost of producing in period  $t$  is the real wage  $\frac{W_t}{P_t}$  multiplied by the quantity of labor demanded  $L_t(z)$ . The last term is the firm's menu cost as firm's must hire an additional  $K$  units of labor to change its price.  $I_t(z)$  is an indicator variable that is equal to one if the retailer changes its price in period  $t$  and zero otherwise. Thus, the firm only pays the menu cost if they change their price.

Assume that the demand for the firm's good,  $c_t(z)$ , is proportional to its relative price:

$$c_t(z) = C \left( \frac{p_t(z)}{P_t} \right)^{-\theta} \quad (2.4)$$

where  $C$  is a constant which determines the size of the market. The firm produces its good using a linear technology:

$$y_t(z) = A_t(z) L_t(z) \quad (2.5)$$

where  $y_t(z)$  denotes the output of the firm in period  $t$  and  $A_t(z)$  denotes the productivity of the firm. Markets clear in equilibrium, so that  $y_t(z) = c_t(z)$ . Using equations (2.4) and (2.5), we have that  $L_t(z) = c_t(z)/A_t(z)$ .

Following Nakamura and Steinsson (2008), I assume that the real wage is constant and equal to  $\frac{W_t}{P_t} = \frac{\theta-1}{\theta}$ . Substituting the real wage, firm demand, and market clearing conditions into (2.3) yields

$$\Pi_t(z) = C \left( \frac{p_t(z)}{P_t} \right)^{-\theta} \left( \frac{p_t(z)}{P_t} - \frac{\theta-1}{\theta} \frac{1}{A_t(z)} \right) - K \frac{\theta-1}{\theta} I_t(z) \quad (2.6)$$

The firm then chooses its price at time  $t$  to maximize discounted profits:

$$V(p_{t-1}(z)/P_t, A_t(z)) = \max_{p_t(z)} [\Pi_t(z) + \beta E_t V(p_t(z)/P_{t+1}, A_{t+1}(z))] \quad (2.7)$$

where  $V(\cdot)$  is the firm's value function and  $\beta$  is the discount factor. The firm's state variables are its relative price  $p_{t-1}/P_t$  and productivity level  $A_t(z)$  as evident from (2.6).

Uncertainty arises from aggregate shocks to the price level and idiosyncratic productivity shocks. The process for the price level follows:

$$\log P_t = \mu + \log P_{t-1} + \eta_t \quad (2.8)$$

where  $\eta_t \sim N(0, \sigma_\eta^2)$ . Productivity follows an AR(1) process:

$$\log A_t(z) = \rho \log A_{t-1}(z) + \varepsilon_t(z) \quad (2.9)$$

where  $\varepsilon_t(z) \sim N(0, \sigma_\varepsilon^2)$ .

#### 2.4.2 Retail Chain Extension

I extend the menu cost model to account for retail chain pricing by including a common retail component to the store's productivity process. Thus, in the extended model, a store  $z$  which belongs to retail chain  $r$  follows the productivity process:

$$\log A_t(z) = \rho \log A_{t-1}(z) + \varepsilon_t(z) \quad (2.10)$$

$$= \rho \log A_{t-1}(z) + \varepsilon_t(r) + \varepsilon_t(z) \quad (2.11)$$

where  $\varepsilon_t(r) \sim N(0, \sigma_{\varepsilon_r}^2)$  and  $\varepsilon_t(z) \sim N(0, \sigma_{\varepsilon_z}^2)$  are independent. Similarly, retail chain  $r$  has a productivity process that follows  $\log A_t(r) = \rho \log A_{t-1}(r) + \varepsilon_t(r)$ .

I also assume that the retailer sets the price of store  $z \in r$  in period  $t$  with probability  $\lambda$ , and store  $z$  sets its optimal price with probability  $1 - \lambda$ . When determining store  $z$ 's price, the retailer has the added restrictions that they (1) can set only one price for all stores belonging to  $r$  and (2) observe only chain-level state variables. The profit function of the retail chain can then be written as:

$$\Pi_t(r) = \sum_{z \in r} \left( C \left( \frac{p_t(r)}{P_t} \right)^{-\theta} \left( \frac{p_t(r)}{P_t} - \frac{\theta - 1}{\theta} \frac{1}{\hat{E}_t A_t(z)} \right) - K \frac{W_t}{P_t} I_t(r) \right) \quad (2.12)$$

Table 2.3: Benchmark Parameters

<b>Internally Calibrated (Means)</b>	
Probability Retailer sets Price	$\lambda = 0.499$
Retailer Productivity Shock Std. Dev.	$\sigma_{\varepsilon_r} = 0.035$
Store Productivity Shock Std. Dev.	$\sigma_{\varepsilon_z} = 0.072$
Menu Cost	$K/C = 0.018$
<b>Remaining Parameters</b>	
Discount Factor	$\beta = 0.96^{1/12}$
Elasticity of Demand	$\theta = 4$
Persistence of Productivity	$\rho = 0.7$
Mean Price Level Growth	$\mu = 0.0022$
Standard Deviation of Price Level Growth	$\sigma_\eta = 0.0028$

Note: This table presents the parameters used in the benchmark model. The menu cost, retailer productivity shock volatility, store-level productivity shock volatility, and probability that the retailer sets the store price are internally calibrated to match the mean fraction of adjusted prices, mean absolute size of a (non-zero) price change, the fraction of small price changes, and the chain-week component of the variance decomposition in Section 2.3. The remaining parameters are set similar to either Nakamura and Steinsson (2008) or Nakamura and Steinsson (2010).

where  $\hat{E}_t$  denotes the chain's expectation operator. Using Equation (2.11), we have  $\hat{E}_t A_t(z) = A_t(r)$ .<sup>10</sup> These assumptions simplify the retailer's problem and allow the retailer to behave similarly to an individual store with the following profit and value functions:

$$\Pi_t(r) = C \left( \frac{p_t(r)}{P_t} \right)^{-\theta} \left( \frac{p_t(r)}{P_t} - \frac{\theta-1}{\theta} \frac{1}{A_t(r)} \right) - K \frac{\theta-1}{\theta} I_t(r) \quad (2.13)$$

$$V(p_{t-1}(r)/P_t, A_t(r)) = \max_{p_t(r)} [\Pi_t(r) + \beta E_t V(p_t(r)/P_{t+1}, A_{t+1}(r))] \quad (2.14)$$

### 2.4.3 Calibration

I calibrate the model to match four empirical moments separately for each of the 655 goods in the sample. These moments are the (1) mean fraction of adjusted prices, (2) mean absolute size of a (non-zero) price change, (3) the fraction of small price changes, and (4) the chain-week component of the variance decomposition in Section 2.3.<sup>11,12</sup> I match these moments using the menu cost ( $K/C$ ), the volatility of the retailer's productivity shock ( $\sigma_{\varepsilon_r}$ ), the volatility of the store's total productivity shock ( $\sigma_{\varepsilon_z}$ ), and the probability that the retailer sets an individual store's price ( $\lambda$ ).

<sup>10</sup>Recall that  $A_t(z) = \rho A_{t-1}(z) + \varepsilon_t(z)$ . Consider the  $MA(\infty)$  representation  $A_t(z) = \sum_{i=0}^{\infty} \rho^i \varepsilon_{t-i}(z) = \sum_{i=0}^{\infty} \rho^i [\varepsilon_{t-i}(r) + \varepsilon_{t-i}(z)]$ . Separating the retail and store terms and taking expectations yields:  $\hat{E}_t A_t(z) = \hat{E}_t [\sum_{i=0}^{\infty} \rho^i \varepsilon_{t-i}(r) + \sum_{i=0}^{\infty} \rho^i \varepsilon_{t-i}(z)] = A_t(r) + \sum_{i=0}^{\infty} \rho^i \hat{E}_t [\varepsilon_{t-i}(z)] = A_t(r)$ .

<sup>11</sup>I use the monthly regular/non-sale price empirical statistics when calibrating the model.

<sup>12</sup>Using aggregate data, it is common to use the median fraction of adjusted prices and median size of price changes due to aggregation issues (Baharad and Eden, 2004; Bills and Klenow, 2004). The issues regarding the mean and median are less pronounced when using sector-level data. I follow Nakamura and Steinsson (2010) and Carvalho and Kryvtsov (2021) and use the mean for each good.

Table 2.3 presents the set of calibrated parameters. Means are presented for the internally calibrated parameters which vary across goods. The mean estimated parameters are  $K/C = 0.018$ ,  $\sigma_{\varepsilon_r} = 0.035$ ,  $\sigma_{\varepsilon_z} = 0.072$ , and  $\lambda = 0.499$ . The remaining parameters ( $\beta, \theta, \mu, \sigma_\eta, \rho$ ) are set similar to Nakamura and Steinsson (2008) or Nakamura and Steinsson (2010) and do not vary across goods. I set the discount factor to  $\beta = 0.96^{1/12}$ , the elasticity of demand to  $\theta = 4$ , the persistence of both retailer and store-level productivity to  $\rho = 0.7$ . The inflation process follows  $\mu = 0.0022$  and  $\sigma_\eta = 0.0028$ .<sup>13</sup>

#### 2.4.4 Model Intuition

Although the retail pricing assumptions initially appear strong, they lend themselves to an intuitive interpretation of pricing.<sup>14</sup> Pricing in this environment can be seen as (1) the retailer sets a target price for all stores belonging to its chain using chain-level state variables. This target price can reflect shocks to all stores in the chain such as warehousing/production costs. (2) Stores choose to keep the retail chain price or set their own price. Allowing stores to deviate from the retail-level price with probability  $1 - \lambda$  serves as a reduced form modelling approach for the extra cost that a retailer needs to pay to observe idiosyncratic demand/supply shocks or the cost a store pays for deviating from the chain price.<sup>15</sup>

Figure 2.3 provides an example pricing path for a given retailer-store combination. Each panel corresponds to the same model sample period. The black line plots the aggregate price level in all three panels. Panel (a) plots the retailer’s target price in red and its inverse productivity in blue. Corresponding with Step (1), we see that the retailer adjusts its price to account for both its productivity level and the aggregate price level.

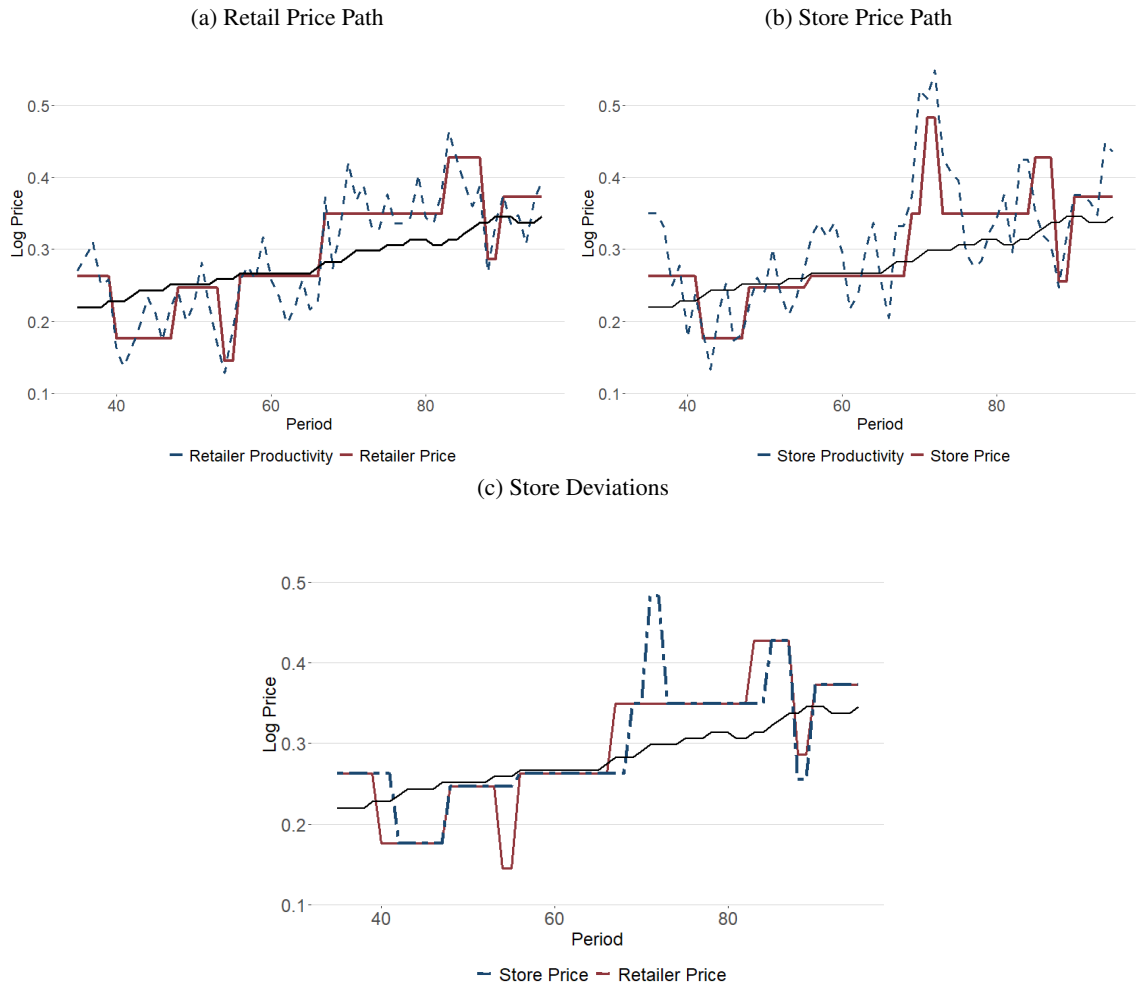
Panel (b) plots the price path for an individual store belonging to the retailer in panel (a). Panel (c) helps highlight the main intuition of the model. The red line in panel (c) plots the retailer’s price. The blue line highlights store-level deviations from the chain’s price. The store sets its price at the chain-level price for most of the sample. Corresponding with Step (2), we see that the store chooses a different price when its idiosyncratic productivity deviates enough from the chain’s productivity. This is particularly evident in period seventy where the store’s inverse productivity is much larger than the chain’s productivity. Overall, the model does well in replicating price paths similar to the data where store deviations from the retail price occur infrequently. Furthermore, when deviations do occur they tend to coincide with periods in which the

<sup>13</sup>The inflation parameters are not taken directly from Nakamura and Steinsson (2008), but rather I follow their procedure. I calibrate  $\mu$  and  $\sigma_\eta$  using CPI data from 2001-2007 to correspond with the sample period in this paper.

<sup>14</sup>The retail chain assumptions of (1) correlated productivity shocks and (2) the retailer constraint are built upon previous studies (in addition to the results in Section 2.3 in this paper). For example, Nakamura and Steinsson (2013) state that productivity shocks may “stand in for other, unmodeled sources of variation in firms’ desired prices” in menu cost models. Hitsch et al. (2019) show that stores belonging to the same retail chain have similar demand in the same market. Furthermore, costs such as wholesale prices are likely similar for stores belonging to the same retailer which play a significant role in price determination (Anderson et al., 2017).

<sup>15</sup>I tested an alternative version of the model where stores pay a cost for deviating from the chain price. The model yielded similar pricing decisions.

Figure 2.3: Example Pricing Decision



Note: This figure presents a standard price path for both a retailer and a store. The black line plots the aggregate price level in both panels. Panel (a) plots the retailer's target price in red and its inverse productivity in blue. Panel (b) is analogous to panel (a) with observations at the store level. Panel (c) plots the retailer price in red and store-level deviations from the retailer price in blue.

retail chain changes its price. The store then resets to the chain price within several periods similar to the price paths seen in Figure 2.1.

## 2.5 Selection Effects in Retail Chain Pricing

The previous sections have analyzed how retail chains are an important determinant in store-level pricing decisions. This section aims to relate the effect of uniform pricing within retail chain chains to the macroeconomy. Specifically, this section analyzes the selection effect—the concept that firms time their price changes optimally in response to aggregate shocks rather than change their prices at random. The first subsection briefly describes the intuition behind selection effects and the estimation procedure. The second subsection discusses the results.

### 2.5.1 Quantifying Selection Effects

To analyze how retail chains affect price selection, I simulate the calibrated retail chain model in Section 2.4. For each good, I compute the total number of chains in the IRI dataset as well as the average number of stores per chain rounded to the nearest whole number. The total number of stores in the simulation is then given by the total number of chains multiplied by the average number of stores per chain.<sup>16</sup> The same inflation process is drawn in each simulation for all 655 goods. The model is simulated for 300 periods with a burn-in sample of 100 periods for 400 periods in total.

After simulating the retail chain model, I perform a counterfactual analysis where stores do not adhere to the retail-chain constraint. Each store faces the same inflation process and analogous productivity process in the retail-chain simulation. Thus, differences between the models are not driven by different productivity draws and solely by every store selecting its unconstrained optimal price in each period.

#### 2.5.1.1 Specification

After simulating each model, I perform the following regression for each good:

$$\Delta \log p_{jt} = \alpha + \beta_P \Delta \log P_t + \beta_A \Delta \log A_{jt} + \varepsilon_{jt} \quad (2.15)$$

where  $\Delta \log p_{jt}$  is the change of the store  $j$ 's log nominal price in period  $t$ ,  $\Delta \log P_t$  is the change of the log price level, and  $\Delta \log A_{jt}$  is the change in the store's nominal cost.

In this circumstance,  $\beta_P$  represents my measure of the selection effect.<sup>17</sup> As the specification is in logs,  $\beta_P$  represents the effect of a 1% increase in the aggregate price level on a store's price on average. For example, if  $\beta_P = 0.2$ , then a 1% increase in the aggregate price level leads to a 0.2% increase in a store's price on average.

#### 2.5.1.2 Intuition

The specification in equation (2.15) lends itself to an intuitive interpretation of the selection effect. To see this, consider reducing the variance of the store's idiosyncratic productivity process in the standard menu cost model. Nakamura and Steinsson (2010) show reducing this variance leads nominal shocks to have less real effects on the economy. This is a result of the average inflation rate becoming a more important determinant in stores' pricing decisions. This is similar to the result in Golosov and Lucas (2007) who show in the absence of idiosyncratic shocks that their model converges to Caplin and Spulber (1987) in which nominal shocks


---

<sup>16</sup>Given this calibration, the total number of stores may differ slightly in the model compared to the data. Overall, these differences tend to be small and are unlikely to affect the results.

<sup>17</sup>Note that this is not a structural definition of the selection effect as in Caballero and Engel (2007); Carvalho and Schwartzman (2015); Dotsey and Wolman (2018); Karadi and Reiff (2019). Instead, this specification serves as a reduced form manner to estimate the responsiveness of micro prices to aggregate shocks.



Figure 2.4: Selection Effects with Retail Chain Pricing



Note: This figure plots the distributions of the average percent response of a store's price to a 1% increase in the aggregate price level ( $\beta_P$ ) from equation (2.15) for the unconstrained model in red and the retail-chain model in blue. A KS-test suggests that the distribution are significantly different with a maximum difference of 0.75.

have no real effects on the economy. In this setup, the coefficients  $\beta_P$  and  $\beta_A$  can be loosely interpreted as the weight that a store places on the aggregate price level and its idiosyncratic productivity, respectively, when choosing to change its price. Thus, as  $\beta_P$  increases, stores place more weight on the aggregate price level when timing and deciding the size of their price changes. As a result, the real effects of nominal shocks would decrease.

### 2.5.2 Results

Figure 2.4 presents the results of equation (2.15). The bar graphs in red and blue plot the distribution of  $\beta_P$  over goods for the unconstrained and constrained retail-chain models, respectively. A KS-test suggests that the distributions are significantly different with a maximum distance of 0.75. The weighted mean  $\beta_P$  for the constrained and unconstrained models are 0.25 and 0.09, respectively. This difference suggests that the standard menu cost model without retail chains significantly underestimates the degree of selection. Consequently, this suggests the standard menu cost model overestimates the degree of monetary non-neutrality by not accounting for price synchronization within retail chains.

Table 2.4: Selection Effects in Retail Chain Pricing

	Baseline Model			CalvoPlus		
	Retail Chain (1)	Standard (2)	(1)/(2) (3)	Retail Chain (4)	Standard (5)	(4)/(5) (6)
Price Level ( $\beta_P$ )	0.25% [0.09, 0.41]	0.09% [0.008, 0.18]	2.72	0.24% [0.09, 0.38]	0.1% [0.03, 0.18]	2.26
Productivity ( $\beta_A$ )	0.44% [0.21, 0.67]	0.61% [0.32, 0.88]	0.74	0.40% [0.18, 0.61]	0.55% [0.32, 0.77]	0.72

Note: This table presents the results of estimating Equation (2.15). The first row presents the weighted mean (over goods) elasticity of store prices with respect to the aggregate price level ( $\beta_P$ ). The second row presents the weighted mean elasticity with respect to the store's idiosyncratic productivity ( $\beta_A$ ). Columns (1) and (2) present the estimated elasticities in the baseline model with and without retail chains, respectively. Column (3) presents the estimated retail chain elasticity relative to the standard model. Columns (4)-(6) present similar results with the CalvoPlus extension of random free price changes.

### 2.5.3 CalvoPlus Model Extension

I show that the baseline results hold in a richer setting using the CalvoPlus model of Nakamura and Steinsson (2010). First-generation menu cost models such as the baseline model in this paper fail to match several empirical facts about price changes (Klenow and Kryvtsov, 2008; Midrigan, 2011). In light of this, it has become common to incorporate a random menu cost component as in the CalvoPlus model (Alvarez et al., 2016; Hee Hong et al., 2021; Carvalho and Kryvtsov, 2021).<sup>18</sup>

The CalvoPlus model incorporates free price changes with probability  $\alpha$ . When calibrating the CalvoPlus model, I set the probability of a free price change equal to the empirical frequency of small price changes for each good. Calibration then proceeds as in Section 2.4. These parameters are presented in Appendix Table 2.5.

Table 2.4 present the results of estimating Equation (2.15). The first row presents the weighted mean (over goods) elasticity of store prices with respect to the aggregate price level ( $\beta_P$ ). The second row presents the weighted mean elasticity with respect to the store's idiosyncratic productivity ( $\beta_A$ ). The results of the CalvoPlus model with and without retail chains are presented in columns (4) and (5), respectively. For comparison, the results of the baseline model are presented in columns (1) and (2). Columns (3) and (6) presents the estimate for the retail chain model divided by estimate for the model without retail chains. Confidence intervals at the 95% level are presented in brackets.<sup>19</sup>

Comparing columns (1) and (4), we see that the baseline model and CalvoPlus model yield similar results

<sup>18</sup>Other methods have been used to help match the empirical price statistics such as altering the productivity shock process (Gertler and Leahy, 2008) and introducing economies of scope (Midrigan, 2011).

<sup>19</sup>Confidence intervals are constructed by calculating the weighted standard deviation over goods for each estimate.

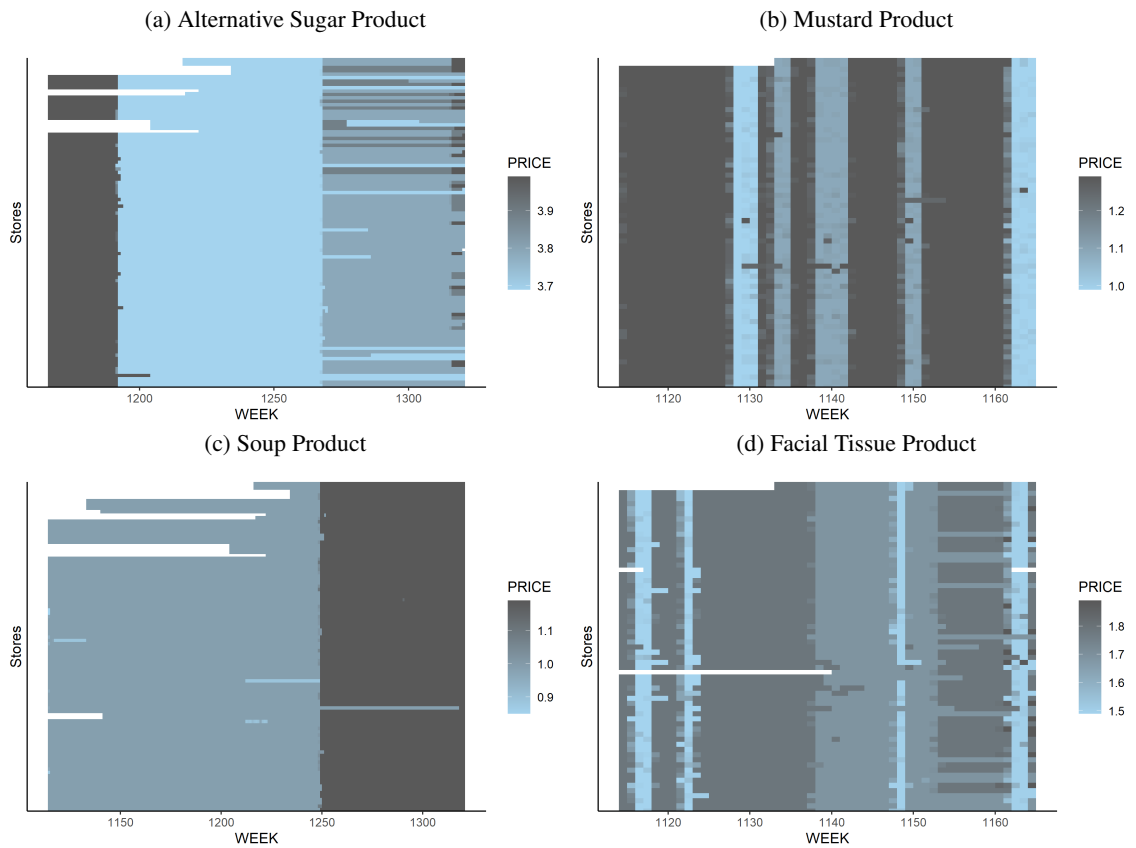
when accounting for retail chains. The weighted mean aggregate price elasticities are 0.25% and 0.24%, respectively. The confidence intervals suggest that these estimates are statistically significant over goods. Overall, the CalvoPlus model typically estimates lower elasticities compared to the baseline model. However, these differences tend to be small, and they have little effect on the *relative* elasticities. The elasticity estimated with retail chains relative to that estimated without chains is  $0.24\%/0.1\% = 2.26$ . Thus, both the baseline model and the CalvoPlus model suggest that the inclusion of retail chains more than doubles the degree of selection effects.

## 2.6 Conclusion

This paper examines selection effects in retail chain pricing. Using scanner-level data, I find that retail chains synchronize the timing and magnitude of their price changes across stores. A variance decomposition suggests that retail chains account for almost two-thirds of stores' relative price dispersion on average. This relationship is prevalent across almost all goods in my sample with the retail chain component of the decomposition accounting for at least half of price variation for 625 of the 655 goods. I develop a menu cost model with retail chain price synchronization to account for this finding. After calibrating the model separately for all 655 goods in my sample, I then measure selection by regressing the change in a store's log price on the change in the aggregate price level. My estimates suggest that a 1% increase in the aggregate price level leads to a 0.25% increase in a store's price on average. This effect is more than double suggested by an analogous simulation of the standard menu cost model without retail chains. This relationship suggests that the standard menu cost model overestimates the degree of monetary non-neutrality by ignoring synchronization in retail chain pricing.

2.7 Appendix  
2.7.1 Empirics

Figure 2.5: Retail Chain Price Synchronization



Note: This figure plots examples of uniform pricing within a retail chain. Each point on the y-axis represents an individual store belonging to the retail chain. Darker (lighter) shades represent higher (lower) prices. Missing values are represented by white space. Panel (a) presents an example for alternative sugar product compared to Figure 2.1. Panels (b), (c), and (d) present an example for a mustard, soup, and facial tissue product, respectively.

### 2.7.1.1 Variance Decomposition

The variance decomposition in Section 2.3 began by modelling prices as

$$p_{j,r,t} = \alpha_t + \delta_j + \gamma_{r,t} + \varepsilon_{j,r,t} \quad (2.16)$$

where  $\alpha_t$  was a week fixed effect,  $\delta_j$  was a store fixed effect,  $\gamma_{r,t}$  was a chain-by-week fixed effect, and  $\varepsilon_{j,r,t}$  was the residual. Due to the size of the dataset, I used the method of iterative means to estimate the fixed effects following Darulich and Kozlowski (2021) and Kaplan et al. (2016). The order of the estimation is (1)  $\delta_j$ , (2)  $\alpha_t$ , (3)  $\gamma_{r,t}$ , and (4)  $\varepsilon_{i,j,t}$ . Thus, each fixed effect is estimated as (good index  $i$  is omitted as estimation is conducted separately for each good):

$$\hat{\delta}_j = \frac{1}{T_j} \sum_t p_{j,r,t} \quad \text{Store Component} \quad (2.17)$$

$$\hat{\alpha}_t = \frac{1}{N_{j,t}} \sum_j (p_{j,r,t} - \hat{\delta}_j) \quad \text{Week Component} \quad (2.18)$$

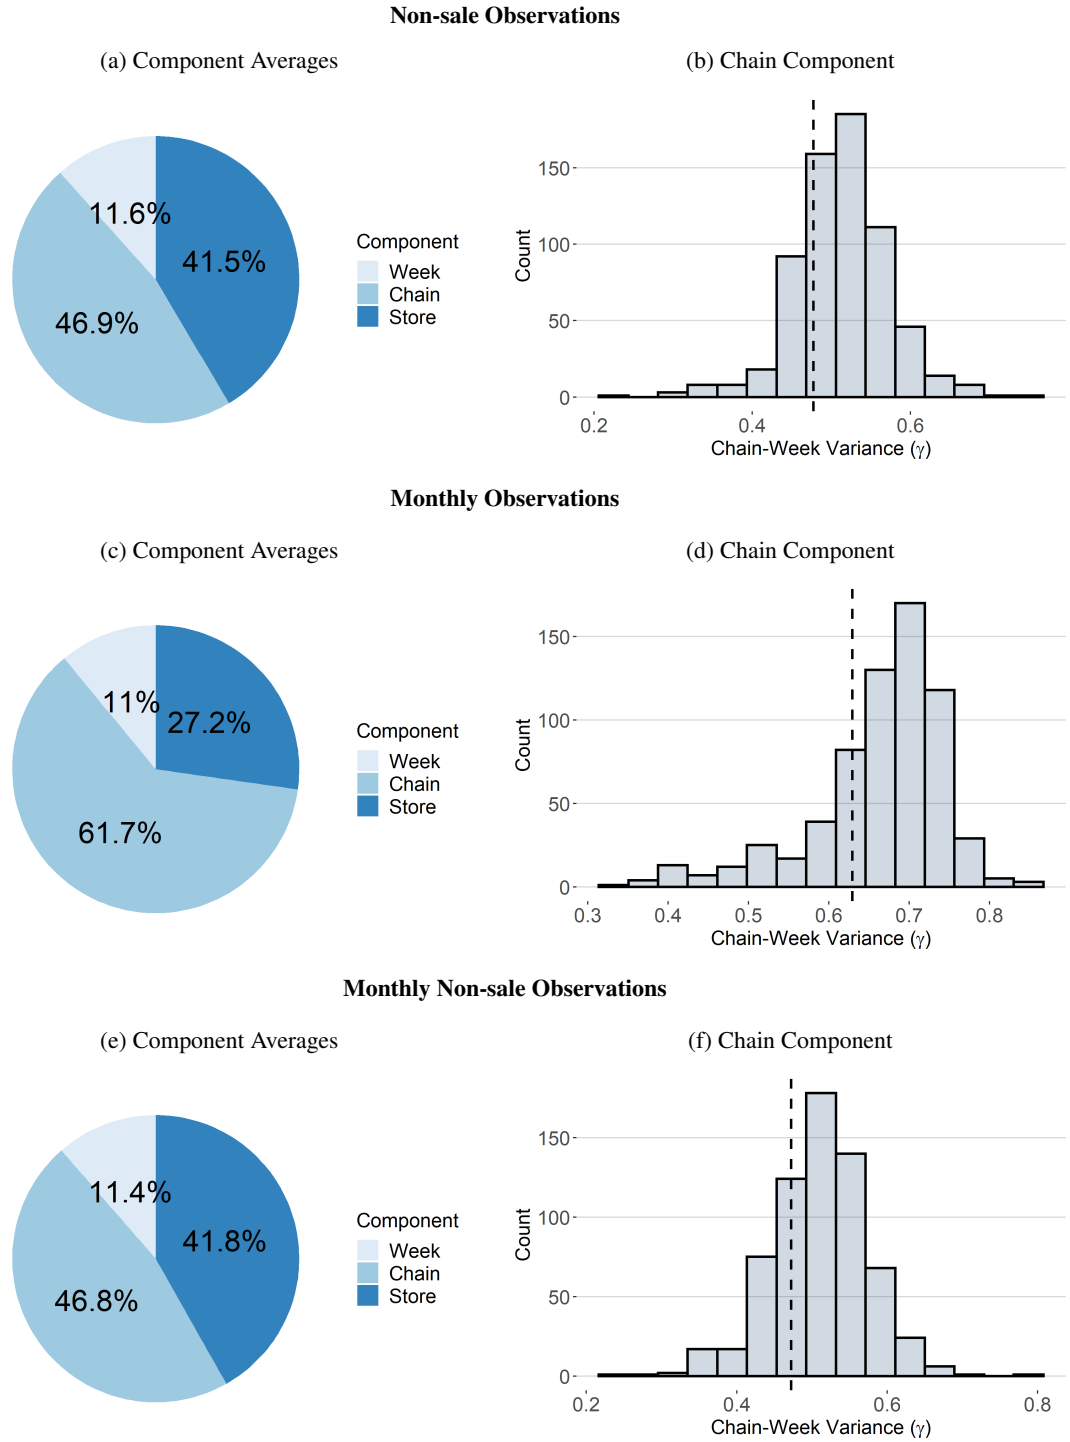
$$\hat{\gamma}_{r,t} = \frac{1}{N_{j \in r,t}} \sum_{j \in r,t} (p_{j,r,t} - \hat{\delta}_j - \hat{\alpha}_t) \quad \text{Chain-Week Component} \quad (2.19)$$

$$\hat{\varepsilon}_{j,r,t} = p_{j,r,t} - \hat{\delta}_j - \hat{\alpha}_t - \hat{\gamma}_{r,t} \quad \text{Residual Component} \quad (2.20)$$

where  $T_j$  is the number of weeks that store  $j$  has a posted price,  $N_{j,t}$  is the total number stores in week  $t$ , and  $N_{j \in r,t}$  is the total number of store observations that belong to retail chain  $r$  in week  $t$ .

After estimating the fixed effects, the following variance decomposition of relative prices was performed  $\text{Var}(p_{j,r,t} - \hat{\delta}_j) = \text{Var}(\hat{\alpha}_t) + \text{Var}(\hat{\gamma}_{r,t}) + \text{Var}(\hat{\varepsilon}_{j,r,t})$ . Several assumptions lead to the absence of covariance terms in this equation. First,  $E[\alpha_t] = 0$ . Thus, deviations from a store's average price are zero in expectation. Second, in a given week, deviations from a store's relative price are zero in expectation after accounting for the week component,  $E[\gamma_{r,t}|t] = 0$ .

Figure 2.6: Variance Decomposition (Robustness Checks)



Note: This figure presents the results of the variance decomposition in Equation (2.2). The variance decomposition is conducted separately for each good. Panel (a) presents the mean estimate for each component of the decomposition. Panel (b) plots the distribution of the chain-week variance ( $\gamma$ ) over all goods. The vertical dashed line represents the cutoff for the first quartile.

## 2.7.2 Model

Table 2.5: CalvoPlus Parameters

<b>CalvoPlus Parameter (Mean)</b>	
Probability of Free Price Change	$\alpha = 4.8\%$
<b>Internally Calibrated (Means)</b>	
Probability Retailer sets Price	$\lambda = 0.498$
Retailer Productivity Shock Std. Dev.	$\sigma_{\varepsilon_r} = 0.034$
Store Productivity Shock Std. Dev.	$\sigma_{\varepsilon_z} = 0.070$
Menu Cost	$K/C = 0.019$
<b>Remaining Parameters</b>	
Discount Factor	$\beta = 0.96^{1/12}$
Elasticity of Demand	$\theta = 4$
Persistence of Productivity	$\rho = 0.7$
Mean Price Level Growth	$\mu = 0.0022$
Standard Deviation of Price Level Growth	$\sigma_\eta = 0.0028$

Note: This table presents the parameters used in the CalvoPlus model. The CalvoPlus parameter is set to the frequency of small price changes. The menu cost, retailer productivity shock volatility, store-level productivity shock volatility, and probability that the retailer sets the store price are internally calibrated to match the mean fraction of adjusted prices, mean absolute size of a (non-zero) price change, the fraction of small price changes, and the chain-week component of the variance decomposition in Section 2.3. The remaining parameters are set similar to either Nakamura and Steinsson (2008) or Nakamura and Steinsson (2010).

## CHAPTER 3

### Temporary Sales in Response to Demand Shocks

#### 3.1 Introduction

Temporary sales – defined as large price drops which quickly rebound – are commonly viewed as a tool for price discrimination that cancels out at the aggregate and contains no information on macroeconomic conditions (Chevalier and Kashyap, 2019; Guimaraes and Sheedy, 2011; Salop and Stiglitz, 1977; Shilony, 1977; Varian, 1980). In light of this, it has become standard practice in macroeconomics to ignore sales and focus on the behavior of regular/non-sale prices (Nakamura and Steinsson, 2008; Eichenbaum et al., 2011; Kehoe and Midrigan, 2015). Here, we examine this common view and argue that temporary sales play an important role in the reaction of stores to demand shocks.

We start by asking whether it is indeed the case that sales cancel out in the aggregate. We find that sales do not cancel out. On average, a product is not on sale in any store in over 40% of weeks. We also find that the frequency of sales increases following negative demand shocks. We then illustrate the importance of temporary sales by comparing the average price reaction to a demand shock with the reaction of the average non-sale price. The average price decreases by almost 1.4% following a negative demand shock. Ignoring sales and analyzing the average non-sale price mitigates this response to 0.7%. Thus, sales account for about half of the average price decline following demand shocks.

We interpret these findings in the context of a model in which temporary sales are reactions to unwanted inventories. In the model, inventories are accumulated after negative demand shocks. The main insight of the model is that, if storage is associated with depreciation (as is the case with perishable goods that have expiration dates), then sharp, temporary reductions in prices can occur even in response to moderate shocks.

Our model is a flexible price version of Prescott (1975) hotels model: The Uncertain and Sequential Trade (UST) model in Eden (1990).<sup>1</sup> Most closely related is Bental and Eden (1993) that allows for storage and assumes exponential decay. In their model, there are demand and supply shocks, and the equilibrium price distribution depends on the current cost shock and the beginning of period level inventories. Inventories are accumulated when demand in the previous period was low. The accumulation of inventories leads to a reduction in prices (the entire price distribution shifts to the left) and as a result the quantity sold increases on average. Roughly speaking, the reduction in prices lasts until inventories are back to their “normal” level.

We adopt here the feature emphasized by Eden (2018), who assumes that units close to their expiration date are offered at a low price to minimize the probability that they will reach the expiration date before

---

<sup>1</sup>For rigid price versions of the model, see Dana (1998, 1999) and Deneckere and Peck (2012).



being sold. A store may therefore start at a relatively high “regular price” and then if it fails to make a sale switch to a low price until the level of inventories get back to “normal”. This mechanism is motivated by Aguirregabiria (1999) who found a significant and robust effect of inventories at the beginning of the month on the current price using a unique data set from a chain of supermarket stores in Spain.

In practice, temporary sale decisions are often negotiated between the chain’s headquarters and its suppliers. Aguirregabiria (1999) describes that the toughest part of the negotiation with suppliers is about the number of weeks during the year that the brand will be under promotion, and about the percentage of the cost of sales promotions that will be paid by the wholesaler (e.g. cost of posters, mailing, price labels). A similar description is in Anderson et al. (2017) who present institutional evidence that sales are complex contingent contracts that are determined substantially in advance. However, there is also some flexibility. For many promotions, manufacturers allow for a “trade deal window” of several weeks where the seller can execute the promotion. The flexibility in the timing of sales—as evidenced in our empirical framework—may also reflect the need to respond to inventories that were accumulated as a result of demand shocks.

Chevalier and Kashyap (2019) raise the issue of close substitutes when aggregating prices. We focus on the average posted price at the product level rather than an aggregate price index (Chevalier and Kashyap, 2019) or the average price paid by the consumer (Coibion et al., 2015; Gandon, 2018). As our focus is thus on the observed behavior of the store, this allows us to avoid the concerns of close substitutes in measuring cost-of-living or consumer welfare.

Kryvtsov and Vincent (2021) find a robust relationship between the frequency of sales and the rate of unemployment. We use weekly data to study the question of price rigidity and focus on shocks to the demand for a narrowly defined product. We do not attempt to study aggregate shocks, like shocks to the money supply, which affect the demand for all goods. Nevertheless, we think that our study is relevant for the study of aggregate shocks. For example, Bental and Eden (1996) and Eden (1994) develop monetary versions of the UST model. These monetary models illustrate that it does not matter much if the uncertainty is about the number of buyers that will arrive at the marketplace or the number of dollars that will arrive. Given this, it is likely that shocks which affect the demand for all goods and stores will work in a similar manner as a good-specific demand shock that affects all stores.

The rest of the paper proceeds as follows. Section 3.2 describes the data and summary statistics of temporary sales. Section 3.3 presents several stylized facts about store and product-level decision making of temporary sales. Section 3.4 analyzes the impact of temporary sales on the average posted price in response to demand shocks. Section 3.5 presents the Uncertain and Sequential Trade (UST) model in light of the empirical results. Section 3.6 concludes.

## 3.2 Data

Our primary analysis covers grocery stores across the US from 2004-2005 using the Information Resources Inc. (IRI) retail scanner dataset.<sup>2</sup> Observations are recorded weekly at the store-product level. Products are defined by their Universal Product Code (UPC) and belong to specific product categories (e.g. Beer, Hot Dogs). Stores in our sample are located in different markets across the US. A market is sometimes classified as a city (Chicago, Los Angeles) or a state (Mississippi). For each store-product-week observation, we see the total revenue ( $Rev_{ijt}$ ) and total quantity sold ( $Q_{ijt}$ ). We then compute the average price of product  $i$  in store  $j$  for the week  $t$  as the total revenue divided by quantity,  $P_{ijt} = \frac{Rev_{ijt}}{Q_{ijt}}$ .

### 3.2.1 Filtering Process

In order to analyze the effect of temporary sales, we apply the following filter in a sequential manner by market area:

1. We drop all UPC-Store cells that do not have strictly positive quantities in all weeks.
2. We drop all UPCs that are sold by fewer than 11 stores.
3. We drop all categories with less than 10 UPCs.

The first exclusion is applied because we cannot distinguish between product stockouts (product is not carried in the store) and periods of low demand (product is carried by the store but not purchased). These two scenarios carry different supply and demand implications which may obscure our analysis. This step is also required for identifying temporary sale prices. The second exclusion is aimed at reliable measures of the average cross-sectional price distribution. This requirement also leads to a sample of fairly popular brands.<sup>3</sup> The third economizes on the number of category dummies. After applying our filter, we obtain a semi-balanced panel in which the number and identities of stores may vary across UPCs, but does not vary over weeks for a given UPC.

Panel A of Table 3.1 provides summary statistics for our filtering process. The original sample contains over 1,500 unique products across 56,000 stores for a total of almost 400 million store-product-week observations. Observations span 31 categories and 50 markets across the US in the full sample. Our final sample reduces the total number of observations to under 10 million. The restriction that products are continuously sold accounts for over 90% of this reduction. The filtering process reduced the total number of stores to 546 located across 26 markets. The final sample spans 1,686 unique products which belong to 22 product categories.

---

<sup>2</sup>A complete description of the dataset can be found in Bronnenberg et al. (2008).

<sup>3</sup>This is not unique to this paper. Sorenson (2000) has collected data on 152 top selling drugs. Lach (2002) excluded products that were sold by a small number of stores. Kaplan and Menzio (2015) exclude UPCs with less than 25 reported transactions during a quarter in a given market.

Table 3.1: Summary Statistics of IRI Data

<b>Panel A: Filtering Process</b>					
	Stores	Products	Categories	Markets	Observations
Pre-filter	1,589	56,342	31	50	395,756,478
Product Continuously Sold	693	12,114	31	50	35,718,904
Product $\geq$ 11 Stores	546	1,808	25	26	11,440,832
Category $\geq$ 10 UPCs	546	1,686	22	26	9,828,520
<b>Panel B: Final Sample Composition</b>					
Market Name	Percent of Sample	Product Category	Percent of Sample		
New York, NY	16.7%	Carbonated Beverages	20.6%		
Los Angeles, CA	12.0%	Yogurt	15.8%		
New England	8.4%	Salted Snacks	11.5%		
Seattle/Tacoma, WA	6.9%	Cold Cereal	10.6%		
Dallas, TX	5.3%	Milk	10.4%		
St. Louis, MO	5.2%	Beer	6.9%		
Chicago, IL	5.0%	Soup	6.8%		
Philadelphia, PA	4.6%	Margarine/Butter	5.6%		
San Francisco, CA	4.4%	Mayonnaise	1.9%		
Raleigh/Durham, NC	3.6%	Peanut Butter	1.5%		
Portland, OR	3.2%	Mustard/Ketchup	1.4%		
Phoenix, AZ	3.2%	Toilet Tissue	1.4%		
Washington D.C.	2.6%	Frozen Dinner Entrees	1.1%		
Houston, TX	2.5%	Spaghetti Sauce	1.1%		
Richmond/Norfolk, VA	2.4%	Hot Dogs	0.8%		
Roanoke, VA	2.3%	Paper Towels	0.5%		
Kansas City	2.0%	Coffee	0.5%		
San Diego, CA	1.8%	Cigarettes	0.4%		
Knoxville, TN	1.6%	Frozen Pizza	0.4%		
Charlotte, NC	1.6%	Laundry Detergent	0.3%		
Buffalo/Rochester, NY	1.3%	Household Cleaner	0.2%		
Harrisburg/Scranton, PA	1.2%	Facial Tissues	0.2%		
Salt Lake City, UT	0.8%				
Boston, MA	0.6%				
Syracuse, NY	0.5%				
Sacramento, CA	0.2%				

Note: This table presents summary statistics for our data filtering process and the final sample composition. Observations are recorded at the store-product-week level. We require that (1) products are continuously sold from 2004-2005, (2) Products are sold by at least 11 stores in a given market, and (3) Categories contain at least 10 products. Totals may not sum to 100% due to rounding error.

Panel B of Table 3.1 describes the composition of our final sample by market area and product category. We see that New York, NY contains the most observations and comprises 16.7% of our panel. Los Angeles, CA is the only other market to exceed 10% of our sample. Carbonated beverages is our largest product category at 20.6%. Four other categories exceed 10% of observations—Yogurt, Salted Snacks, Cold Cereal, and Milk.

### 3.2.2 Temporary Sales

Table 3.2: Summary Statistics for Sales

<b>Panel A: Sale Frequency by UPC-Store</b>					
	1st Quartile	Median	Mean	3rd Quartile	Maximum
IRI Definition	9.6%	24.0%	28.0%	43.3%	100.0%
10% Temporary Reduction (Baseline)	1.9%	11.5%	14.4%	24.0%	60.6%
Combined Definition (Robustness Checks)	1.0%	10.6%	13.2%	22.1%	58.7%

<b>Panel B: Market and Category Heterogeneity</b>			
Market Name	Sales Frequency	Product Category	Sales Frequency
San Francisco, CA	21.6%	Hot Dogs	22.7%
Chicago, IL	20.4%	Salted Snacks	20.0%
San Diego, CA	20.4%	Yogurt	19.1%
Washington D.C.	19.0%	Carbonated Beverages	18.7%
Harrisburg/Scranton, PA	18.7%	Cold Cereal	14.1%
Sacramento, CA	18.6%	Coffee	13.6%
Boston, MA	17.1%	Margarine/Butter	12.5%
Los Angeles, CA	16.8%	Frozen Dinner Entrees	12.3%
Seattle/Tacoma, WA	16.8%	Frozen Pizza	11.5%
Portland, OR	16.6%	Facial Tissues	10.7%
Phoenix, AZ	16.5%	Peanut Butter	10.5%
Philadelphia, PA	15.2%	Toilet Tissue	10.4%
St. Louis, MO	14.8%	Spaghetti Sauce	10.3%
Syracuse, NY	13.9%	Household Cleaner	9.6%
Salt Lake City, UT	13.8%	Beer	8.6%
Charlotte, NC	13.8%	Soup	8.5%
Buffalo/Rochester, NY	13.4%	Mustard/Ketchup	7.0%
New York, NY	13.2%	Mayonnaise	6.6%
Dallas, TX	12.9%	Milk	6.4%
Roanoke, VA	12.8%	Paper Towels	6.1%
Richmond/Norfolk, VA	12.5%	Laundry Detergent	5.3%
Houston, TX	11.9%	Cigarettes	0.2%
Raleigh/Durham, NC	10.2%		
Knoxville, TN	9.4%		
Kansas City	8.2%		
New England	5.0%		

Note: This table presents summary statistics for temporary sales. Observations are recorded at the store-product-week level. We define a sale as a temporary reduction in price of at least 10% which returns to the pre-sale price or greater within four weeks. Sales frequency is then defined as the fraction of total prices which are temporary sales. Totals in Panel B may not sum to 100% due to rounding error.

The IRI provides an indicator for sale prices based on a proprietary algorithm. The IRI definition of sales is rather obscure and may include promotion activities in addition to the behavior of the price. We use another definition that focuses on the behavior of prices as our baseline definition and combine it with the IRI definition as a robustness check. We assume that a temporary sale occurs when a price drop of at least 10% is followed by a price equal to or above the pre-sale price within four weeks. This definition provides the extra benefit that it can be used for comparability across datasets such as Nielson's Retail Scanner, and it is similar to other definitions of temporary sales (Coibion et al., 2015; Nakamura and Steinsson, 2008). As a

robustness check, we add the requirement that the sales prices also satisfy the IRI definition.

Panel A of Table 3.2 compares the three measures of temporary sales. We define the sales frequency as the fraction of total prices that are a temporary sale. At the store-product level, this is denoted as  $sale\_freq_{ij} = \frac{1}{T} \sum_{t=1}^T I(sale_{ijt} = 1)$ , where  $I(sale_{ijt} = 1) = 1$  if the price (of product  $i$  in store  $j$  at week  $t$ ) is a sale price and zero otherwise. We see that the sales frequency for the IRI definition is greater than our definition of a 10% temporary price reduction across the entire distribution of UPC-Store combinations. For example, the average sales frequency across all UPC-Store combinations is 14.4% for our definition of a 10% temporary price reduction. This is about half of the sales frequency of 28% using the IRI definition of a sale. The IRI definition also contains several observations in which a product is always on sale in a given store. It may be the case that the wider definition of sales adopted by the IRI reflects the need to satisfy supplier imposed requirements (Aguirregabiria, 1999; Anderson et al., 2017). We are primarily interested in the behavior of store-level pricing decisions. For this reason, we adopt our definition based on the behavior of prices for our baseline results. We use the combined definition that a sale requires a 10% temporary price reduction and the IRI flag for robustness checks throughout the paper. This added restriction does not significantly affect the distribution of sales frequency. Overall, the correlation between our baseline and combined definition is 0.99.


Panel B of Table 3.2 compares the overall sales frequency by market area and product category. We see that San Francisco has the highest sales frequency of 21.6%. The New England market has the lowest sales frequency at 5.0%. Sales frequency also varies across product categories. Cigarettes are the least likely products to have a sale with a frequency of 0.2%. Hot dogs have the highest sales frequency of 22.7%.

Figure 3.1 provides a more in depth view of the variation in sales frequency across stores and products. The average sales frequency over stores for a given product is plotted in blue, and the standard deviation in red. Similar to product categories, we see that the sales frequency varies widely across products with a range from 0% to almost 50%. The correlation between the mean and the standard deviation is 0.66. Thus, UPCs with higher frequency of sales tend to have more variation across stores. The large standard deviation measures suggest that store managers have a significant role in choosing temporary sales.

### 3.3 Stylized Facts

Here we advance the hypothesis that stores set sales in response to cost or demand shocks. An alternative hypothesis assumes that stores use a mixed strategy to determine sales (Sheremirov, 2020; Guimaraes and Sheedy, 2011). The first subsection provides evidence against the alternative hypothesis. The second subsection analyzes the relationship between temporary sales and demand uncertainty, and the third subsection analyzes the dynamic relationship between sales and demand shocks.

Figure 3.1: Variation in Sales Frequency



Note: This figure plots the average sale frequency over stores for a given product in blue. The red dots represent the standard deviation of the sales frequency over stores for a given product.

### 3.3.1 Do stores use a mixed strategy?

The mixed strategy hypothesis is that the price distribution is stable across time, but that individual stores change their place within the distribution. This hypothesis predicts that the fraction of weeks in which a product was not on sale in any of the stores which carried it is small. To test this prediction, we define the fraction of weeks with no sales in any of the stores as follows:

$$num\_sale_{it} = \sum_{j=1}^{J_i} I(sale_{ijt} = 1) \quad (3.1)$$

$$no\_sale_i = \frac{1}{T} \sum_{t=1}^T I(num\_sale_{it} = 0) \quad (3.2)$$

where  $J_i$  represents the total number of stores that sell product  $i$ . Thus,  $no\_sale_i$  is the fraction of weeks in which product  $i$  was not on sale in any of the stores.

Panel A of Table 3.3 presents statistics for these variables. On average, a product is not on sale in any of the stores which sold it for 46% of weeks. This equates to about 48 weeks in our sample. If stores use a mixed strategy with the same sale frequency, the implied probability of  $no\_sale_i$  can be written as  $(1 - sale\_freq_i)^{J_i}$ . In our sample, 56 stores carry a product on average. Using the average sale frequency in the data, the implied mixed strategy estimate of  $no\_sale_i$  is then  $(1 - .144)^{56} = 0.016\%$ . This differs significantly from the percentage found in the data. This result may be driven by heterogeneity in the sales frequency over

Table 3.3: Behavior of Temporary Sales

<b>Panel A: Sale Behavior by UPC</b>					
	1st Quartile	Median	Mean	3rd Quartile	Maximum
Sale in No Store ( $no\_sale_i$ )	12.50%	45.19%	46.43%	75.96%	100%
<b>Panel B: Heterogeneity by Frequency of Sale</b>					
	0-5%	5-10%	10-15%	15-20%	>20%
Percent of Sample	32.1%	16.8%	16.3%	14.4%	20.4%
Sale Frequency ( $\overline{sale\_freq}$ )	1.7%	7.6%	12.3%	17.4%	25.8%
Number of Stores ( $\bar{J}$ )	34.9	50.7	56.9	71.7	82.1
Sale in No Store ( $\overline{no\_sale}$ )	83.6%	47.1%	31.5%	22.8%	15.8%
Mixed Strategy ( $(1 - \overline{sale\_freq})^{\bar{J}}$ )	54.7%	1.8%	0.1%	0.0001%	2.2e-9%

Note: This table presents statistics about the fraction of weeks in which a product is not on sale in any of the stores ( $no\_sale_i$ ). Panel A presents summary statistics for this variable. Panel B presents similar statistics grouping products by their sale frequency. Means are denoted by bars.

products. For example, the phenomenon of no sale in any store may occur primarily in products with low sale probabilities.

Panel B of Table 3.3 examines the possibility that our results are driven by this heterogeneity. We divide UPCs into five bins according to their frequency of sale: [0, 5%); [5, 10%); [10, 15%); [15, 20%); [20, 100%]. We present results using the mean of variables in their respective bin which are denoted by bars. We see that a mixed strategy may provide a close approximation for products with a low probability of sale. For products in the [0, 5%) bin, a mixed strategy implies that no stores list a sale price in 54.7% of weeks. The observed frequency in the data is 83.6%. However, these products only comprise 32% of our sample. A mixed strategy provides a poor approximation for all other bins. Products experience no sale prices in all stores in 47% of weeks on average in the [5, 10%) bin. The estimated probability is 1.8% for a mixed strategy. The mixed strategy probability converges towards zero in the higher frequency bins. Although the probability of no sale in the data does decrease with the frequency of sale, this probability remains significantly different than zero. In the > 20% bin, the fraction of weeks with no sales in any of the stores is 15.8%.

### 3.3.2 Temporary sales and demand uncertainty

The main point of the paper is that temporary sales are a reaction to demand shocks and not merely a discrimination device. If temporary sales are a reaction to demand shocks, then stores and products that face more demand uncertainty should have more sales. We start by establishing these correlations between demand

uncertainty and the frequency of temporary sales.<sup>4</sup>

### 3.3.2.1 Empirical Strategy

#### 3.3.2.1.1 Variable Definitions

We use the standard deviation of log units over time as a proxy for demand uncertainty for each product-store combination,  $SDU_{ij}$ . However, sale prices can both respond to and cause changes in demand uncertainty which may bias our estimates. To help mitigate the issue of reverse causality, we calculate  $SDU$  using only observations in which a non-sale price is listed. Representing the quantity sold of product  $i$  at store  $j$  in week  $t$  as  $Q_{ijt}$ , this can be calculated as follows:

$$q_{ijt} = \log(Q_{ijt}) \quad (3.3)$$

$$q_{ij} = \frac{1}{T_{ij}^{reg}} \sum_{t=1}^T q_{ijt} I(\text{sale}_{ijt} = 0) \quad (3.4)$$

$$SDU_{ij} = \sqrt{\frac{\sum_{t=1}^T ((q_{ijt} - q_{ij}) I(\text{sale}_{ijt} = 0))^2}{T_{ij}^{reg} - 1}} \quad (3.5)$$

where  $T_{ij}^{reg}$  represents the number of periods in which product  $i$  in store  $j$  is sold at a regular (non-sale) price.  $I(\text{sale}_{ijt} = 0)$  is an indicator function equal to one if the price of product  $i$  in store  $j$  at week  $t$  is a regular (non-sale) price and zero otherwise.

We follow a similar procedure for prices to define the average log price,  $p_{ij}$ , and the standard deviation of log price,  $SDP_{ij}$ . We then define store-level observations as the average over UPCs:

$$\text{sale\_freq}_j = \frac{1}{I_j} \sum_{i=1}^{I_j} \text{sale\_freq}_{ij} \quad (3.6)$$

$$SDU_j = \frac{1}{I_j} \sum_{i=1}^{I_j} SDU_{ij} \quad (3.7)$$

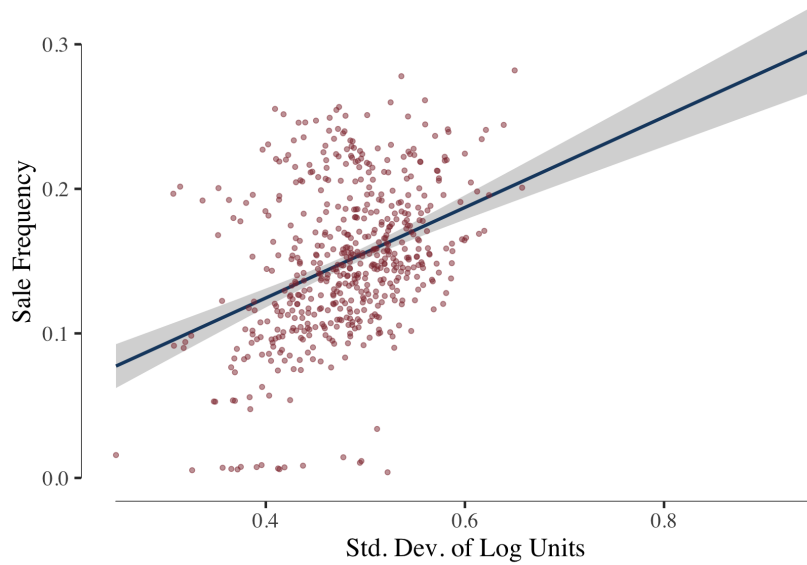
where  $I_j$  is the number of products sold by store  $j$  in our sample. Using the subsample of regular prices, we calculate the average standard deviation of log prices ( $SDP_j$ ), average log units ( $q_j$ ), and the average log price ( $p_j$ ) in a similar manner.

---

<sup>4</sup>Using only the Chicago market, Eden (2018) provides evidence about the correlation between demand uncertainty and the frequency of temporary sales only at the product level.



Figure 3.2: Demand Uncertainty and Temporary Sales



Note: This figure presents a scatterplot of temporary sales frequency (y-axis) and the standard deviation of log units (x-axis) at the store level. A regression fit on the data points is plotted in blue.

### 3.3.2.1.2 Specification

After computing our proxy for demand uncertainty, we estimate:

$$sale\_freq_k = \alpha + \beta SDU_k + \Gamma X_k + \varepsilon_k \quad (3.8)$$

where  $X_k = (SDP_k, q_k, p_k)$  is a vector of controls. We estimate Equation (3.8) for the store level ( $k = j$ ) and the product level ( $k = i$ ).<sup>5</sup> Our coefficient of interest,  $\beta$ , represents the effect of a 1 percentage point increase in the standard deviation of units on the sales frequency.

### 3.3.2.2 Results

Figure 3.2 presents a scatterplot of temporary sales frequency and the standard deviation of log units at the store level. The blue line plots an estimate of Equation (3.8) without other covariates. There is a clear positive relationship between the sale frequency and standard deviation of units. Specification (1) of Table 3.4 shows that a 1 percentage point increase in the standard deviation of units is associated with a 0.3 percentage point increase in the sale frequency for stores on average. The estimate of  $\beta$  reduces to 0.07 after controlling for the average of log units sold, the average of log price and the standard deviation of log price. This estimate remains statistically significant at the 95% level. Much of this effect appears to be captured by the standard

<sup>5</sup>Product-level variables are computed in a similar manner as the store-level variables.

Table 3.4: Demand Uncertainty and Temporary Sales

	<i>Dependent variable:</i>					
	Store-level ( <i>sale_freq_j</i> )		Product level ( <i>sale_freq_i</i> )		Store-product level ( <i>sale_freq_ij</i> )	
	(1)	(2)	(3)	(4)	(5)	(6)
SDU	0.313*** (0.032)	0.070** (0.033)	0.422*** (0.017)	0.245*** (0.019)	0.138*** (0.005)	0.081*** (0.006)
SDP		1.952*** (0.097)		1.162*** (0.052)		0.468*** (0.020)
Ave. ln(Units)		-0.001 (0.005)		0.010*** (0.003)		0.012*** (0.001)
Ave. ln(Price)		-0.006 (0.010)		-0.005** (0.002)		0.180*** (0.007)
UPC Fixed Effects					X	X
Store Fixed Effects					X	X
Observations	546	546	1,686	1,686	94,505	94,505
Adjusted R <sup>2</sup>	0.148	0.525	0.271	0.497	0.619	0.650

Note: This table presents regression estimates of Equations (3.8)-(3.9). SDU and SDP represent the standard deviations of log units and log prices, respectively. Standard errors are clustered at the store level for the product-store regressions in columns (5)-(6).

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01

deviation of prices. This suggests a correlation between *SDU* and *SDP*. Thus, stores may react to demand shocks by changing regular prices and by temporary sales.

Columns (3) and (4) present the estimates of Equation (3.8) at the product level. The product-level results also suggest a positive relationship between temporary sales and demand uncertainty. The coefficient for *SDU* is larger at the product level when compared to the store-level estimates. The product-level results suggest that stores issue more sales for products that experience more demand uncertainty.

One concern is that products may be on sale more often for reasons that are not related to demand uncertainty. Another concern is that differences in the characteristics and demographics of stores may lead to this result. To account for these possibilities, we estimate

$$sale\_freq_{ij} = \alpha_i + \gamma_j + \beta SDU_{ij} + \Gamma X_{ij} + \varepsilon_{ij} \quad (3.9)$$

where  $\alpha_i$  are product fixed effects and  $\gamma_j$  are store fixed effects. The product fixed effects control for differences in the means across products. The inclusion of store fixed effects allows for different pricing strategies across stores. Columns (5) and (6) of Table 3.4 present the results of Equation (3.9). Standard errors are clustered at the store level. The estimate of  $\beta$  is now 0.138 without controls and 0.081 with controls, and both remain statistically significant.

To address potential endogeneity problems, we re-estimate Equations (3.8) and (3.9) using the 2005 sample to measure the dependent variable and the 2004 sample to measure the independent variables. These results are presented in Appendix Table 3.7. The results are very similar to those in Table 3.4 suggesting that endogeneity is not a problem. Overall, our results suggest that demand uncertainty is positively correlated with temporary sales.

### 3.3.3 Temporary sales and demand shocks

The previous section provided evidence that stores with more demand uncertainty also issue more temporary sales. This suggests that stores use sales to react to demand shocks. We examine this hypothesis by exploiting the dynamic nature of the data to estimate several panel vector autoregressions (PVAR).

#### 3.3.3.1 Empirical Strategy

##### 3.3.3.1.1 Specification

We estimate a panel vector auto-regression of order two:


$$y_{ijt} = \alpha_{ij} + \gamma_{it} + A_1 y_{ijt-1} + A_2 y_{ijt-2} + \varepsilon_{ijt} \quad (3.10)$$

where  $y_{ijt}$  is a vector consisting of log price ( $p_{ijt}$ ), a sale indicator ( $sale_{ijt}$ ), and log quantity ( $q_{ijt}$ ) in that order for a store-product-week combination. The inclusion of product-store fixed effects,  $\alpha_{ij}$ , controls for differences in the average price, sales frequency, and units sold across stores for a given product. The product-time fixed effects,  $\gamma_{it}$ , control for product-specific seasonality and shocks.

##### 3.3.3.1.2 Interpretation

We use the estimated coefficients of Equation (3.10) to compute orthogonalized impulse response functions (IRFs) and simulate the effect of a quantity shock on temporary sales. The inclusion of the product-time fixed effects implies that the error term in (3.10) does not reflect shocks to the aggregated demand for the product and is due to store-specific shocks. Our IRFs therefore describe the response to a store-specific shock (that may vary across products). We will discuss the effects of shocks to the aggregated demand for products in the next section. In this context, the impulse response function  $k$  periods after a shock at time  $t$  ( $IRF_{t+k}$ ) represents how much *more* likely a given store is to issue a sale. For example, a store that experiences a one standard deviation negative demand shock in week  $t$  is  $IRF_{t+k}$  percentage points more likely to have a sale in period  $t+k$ .

Figure 3.3: Impulse Response of Temporary Sales to Demand Shocks



Note: This figure plots the orthogonalized impulse response function of temporary sales in response to a one standard deviation negative quantity shock. Confidence intervals are bootstrapped at the 95% confidence level.

### 3.3.3.2 Results

Figure 3 plots the orthogonalized impulse response function of temporary sales in response to a one standard deviation negative quantity shock. It describes the effect of a store-specific demand shock on the probability that a store will issue a temporary sale. Confidence intervals are bootstrapped at the 95% confidence interval. We see positive effects on the probability of sale in the first three weeks after a demand shock. The impulse response function shows that a store is 0.25 percentage points more likely to place a product on sale one week after a store-specific demand shock. This response peaks at 0.3pp two weeks after the shock, and begins to subside to about 0.1pp at three weeks. After three weeks, the effect converges to zero.

We provide several robustness checks for these results. We estimate Equation (3.10) with four lags rather than two. We change  $\gamma_{it}$  in Equation (3.10) to product-market-time fixed effects ( $\gamma_{imt}$ ). Lastly, we provide estimates for the alternative definition of a sale as a temporary price reduction that coincides with the IRI flag. In Table 3.5, rows represent the specification, and columns the period after shock. The last column presents the maximum cumulative effect for each specification.<sup>6</sup> Upper and lower forecast confidence bands bootstrapped at the 95% confidence level are presented in brackets.

We find similar estimates in the first two weeks after a demand shock using two and four lags when comparing the first and second rows of Table 3.5. However, at three weeks, we estimate a 0.29 percentage

<sup>6</sup>This is typically the fourth week after a demand shock, but may vary across specifications. The cumulative effect can be found by summing the individual response estimates.

Table 3.5: Impulse Response of Temporary Sales to Demand Shocks


	Weeks after Shock						Max. Cumulative Effect
	1	2	3	4	5	6	
<b>Baseline</b>							
2 Lags	0.26 [0.24, 0.27]	0.31 [0.29, 0.33]	0.13 [0.12, 0.14]	0.02 [0.01, 0.02]	-0.02 [-0.02, -0.01]	-0.02 [-0.02, -0.02]	0.71 [0.67, 0.76]
4 Lags	0.26 [0.24, 0.29]	0.39 [0.37, 0.41]	0.34 [0.32, 0.36]	0.05 [0.03, 0.07]	-0.02 [-0.03, 0]	-0.03 [-0.03, -0.02]	1.04 [1.00, 1.09]
Market Fixed Effects	0.22 [0.2, 0.24]	0.26 [0.24, 0.27]	0.10 [0.09, 0.11]	0.01 [0.01, 0.02]	-0.02 [-0.02, -0.01]	-0.02 [-0.02, -0.02]	0.59 [0.55, 0.63]
Market FEs & 4 Lags	0.23 [0.21, 0.25]	0.31 [0.29, 0.33]	0.28 [0.27, 0.30]	0.04 [0.02, 0.06]	-0.01 [-0.02, -0.01]	-0.03 [-0.03, -0.02]	0.87 [0.83, 0.92]
<b>10% Reduction &amp; IRI Flag</b>							
2 Lags	0.35 [0.33, 0.37]	0.33 [0.31, 0.35]	0.12 [0.12, 0.13]	0.02 [0.01, 0.02]	-0.02 [-0.02, -0.01]	-0.02 [-0.02, -0.02]	0.82 [0.78, 0.86]
4 Lags	0.36 [0.34, 0.38]	0.40 [0.38, 0.42]	0.32 [0.3, 0.34]	0.02 [0, 0.04]	-0.02 [-0.03, -0.01]	-0.03 [-0.03, -0.02]	1.10 [1.05, 1.14]
Market Fixed Effects	0.30 [0.29, 0.32]	0.27 [0.25, 0.29]	0.10 [0.09, 0.1]	0.01 [0.01, 0.01]	-0.01 [-0.02, -0.01]	-0.02 [-0.02, -0.01]	0.68 [0.65, 0.72]
Market FEs & 4 Lags	0.32 [0.30, 0.34]	0.32 [0.31, 0.34]	0.27 [0.25, 0.28]	0.02 [0, 0.03]	-0.02 [-0.03, -0.01]	-0.03 [-0.03, -0.02]	0.92 [0.88, 0.96]

Note: This table plots the orthogonalized impulse response functions of temporary sales in response to a one standard deviation negative quantity shock. Rows represent the specification, and columns specify the period after the shock. The last column presents the maximum cumulative effect for each specification. Forecast confidence intervals bootstrapped at the 95% confidence level are in brackets. The first row plots our baseline results. The second row plots the estimates using a VAR(4). The third and fourth rows are similar to first and second rows, but replace Product-Week fixed effects in Equation (3.10) with Product-Market-Week fixed effects. The bottom panels adds the restriction that a sale is also indicated by the IRI flag.

point increased probability of a sale using four lags compared to the baseline result of 0.14pp. As in the baseline results, impulse responses converge to zero after three weeks. This results in a maximum cumulative effect of 1.04pp with four lags compared to the baseline of 0.71pp. Including market fixed effects does not change the overall dynamics or relative effects in the impulse response functions. However, the magnitude of the effects are slightly mitigated. This results in a maximum cumulative effect of 0.59pp and 0.87pp for the estimates with two and four lag lengths, respectively.

The added restriction that a sale must coincide with the IRI flag increases the probability of a sale price in the first week after a demand shock across all specifications. Without the IRI flag, the peak response always occurred in the second week. With the IRI flag, the peak response varies across specifications between the first and second week. The increased response in the first week also leads to larger cumulative effects compared to our baseline specifications.

Figure 3.4: Impulse Response of Aggregate Sales Frequency to Demand Shocks



Note: This figure plots the orthogonalized impulse response function of temporary sales in response to a one standard deviation negative quantity shock from Equation 3.11. Confidence intervals are bootstrapped at the 95% confidence level.

### 3.4 Shocks to the Aggregated Demand for a Product

The previous section showed that sales respond to store-level demand shocks. We now show that this is also the case when the shock occurs to the aggregated demand for a product. We then assess the role of temporary sales in the response of the average price (across stores) to a demand shock.


#### 3.4.1 Specification

We estimate a panel vector autoregression of order two:

$$y_{it} = \alpha_{iw} + A_1 y_{it-1} + A_2 y_{it-2} + \varepsilon_{it} \tag{3.11}$$

where  $t$  indexes times and  $w$  represents the week of the year that period  $t$  belongs to. In our sample we have 104 weeks, so  $t \in \{1, \dots, 104\}$  and  $w \in \{1, \dots, 52\}$  since there are 52 weeks in a year. Thus,  $w = 1$  for weeks with time index  $t = 1$  and  $t = 53$ ,  $w = 2$  for  $t = 2$  and  $t = 54$  and so on. The product/week-of-year fixed effects,  $\alpha_{iw}$ , control for seasonality. The vector  $y_{it}$  consists of the log average price over stores ( $p_{it}$ ), the average sales frequency over stores ( $sale_{it}$ ), and log of the total quantity sold over stores ( $q_{it}$ ) in that order.<sup>7</sup>

Figure 3.5: Impulse Response of the Average Price to Demand Shocks



Note: This figure plots the orthogonalized impulse response function of all prices (blue) and non-sale prices (red) in response to a one standard deviation negative quantity shock from Equation 3.11. Confidence intervals are bootstrapped at the 95% confidence level.

### 3.4.2 Sales Response

Figure 3.4 presents the response of the aggregate sales frequency on average in response to a one standard deviation shock to the total units sold over all stores. The aggregate sales frequency increases by 0.3 percentage points one week after the shock. This effect remains greater than 0 percentage points for three weeks. After this, the response begins to converge to zero. Appendix Table 3.8 shows that a shock to the total demand for a product (over stores) has a maximum cumulative effect of 0.62 percentage points on the average sales probability (over stores). Robustness checks for alternative lag lengths and the definition of sales are also included in this table.

### 3.4.3 Contribution to Price Dynamics

We now explore how this increase in temporary sales affects the response of the average price (across stores) to demand shocks. To do this, we re-estimate Equation (3.11) using the log of the average non-sale/regular price over stores. The resulting impulse response function is then compared to the IRF when using all prices to compute the average price.<sup>8</sup>

Figure 3.5 plots the response of all prices (blue) and non-sale prices (red) to a one standard deviation negative shock to the quantity sold over all stores. Confidence intervals are bootstrapped at the 95% confidence level. As expected, prices decrease after a negative demand shock. However, this response is mitigated for non-sale prices. Non-sale prices experience a trough decline of 0.20% two weeks after a demand shock.

<sup>7</sup>To control for trending factors such as inflation, a product-specific linear trend is removed from  $p_{it}$  and  $q_{it}$  before estimating Equation (3.11).

<sup>8</sup>For these regressions, the sales vector is omitted from the specification.

Table 3.6: Impulse Response of Average Prices to Demand Shocks

	Weeks after Shock						Cumulative Effect (10 Weeks)	Contribution of Sales
	1	2	3	4	5	6		
<b>Baseline</b>								
<i>Two Lags</i>								
All Prices	-0.34 [-0.37, -0.31]	-0.41 [-0.43, -0.38]	-0.27 [-0.29, -0.25]	-0.16 [-0.18, -0.15]	-0.09 [-0.1, -0.08]	-0.05 [-0.06, -0.05]	-1.37 [-1.48, -1.28]	48.8%
Regular Prices	-0.09 [-0.11, -0.07]	-0.20 [-0.22, -0.18]	-0.15 [-0.17, -0.13]	-0.10 [-0.11, -0.09]	-0.06 [-0.07, -0.06]	-0.04 [-0.05, -0.04]	-0.70 [-0.79, -0.63]	
<i>Four Lags</i>								
All Prices	-0.30 [-0.33, -0.27]	-0.30 [-0.32, -0.27]	-0.31 [-0.34, -0.29]	-0.24 [-0.26, -0.21]	-0.19 [-0.21, -0.18]	-0.16 [-0.17, -0.14]	-1.88 [-2.02, -1.76]	44.9%
Regular Prices	-0.07 [-0.09, -0.05]	-0.15 [-0.18, -0.13]	-0.20 [-0.22, -0.18]	-0.16 [-0.18, -0.14]	-0.11 [-0.13, -0.1]	-0.09 [-0.11, -0.08]	-1.04 [-1.16, -0.93]	
<b>10% Reduction &amp; IRI Flag</b>								
<i>Two Lags</i>								
All Prices	-0.34 [-0.37, -0.31]	-0.41 [-0.43, -0.38]	-0.27 [-0.29, -0.25]	-0.16 [-0.18, -0.15]	-0.09 [-0.1, -0.08]	-0.05 [-0.06, -0.05]	-1.37 [-1.48, -1.28]	47.5%
Regular Prices	-0.08 [-0.1, -0.06]	-0.20 [-0.22, -0.18]	-0.15 [-0.17, -0.14]	-0.10 [-0.12, -0.09]	-0.07 [-0.08, -0.06]	-0.05 [-0.05, -0.04]	-0.72 [-0.81, -0.65]	
<i>4 Lags</i>								
All Prices	-0.30 [-0.33, -0.27]	-0.30 [-0.32, -0.27]	-0.31 [-0.34, -0.29]	-0.24 [-0.26, -0.21]	-0.19 [-0.21, -0.18]	-0.16 [-0.17, -0.14]	-1.88 [-2.02, -1.76]	44.4%
Regular Prices	-0.06 [-0.08, -0.04]	-0.16 [-0.18, -0.14]	-0.20 [-0.22, -0.18]	-0.16 [-0.18, -0.14]	-0.12 [-0.13, -0.1]	-0.10 [-0.11, -0.09]	-1.05 [-1.15, -0.95]	

Note: This table plots the orthogonalized impulse response functions of the average price for a product in response to a one standard deviation negative quantity shock. Rows represent the specification, and columns specify the period after the shock. The last two columns present the cumulative price response and the contribution of sales to the cumulative response after ten weeks, respectively. Forecast confidence intervals bootstrapped at the 95% confidence level are in brackets. The first row plots our baseline results. The second row plots the estimates using a VAR(4). The third and fourth rows are similar to first and second rows, but add the restriction that a sale is also indicated by the IRI flag.

Including sale prices estimates an effect of 0.41%. The response of all prices is larger in magnitude than non-sale prices for the first six weeks following the shock. After six weeks, both impulse response functions converge to zero.

Panel (b) shows that the cumulative price response for all prices remains greater (in absolute value) than non-sale prices up to ten weeks following the shock. After 10 weeks, the cumulative response of all prices is 1.37% while the response of regular prices is 0.70%. The relative contribution of sales to the cumulative price response after 10 weeks is then  $\frac{1.37-0.70}{1.37} = 49\%$ . Thus, sales account for almost half of the average price decline for a product following a demand shock.

Table 3.6 presents several robustness checks for these results. We estimate Equation (3.11) with four lags rather than two. We provide estimates for the alternative definition of a sale as a temporary price reduction that coincides with the IRI flag. In Table 3.6, rows represent the specification, and columns the period after shock. The last column presents the maximum cumulative effect for each specification. Upper and lower forecast confidence bands bootstrapped at the 95% confidence level are presented in brackets.

We see that using four lags in the VAR leads to larger cumulative decreases in both the all price and



non-sale price response functions. This leads to a cumulative response after ten weeks of -1.88% and -1.04% for all prices and regular prices, respectively. However, this has a small effect on the relative contribution of sales to the total price response which has changed from 48.8% to 44.9%. Similarly, using the alternative definition of sales does not significantly affect the contribution to the cumulative price response.

Overall, the results suggest that temporary sales play a significant role in the response to demand shocks.

### 3.5 A Model

To account for our empirical findings, we present an Uncertain and Sequential Trade (UST) model in which sales are determined endogenously in response to demand shocks. The first subsection presents the baseline UST model. The second subsection extends the UST model to allow for storage of goods. The final subsection uses the model to account for the stylized facts.

#### 3.5.1 Uncertain and Sequential Trade Model

##### 3.5.1.1 Uncertain

In the UST model, the economy is comprised of many consumers and firms. Each firm makes their production decision at the beginning of the period *before* the number of consumers is known. We assume the number of consumers for good  $i$  in period  $t$  is given by  $N_{it}$  where  $N_{it}$  is an iid random variable. For simplicity of the exposition, we describe the scenario in which the number of consumers can take two possible realizations:


$$N_{it} = \begin{cases} x_i & \text{with prob. } \pi_i, \\ x_i + \Delta_i & \text{with prob. } 1 - \pi_i. \end{cases}$$

All consumers have the same demand function. An individual consumer who buys good  $i$  in store  $j$  at price  $P_{ijt}$  demands  $D_i(P_{ijt})$  units of the good.

##### 3.5.1.2 Sequential

Consumers arrive sequentially in the model. The first group of  $x_i$  consumers arrive early with certainty. The second group of  $\Delta_i$  consumers arrive later with probability  $1 - \pi_i$ . It is useful to think of two hypothetical markets. The first market (*Early*) opens with certainty and serves the group of  $x_i$  consumers. The second market (*Late*) opens with probability  $1 - \pi_i$  and serves the second batch of  $\Delta_i$  consumers if they arrive. Firms take the price in each of the hypothetical markets as given. Thus, firms can sell good  $i$  at the price  $P_{i,Early,t}$  with certainty or at the price  $P_{i,Late,t}$  if the additional consumers arrive. Firms that sell at the early price supply the market with  $Q_{i,Early,t}$  units of good  $i$  while firms that sell at the late price supply  $Q_{i,Late,t}$  units.

Figure 3.6: Prices and Quantities in the Baseline UST Model



Note: This figure plots a possible equilibrium in the baseline Uncertain and Sequential Trade model. Prices and quantities are denoted by  $P$  and  $Q$  respectively. Subscripts denote *Late* and *Early* markets.

### 3.5.1.3 Equilibrium

The cost of production is  $\lambda_i$  per unit. Equilibrium is then defined as a vector of prices and quantities  $(P_{i,Early,t}, P_{i,Late,t}, Q_{i,Early,t}, Q_{i,Late,t})$  such that the expected profits for each unit is 0:

$$P_{i,Early,t} = (1 - \pi_i)P_{i,Late,t} = \lambda_i \quad (3.12)$$

And markets that open are cleared:

$$Q_{i,Early,t} = x_i D_i(P_{i,Early,t}) \quad (3.13)$$

$$Q_{i,Late,t} = \Delta_i D_i(P_{i,Late,t}) \quad (3.14)$$

Thus, in equilibrium, firms are indifferent between the two prices as the expected profits are the same.

Figure 3.6 plots a possible equilibrium solution for a given good. The product and time subscripts are omitted for simplicity. At the price  $\lambda$ , the total demand of the first group of consumers,  $xD(\lambda)$ , is equal to the quantity supplied,  $Q_{Early}$ . If the second group of consumers arrives to the market, they purchase the good at price  $\frac{\lambda}{1-\pi}$ . Again, markets clear and quantity demanded is equal to quantity supplied,  $\Delta D(\frac{\lambda}{1-\pi}) = Q_{Late}$ . Note that, in this simple version of the UST model, prices do not change over time (Equation 3.12). The quantity sold at the low price also does not change over time. However, the quantity sold at the high price

fluctuates between  $Q_{Late}$  and 0 depending on the realization of demand.

### 3.5.2 Extended Model with Temporary Sales

Bental and Eden (1993) develop a UST model that allows for storage. Including storage allows prices to fluctuate as a result of demand shocks. In the model, a negative demand shock leads to the accumulation of inventories and a reduction in future prices.<sup>9</sup> In what follows, we describe the one-good model under the assumptions of a constant per unit cost ( $\lambda$ ) and two states of demand as in the previous subsection. Following Eden (2018) we assume that goods have an expiration date and follow one-hoss shay depreciation which is a natural assumption for our analysis of supermarket goods. For simplicity, we assume that the good can be stored for one period only. Thus, if a good is not sold in the first period of its life, it can still be sold in the second period. After these two periods, the good then has no value.

With the inclusion of storage, the economy at the beginning of period  $t$  can now be in one of two states: inventories ( $Inv$ ) or no inventories ( $NoInv$ ). In state  $Inv$ , the demand in the previous period was low ( $N_{t-1} = x$ ) and the additional  $\Delta$  consumers did not arrive. Thus, inventories are carried from the previous period as a result. In state  $NoInv$ , demand was high ( $N_{t-1} = x + \Delta$ ) and there are no inventories. Prices are now a function of last period's demand,  $P_j(N_{t-1})$ . This extends the equilibrium vector of prices in the baseline UST model to  $(P_{Early}^{Inv}, P_{Early}^{NoInv}, P_{Late}^{Inv}, P_{Late}^{NoInv})$ . Similarly, the equilibrium vector of quantities is now extended to  $(Q_{Early}^{Inv}, Q_{Early}^{NoInv}, Q_{Late}^{Inv}, Q_{Late}^{NoInv})$ . Inventories at the beginning of the period are given by:

$$I = \begin{cases} \Delta D(P_{Late}) & \text{if } N_{t-1} = x, \\ 0 & \text{otherwise.} \end{cases} \quad (3.15)$$

These inventories are allocated over the two hypothetical markets:

$$I = q_{Early}^{Stored} + q_{Late}^{Stored} \quad (3.16)$$

The supply to each of the hypothetical markets is the sum of newly produced and stored units:

$$Q_{Early}^{Inv} = q_{Early}^{Stored} + q_{Early}^{New} \quad (3.17)$$

$$Q_{Late}^{Inv} = q_{Late}^{Stored} + q_{Late}^{New} \quad (3.18)$$

A newly produced unit that was not sold in the current period can be sold in the next period. Since stored units will expire if they are not sold, we assume that stored units are supplied to the *Early* market first. They

<sup>9</sup>This is consistent with Aguirregabiria (1999) who finds that markups are negatively correlated with inventory levels.

are supplied to the *Late* market only if the quantity of inventories exceeds the demand in the *Early* market. Thus, the value of an individual stored unit can be given by  $\beta P_{Early}^{Inv}$  where  $\beta$  is a discount factor that captures discounting, storage costs, and depreciation with the restriction that  $\beta \in (0, 1)$ .

### 3.5.2.1 Equilibrium

The full equilibrium vector is  $(I, P_{Early}^{Inv}, P_{Early}^{NoInv}, P_{Late}^{Inv}, P_{Late}^{NoInv}, Q_{Early}^{Inv}, Q_{Early}^{NoInv}, Q_{Late}^{Inv}, Q_{Late}^{NoInv}, q_{Early}^{Stored}, q_{Early}^{New}, q_{Late}^{Stored}, q_{Late}^{New})$  which must satisfy Equations (3.15)-(3.18) in addition to the following conditions:

1. Prices and quantities in the *Late* consumer market do not depend on the state of inventories:

$$P_{Late} = P_{Late}^{Inv} = P_{Late}^{NoInv} \quad (3.19)$$

$$Q_{Late} = Q_{Late}^{Inv} = Q_{Late}^{NoInv} \quad (3.20)$$

2. If the level of inventories is positive, stored units are supplied to the *Early* consumer market first:

$$q_{Early}^{Stored} = \min \{ Q_{Early}^{Inv}, I \} \quad (3.21)$$

$$P_{Early}^{Inv} \geq (1 - \pi) P_{Late} \quad \text{with equality if } q_{Late}^{Stored} > 0 \quad (3.22)$$

Equation (3.21) states that stored units are supplied to the *Late* consumer market only if the level of inventories is larger than the demand of the *Early* consumer market. Equation (3.22) requires that supplying stored goods to the *Early* consumer market is optimal for firms. When stored goods are supplied to both markets ( $q_{Late}^{Stored} > 0$ ), the price in the *Early* market must equal the expected price in the *Late* market. Otherwise, it is optimal to allocate all stored units to the *Early* market.

3. Price in the *Early* consumer market must be less than or equal to  $\lambda$  with equality if new units are supplied:

$$P_{Early}^{Inv} \leq \lambda \quad \text{with equality if } q_{Early}^{New} > 0 \quad (3.23)$$

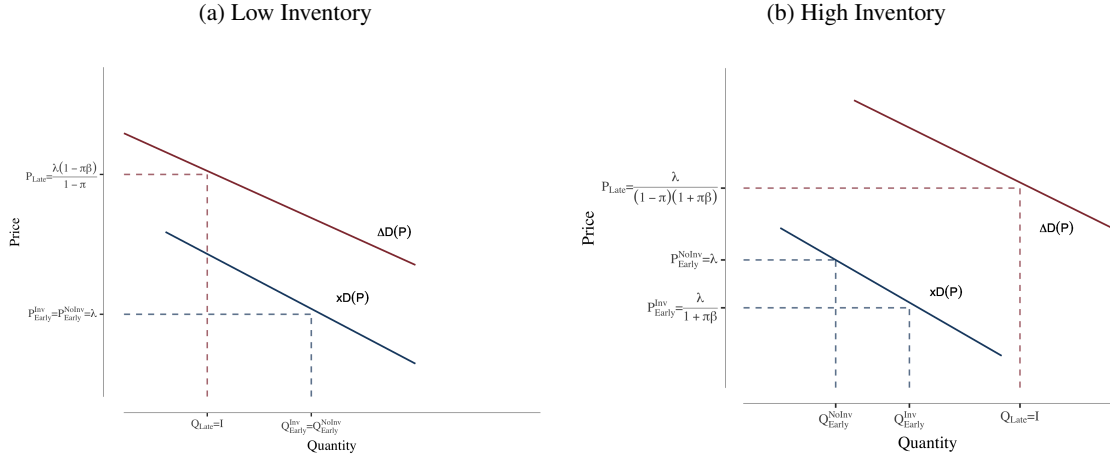
$$P_{Early}^{NoInv} = \lambda \quad (3.24)$$

4. The expected marginal revenue of supplying a newly produced good to the *Late* consumer market is equal to marginal cost:

$$(1 - \pi) P_{Late} + \pi \beta P_{Early}^{Inv} = \lambda \quad (3.25)$$

To gain the intuition, this was simply  $(1 - \pi) P_{Late} = \lambda$  in the UST model without storage. With storage, the firm may still sell the good in the next period if the *Late* market does not open in the current period. Thus, expected revenues also include the value of inventories mentioned earlier,  $\pi \beta P_{Early}^{Inv}$ .

Figure 3.7: Possible Equilibria in the UST Model with Storage



Note: This figure plots possible equilibria in the extended Uncertain and Sequential Trade model with storage. Panels A and B represent scenarios in which the inventory level is lower/higher than the quantity demanded in the *Early* market, respectively. Prices and quantities are denoted by  $P$  and  $Q$  respectively. Subscripts denote *Late* and *Early* markets. Superscripts denote the state of inventories where *Inv* represents a positive amount of inventories and *NoInv* represents no inventories.

#### 5. Markets clear:

$$Q_{Early}^{Inv} = xD(P_{Early}^{Inv}) \quad (3.26)$$

$$Q_{Early}^{NoInv} = xD(P_{Early}^{NoInv}) \quad (3.27)$$

$$Q_{Late} = \Delta D(P_{Late}) \quad (3.28)$$

Figure 3.7 plots two possible equilibria. As in the previous subsection, the blue and red lines depict the *Late* and *Early* consumer markets, respectively. Panel A depicts the case in which both newly produced goods and stored goods are supplied to the *Early* market if inventories are carried. In this case, the price in the *Early* market is equal to marginal cost regardless of the state of inventories,  $P_{Early}^{Inv} = P_{Early}^{NoInv} = \lambda$ . As a result, the quantities in the *Early* market do not depend on the level of inventories,  $Q_{Early}^{Inv} = Q_{Early}^{NoInv}$ . We can determine the price in the *Late* market by solving the equilibrium condition (3.25),  $P_{Late} = \frac{\lambda(1-\pi\beta)}{1-\pi}$ . Lastly, the equilibrium quantity in the *Late* market is less than the quantity in the *Early* market.

Panel B of Figure 3.7 depicts the case in which inventories are greater than the quantity demanded in the *Early* market. Since stored units are supplied to both markets, we must have  $(1-\pi)P_{Late} = P_{Early}^{Inv}$ .

Substituting into Equation 3.25 and solving for prices yields:

$$P_{Early}^{Inv} = \frac{\lambda}{1 + \pi\beta} \quad (3.29)$$

$$P_{Late} = \frac{\lambda}{(1 - \pi)(1 + \pi\beta)} \quad (3.30)$$

The price in the *Early* market is equal to marginal cost if inventories are not carried,  $P_{Early}^{NoInv} = \lambda$ . Thus, unlike Panel A, this equilibrium represents a situation in which the price and quantity in the *Early* market vary with the state of inventories.<sup>10</sup>

### 3.5.3 Accounting for the stylized facts

Sections 3.3 and 3.4 presented several stylized facts about sales. First, the fraction of weeks in which there is no sale in any store is larger than predicted by using a mixed strategy. Second, stores and products that face more demand uncertainty have more sales. Third, sales play an important role in the response to demand shocks. In this section we use the UST model to discuss these findings.

In the UST model, stores that post a high price accumulate unwanted inventories when the realization of demand is low. The accumulated inventories are typically offered at a low price, so that they will be sold before reaching the expiration date. After the store sells all the “unwanted” inventories it may switch to the high price, so the reduction in price may be temporary and may satisfy our definition of temporary sale.

In our model, there are periods of high demand. After a period of high demand, there are no “unwanted” inventories and stores do not reduce prices. This is consistent with the observation that in many weeks there are no temporary sales in any of the stores.

Price dispersion in the UST model arises in response to demand uncertainty. When the distribution of demand is degenerate we get the standard Walrasian equilibrium with no accumulation of “unwanted inventories” and no temporary sales. This is the intuition behind the result in Eden (2018) who showed that when the distribution of aggregate demand is close to uniform, more demand uncertainty leads to more temporary sales.

Here we extend the results in Eden (2018) to the store level. In our model, newly produced goods are typically allocated to both (high and low price) markets. In the no inventories state that occurs after the realization of high demand, stores are indifferent between the two markets and some stores may choose the low price and some may choose the high price. If this choice is correlated over time and some stores consistently choose the high price in the no inventories state and some consistently choose the low price, we

<sup>10</sup>It is also possible that only stored units are allocated to the *Early* market and only new units are supplied to the *Late* market. This is a special scenario of the equilibrium depicted in Panel B.

may find a correlation between demand uncertainty and the frequency of temporary sales also on the store level. Stores that choose the high price for newly produced goods will experience more fluctuations in the amount they sell because the high price market opens only when the realization of demand is high. These stores will also accumulate unwanted inventories when demand is low and will offer the stored goods at a temporary sale price.

In our regressions, average price does not seem to play a significant role in explaining the variation of the frequency of sales over stores. This suggests that variation in demand uncertainty across stores is not fully explained by differences in pricing strategies. We may assume that some buyers do not search over stores, so that each store has a group of non-shoppers that are loyal to it. If the demand of non-shoppers fluctuates in some stores more than in other stores, the stores that experience more fluctuations in the demand of non-shoppers will tend to accumulate unwanted inventories more often and will tend to have temporary sales more often.

The model implies that after a negative demand shock stores will have more unwanted inventories and more sales. Importantly, sales do not cancel out when aggregating over stores. Thus, the average sales frequency over stores also increases following a demand shock. This increase in temporary sales leads to a reduction in the average price. This is consistent with the impulse response functions we estimate.

### **3.6 Conclusion**

It has become increasingly standard in macroeconomics to remove sale observations and analyze regular/non-sale prices. We argue that temporary sales play an important role in the response to demand shocks. Using scanner data, we find that the average price of a product decreases by almost 1.5% following a shock to the total quantity sold across stores. If sale observations are removed and the average regular price is analyzed instead, then the cumulative response of prices is cut in half. Thus, ignoring sales underestimates the price reaction to a demand shock and overestimates the degree of price rigidity.

We then develop a model to account for this finding guided by several stylized facts seen in the data. We find that in about 45% of the weeks, an average item is not on sale in any store. In our model there are no temporary sales in any of the stores after a period of high demand. We also find a correlation between demand uncertainty and the frequency of sales at the store and product levels. In the model, products with more demand uncertainty should have more sales. Furthermore, stores that adopt a strategy of charging a high price for newly produced goods will experience more fluctuations in the quantity sold and will have more sales. We also find that the correlation between the frequency of sales and demand uncertainty holds even after controlling for the average price of the store. This suggests that demand uncertainty by non-shoppers plays an important role.

### 3.7 Appendix

#### 3.7.1 Empirics

Table 3.7: Demand Uncertainty and Temporary Sales (Regressions that address Potential Endogeneity Issues)

	<i>Dependent variable:</i>					
	Store-level ( <i>sale_freq_j</i> )		Product level ( <i>sale_freq_i</i> )		Store-product level ( <i>sale_freq_ij</i> )	
	(1)	(2)	(3)	(4)	(5)	(6)
SDU	0.294*** (0.029)	0.110*** (0.033)	0.432*** (0.018)	0.260*** (0.020)	0.114*** (0.005)	0.072*** (0.005)
SDP		1.392*** (0.092)		1.149*** (0.055)		0.361*** (0.017)
Ave. ln(Units)		-0.001 (0.006)		0.004 (0.004)		0.013*** (0.001)
Ave. ln(Price)		-0.007 (0.011)		-0.010*** (0.002)		0.161*** (0.009)
UPC Fixed Effects					X	X
Store Fixed Effects					X	X
Observations	546	546	1,686	1,686	94,505	94,505
Adjusted R <sup>2</sup>	0.159	0.421	0.256	0.463	0.556	0.577

Note: This table presents regression estimates of Equations (3.8)-(3.9) that address potential endogeneity issues. This table repeats the regressions of Table 3.4, but here we use the 2005 sample to measure the dependent variable and the 2004 sample to measure the independent variables. SDU and SDP represent the standard deviations of log units and prices, respectively. Standard errors are clustered at the store level for the product-store regressions in columns (5)-(6).

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01



Table 3.8: Aggregate Impulse Response of Temporary Sales to Demand Shocks

	Weeks after Shock						Max. Cumulative Effect
	1	2	3	4	5	6	
<b>Baseline</b>							
2 Lags	0.34 [0.29, 0.39]	0.22 [0.17, 0.28]	0.06 [0.02, 0.09]	-0.03 [-0.05, -0.01]	-0.05 [-0.05, -0.04]	-0.04 [-0.04, -0.03]	0.62 [0.52, 0.74]
4 Lags	0.34 [0.29, 0.39]	0.22 [0.17, 0.27]	0.14 [0.09, 0.2]	-0.03 [-0.08, 0.02]	-0.03 [-0.07, 0]	-0.02 [-0.04, 0]	0.70 [0.58, 0.82]
<b>10% Reduction &amp; IRI Flag</b>							
2 Lags	0.41 [0.35, 0.46]	0.22 [0.17, 0.28]	0.05 [0.02, 0.09]	-0.03 [-0.05, -0.02]	-0.04 [-0.05, -0.03]	-0.03 [-0.04, -0.03]	0.68 [0.58, 0.80]
4 Lags	0.40 [0.36, 0.47]	0.23 [0.18, 0.29]	0.13 [0.08, 0.19]	-0.06 [-0.12, -0.01]	-0.04 [-0.07, -0.01]	-0.02 [-0.04, -0.01]	0.77 [0.66, 0.89]

Note: This table plots the orthogonalized impulse response functions of the aggregated sales frequency for a product in response to a one standard deviation negative shock to the total quantity sold. Rows represent the specification, and columns specify the period after the shock. The last column presents the maximum cumulative effect for each specification. Forecast confidence intervals bootstrapped at the 95% confidence level are in brackets. The first row plots our baseline results. The second row plots the estimates using a VAR(4). The third and fourth rows are similar to first and second rows, but add the restriction that a sale is also indicated by the IRI flag.

## References

- Abraham, S. and Sun, L. (2020). Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects. *Journal of Econometrics*.
- Acemoglu, D., Chernozhukov, V., Werning, I., and Whinston, M. (2021). Optimal Targeted Lockdowns in a Multi-Group SIR Model. *American Economic Review: Insights*, Forthcoming:–.
- Adams, B. and Williams, K. R. (2019). Zone pricing in retail oligopoly. *American Economic Journal: Microeconomics*, 11(1):124–156.
- Adda, J. (2016). Economic activity and the spread of viral diseases: Evidence from high frequency data. *Quarterly Journal of Economics*, 131(2):891–941.
- Aguirregabiria, V. (1999). The Dynamics of Markups and Inventories in Retailing Firms. *Review of Economic Studies*.
- Allcott, H., Boxell, L., Conway, J. C., Ferguson, B. A., Gentzkow, M., and Goldman, B. (2020). What Explains Temporal and Geographic Variation in the Early US Coronavirus Pandemic? *NBER Working Paper No. 27965*.
- Alvarez, F., Le Bihan, H., and Lippi, F. (2016). The real effects of monetary shocks in sticky price models: A sufficient statistic approach. *American Economic Review*, 106(10):2817–2851.
- Alvarez, F. and Lippi, F. (2014). Price Setting With Menu Cost for Multiproduct Firms. *Econometrica*, 82(1):89–135.
- Anderson, E., Malin, B. A., Nakamura, E., Simester, D., and Steinsson, J. (2017). Informational rigidities and the stickiness of temporary Sales. *Journal of Monetary Economics*, 90:64–83.
- Angulo, J., Yu, H. L., Langousis, A., Kolovos, A., Wang, J., Madrid, A. E., and Christakos, G. (2013). Spatiotemporal Infectious Disease Modeling: A BME-SIR Approach. *PLoS ONE*, 8(9).
- Antràs, P., Redding, S. J., and Rossi-Hansberg, E. (2020). Globalization and Pandemics. *NBER Working Paper No. 27840*.
- Atkeson, A. (2020). What Will be the Economic Impact of Covid-19 in the US? Rough Estimates of Disease Scenarios. *NBER Working Paper 26867*.
- Baek, C., McCrory, P. B., Messer, T., and Mui, P. (2020). Unemployment Effects of Stay-at-Home Orders: Evidence from High Frequency Claims Data. *Review of Economics and Statistics*, Forthcoming.
- Baharad, E. and Eden, B. (2004). Price rigidity and price dispersion: Evidence from micro data. *Review of Economic Dynamics*, 7(3):613–641.
- Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., and Yannelis, C. (2020). How Does Household Spending Respond to an Epidemic? Consumption during the 2020 COVID-19 Pandemic. *The Review of Asset Pricing Studies*.
- Bartik, A. W., Bertrand, M., Lin, F., Rothstein, J., and Unrath, M. (2020). Measuring the Labor Market at the Onset of the Covid-19 Crisis. *Brooking Papers on Economic Activity*, Forthcoming.
- Beland, L.-P., Brodeur, A., and Wright, T. (2020). COVID-19, Stay-At-Home Orders and Employment: Evidence from CPS Data. *IZA Working Paper*.
- Bental, B. and Eden, B. (1993). Inventories in a Competitive Environment. *Journal of Political Economy*.
- Bental, B. and Eden, B. (1996). Money and inventories in an economy with uncertain and sequential trade. *Journal of Monetary Economics*, 37:445–459.

- Bils, M. and Klenow, P. J. (2004). Some Evidence on the Importance of Sticky Prices. *Journal of Political Economy*, 112(5).
- Binder, C. (2020). Coronavirus Fears and Macroeconomic Expectations. *Review of Economics and Statistics*, Forthcoming.
- Bisin, A. and Moro, A. (2020). Learning Epidemiology by Doing: The Empirical Implications of a Spatial-SIR Model with Behavioral Responses. *NBER Working Paper No. 27590*.
- BLS (2020). Supplemental data measuring the effects of the coronavirus (COVID-19) pandemic on the labor market. <https://www.bls.gov/cps/effects-of-the-coronavirus-covid-19-pandemic.htm>.
- Bognanni, M., Hanley, D., Kolliner, D., and Mitman, K. (2020). Economics and Epidemics: Evidence from an Estimated Spatial Econ-SIR Model. *IZA Working Paper No. 13797*.
- Bonomo, M., Carvalho, C., Kryvtsov, O., Ribon, S., and Rigato, R. (2020). Multi-Product Pricing: Theory and Evidence From Large Retailers. *Working Paper*.
- Bronnenberg, B. J., Kruger, M. W., and Mela, C. F. (2008). The IRI marketing data set. *Marketing Science*, 27(4):745–748.
- Caballero, R. J. and Engel, E. M. (2007). Price Stickiness in Ss Models: New Interpretations of Old Results. *Journal of Monetary Economics*, 54:100–121.
- Callaway, B. and Sant’Anna, P. H. (2020). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, Forthcoming.
- Calvo, G. A. (1983). Staggered Prices in a Utility-Maximizing Framework. *Journal of Monetary Economics*, 12:383–398.
- Caplin, A. S. and Spulber, D. F. (1987). Menu Costs and the Neutrality of Money. *The Quarterly Journal of Economics*, 102(4):703–725.
- Carvalho, C. and Kryvtsov, O. (2021). Price Selection. *Journal of Monetary Economics*, 122:56–75.
- Carvalho, C. and Schwartzman, F. (2015). Selection and monetary non-neutrality in time-dependent pricing models. *Journal of Monetary Economics*, 76:141–156.
- Cavallo, A. (2017). Are online and offline prices similar? Evidence from large multi-channel retailers. *American Economic Review*, 107(1):283–303.
- Cavallo, A. (2019). More Amazon Effects: Online Competition and Pricing Behaviors. *Jackson Hole Economic Symposium Conference Proceedings (Federal Reserve Bank of Kansas City)*.
- Cavallo, A. (2020). Inflation with Covid Consumption Baskets. *NBER Working Paper No. 27352*.
- Chetty, R., Friedman, J. N., Hendren, N., and Stepner, M. (2020). How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data. *NBER Working Paper No. 27431*.
- Chevalier, J. A. and Kashyap, A. K. (2019). Best prices: Price discrimination and consumer substitution. *American Economic Journal: Economic Policy*, 11(1):126–159.
- Coibion, O., Gorodnichenko, Y., and Hong, G. H. (2015). The cyclicity of sales, regular and effective prices: Business cycle and policy implications. *American Economic Review*, 105(3):993–1029.
- Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). The Cost of the Covid-19 Crisis: Lockdowns, Macroeconomic Expectations, and Consumer Spending. *NBER Working Paper No. 27141*.
- Crucini, M. J. and Telmer, C. I. (2020). Microeconomic sources of real exchange rate variation. *Review of Economic Dynamics*, 38:22–40.

- Dana, J. D. (1998). Advance-purchase discounts and price discrimination in competitive markets. *Journal of Political Economy*, 106(2):395–422.
- Dana, J. D. (1999). Equilibrium Price Dispersion under Demand Uncertainty: The Roles of Costly Capacity and Market Structure. *The RAND Journal of Economics*, 30(4):632–660.
- Daruich, D. and Kozlowski, J. (2021). Macroeconomic Implications of Uniform Pricing. *Federal Reserve Bank of St. Louis Working Paper 2019-024*.
- DellaVigna, S. and Gentzkow, M. (2019). Uniform Pricing in US Retail Chains. *Quarterly Journal of Economics*, 134(4):2011–2084.
- Deneckere, R. and Peck, J. (2012). Dynamic Competition With Random Demand and Costless Search: A Theory of Price Posting. *Econometrica*, 80(3):1185–1247.
- Dotsey, M., King, R. G., and Wolman, A. L. (1999). State-Dependent Pricing and the General Equilibrium Dynamics of Money and Output. *The Quarterly Journal of Economics*, 114(2):655–690.
- Dotsey, M. and Wolman, A. (2018). Inflation and Real Activity with Firm Level Productivity Shocks. *FRB of Philadelphia Working Paper*, (18-19).
- Eden, B. (1990). Marginal Cost Pricing When Spot Markets Are Complete. *Journal of Political Economy*, 98(6):1293–1306.
- Eden, B. (1994). The Adjustment of Prices to Monetary Shocks when Trade is Uncertain and Sequential. *Journal of Political Economy*.
- Eden, B. (2018). Price Dispersion and Demand Uncertainty: Evidence from U.S. scanner data. *International Economic Review*, 59(3):1035–1075.
- Eden, B., Eden, M., and O’Flaherty, O. (2021). Temporary Sales in Response to Demand Shocks. *Working Paper*.
- Eichenbaum, M., Jaimovich, N., and Rebelo, S. (2011). Reference prices, costs, and nominal rigidities. *American Economic Review*, 101(1):234–262.
- Eichenbaum, M. S., Rebelo, S., and Trabandt, M. (2020). The Macroeconomics of Epidemics. *NBER Working Paper No. 26882*.
- Fajgelbaum, P., Khandelwal, A., Kim, W., Mantovani, C., and Schaal, E. (2020). Optimal Lockdown in a Commuting Network. *American Economic Review: Insights*, Forthcoming.
- Fox News (2020). Coronavirus shutdown: What states have seen protests against stay-at-home orders. <https://www.foxnews.com/us/coronavirus-shutdown-what-states-have-proteted-against-stay-at-home-orders>.
- Gagnon, E. and López-Salido, D. (2020). Small Price Responses to Large Demand Shocks. *Journal of the European Economic Association*, 18(2):792–828.
- Gertler, M. and Leahy, J. (2008). A Phillips curve with an Ss foundation. *Journal of Political Economy*, 116(3):533–572.
- Glandon, P. J. (2018). Sales and the (Mis)measurement of price level fluctuations. *Journal of Macroeconomics*, 58:60–77.
- Golosov, M. and Lucas, R. E. (2007). Menu costs and phillips curves. *Journal of Political Economy*, 115(2):171–199.
- Goodman-Bacon, A. (2018). Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics*, Forthcoming.

- Goodman-Bacon, A. and Marcus, J. (2020). Using Difference-in-Differences to Identify Causal Effects of COVID-19 Policies. *Survey of Research Methods*, 14(2).
- Goolsbee, A. and Syverson, C. (2020). Fear, Lockdown, and Diversion: Comparing Drivers of Pandemic Economic Decline 2020. *Journal of Public Economics*, 193.
- Guimaraes, B. and Sheedy, K. D. (2011). Sales and monetary policy. *American Economic Review*, 101(2):844–876.
- Gupta, S., Montenegro, L., Nguyen, T. D., Rojas, F. L., Schmutte, I. M., Simon, K. I., Weinberg, B. A., and Wing, C. (2020). Effects of Social Distancing Policy on Labor Market Outcomes. *NBER Working Paper No. 27280*.
- Hee Hong, G., Klepacz, M., Pasten, E., and Schoenle, R. (2021). The Real Effects of Monetary Shocks: Evidence from Micro Pricing Moments. *Working Paper*.
- Hitsch, G. J., Hortaçsu, A., and Lin, X. (2019). Prices and Promotions in U.S. Retail Markets: Evidence from Big Data. *NBER Working Paper No. 26306*.
- Jones, C. J., Philippon, T., and Venkateswaran, V. (2020). Optimal Mitigation Policies in a Pandemic: Social Distancing and Working from Home. *NBER Working Paper*.
- Kaplan, G., Menzies, G., Rudanko, L., and Trachter, N. (2016). Relative Price Dispersion: Evidence and Theory. *American Economic Journal: Microeconomics*, 11(3):68–124.
- Karadi, P. and Reiff, A. (2019). Menu Costs, Aggregate Fluctuations and Large Shocks. *American Economic Journal: Macroeconomics*, 11(2):111–146.
- Kehoe, P. and Midrigan, V. (2015). Prices are sticky after all. *Journal of Monetary Economics*, 75:35–53.
- Kermack, W. O. and McKendrick, A. G. (1927). A Contribution to the Mathematical Theory of Epidemics. *Proceedings of the Royal Society of London*.
- Klenow, P. J. and Kryvtsov, O. (2008). State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation? *Quarterly Journal of Economics*, 123(3):863–904.
- Kong, E. and Prinz, D. (2020). Disentangling policy effects using proxy data: Which shutdown policies affected unemployment during the COVID-19 pandemic? *Journal of Public Economics*, 189.
- Kryvtsov, O. and Vincent, N. (2021). The Cyclicity of Sales and Aggregate Price Flexibility. *The Review of Economic Studies*, 88(1):334–377.
- Lin, Z. and Meissner, C. (2020). Health vs. Wealth? Public Health Policies and the Economy During Covid-19. *NBER Working Paper No. 27099*.
- Midrigan, V. (2011). Menu Costs, Multiproduct Firms, and Aggregate Fluctuations. *Econometrica*, 79(4):1139–1180.
- Nakamura, A. O., Nakamura, E., and Nakamura, L. I. (2011). Price dynamics, retail chains and inflation measurement. *Journal of Econometrics*, 161(1):47–55.
- Nakamura, E. (2008). Pass-Through in Retail and Wholesale. *American Economic Review: Papers & Proceedings*, 98(2):430–437.
- Nakamura, E. and Steinsson, J. (2008). Five Facts About Prices: A reevaluation of menu cost models. *Quarterly Journal of Economics*, 123(4):1415–1464.
- Nakamura, E. and Steinsson, J. (2010). Monetary Non-neutrality in a Multisector Menu Cost Model. *Quarterly Journal of Economics*, 125(3):961–1013.
- Nakamura, E. and Steinsson, J. (2013). Price rigidity: Microeconomic evidence and macroeconomic implications. *Annual Review of Economics*, 5:133–163.

- NBC News (2020). 'Fire Gruesome Newsom!': Stay-at-home protests in California and across the country. <https://www.nbcnews.com/politics/politics-news/re-open-now-stay-home-protests-held-across-country-n1198121>.
- New York Times (2020). Travel From New York City Seeded Wave of U.S. Outbreaks. <https://www.nytimes.com/2020/05/07/us/new-york-city-coronavirus-outbreak.html>.
- NJ.com (2020). 'Open New Jersey now!' Protesters in Trenton demand Gov. Murphy lift coronavirus lockdowns despite rising death toll. <https://www.nj.com/coronavirus/2020/04/open-jersey-now-protestors-in-trenton-call-on-gov-murphy-to-lift-coronavirus-lockdowns.html>.
- Politico (2020). Fauci endorses national stay-at-home order: 'I just don't understand why we're not doing that'. <https://www.politico.com/news/2020/04/03/fauci-endorses-national-stay-at-home-order-162794>.
- Prescott, E. C. (1975). Efficiency of the Natural Rate. *The Journal of Political Economy*, 83(6):1229–1236.
- Press Briefing (2020). Remarks by President Trump, Vice President Pence, and Members of the Coronavirus Task Force in Press Briefing. <https://www.whitehouse.gov/briefings-statements/remarks-president-trump-vice-president-pence-members-coronavirus-task-force-press-briefing-16/>.
- Raifman, J., Nocka, K., Jones, D., Bor, J., Lisbon, S. K., Jay, J., and Chan, P. (2020). Non-essential Closure Dates. <https://tinyurl.com/statepolicysources>.
- Rojas, F. L., Jiang, X., Montenovo, L., Simon, K. I., Weinberg, B. A., and Wing, C. (2020). Is the Cure Worse than the Problem Itself? Immediate Labor Market Effects of COVID-19 Case Rates and School Closures in the US. *NBER Working Paper*.
- Salop, S. and Stiglitz, J. (1977). Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion. *Review of Economic Studies*.
- Sheremirov, V. (2020). Price dispersion and inflation: New facts and theoretical implications. *Journal of Monetary Economics*.
- Sheshinski, E. and Weiss, Y. (1977). Inflation and Costs of Price Adjustment. *Review of Economic Studies*, 44(2):287–303.
- Shilony, Y. (1977). Mixed pricing in oligopoly. *Journal of Economic Theory*.
- USA Today (2020). Partisan divide over economy grows, but Americans more worried about their health: Exclusive poll. <https://www.usatoday.com/story/news/nation/2020/05/29/coronavirus-poll-divided-over-economy-americans-worry-health/5264864002/>.
- Varian, H. R. (1980). A Model of Sales. *The American Economic Review*, 70(4):651–659.
- Zakary, O., Rachik, M., and Elmouki, I. (2017). On the analysis of a multi-regions discrete SIR epidemic model: an optimal control approach. *International Journal of Dynamics and Control*, 5(3):917–930.