

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Essays on Household Consumption and Relative Prices

Permalink

<https://escholarship.org/uc/item/06t8q0hn>

Author

Orchard, Jacob D

Publication Date

2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Essays on Household Consumption and Relative Prices

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Economics

by

Jacob Douglas Orchard

Committee in charge:

Professor Valerie Ramey, Chair
Professor Joseph Engelberg
Professor Munseob Lee
Professor Marc Muendler
Professor Johannes Wieland

2022

Copyright

Jacob Douglas Orchard, 2022

All rights reserved.

The Dissertation of Jacob Douglas Orchard is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2022

DEDICATION

To my wonderful parents, Tanja and Doug. They intoxicated me with a thirst for knowledge that has never been quenched.

EPIGRAPH

“Today you can buy an iPad 2 that costs the same as an iPad 1 that is twice as powerful, You have to look at the prices of all things.”

Federal Reserve Bank of New York President William Dudley (March 2011)

“I can’t eat an iPad!”

Working Class Audience Member

TABLE OF CONTENTS

Dissertation Approval Page	iii
Dedication	iv
Epigraph	v
Table of Contents	vi
List of Figures	ix
List of Tables	xii
Acknowledgements	xiv
Vita	xv
Abstract of the Dissertation	xvi
Chapter 1 Cyclical Demand Shifts and Cost of Living Inequality	1
1.1 Introduction	1
1.2 Data	7
1.3 Two Facts	12
1.3.1 Fact 1: Relative Spending on Necessities is Counter-Cyclical	12
1.3.2 Fact 2: Counter-cyclical Necessity Prices	16
1.4 A Static Model of Relative Supply and Demand	24
1.4.1 Firms	24
1.4.2 Households and Intra-temporal Substitution	25
1.5 Empirical Strategy	27
1.5.1 Results	29
1.5.2 Robustness	30
1.6 New Keynesian Model with Non-homothetic consumption preferences	33
1.6.1 Households	33
1.6.2 Firms	38
1.6.3 Equilibrium	39
1.6.4 Calibration	39
1.6.5 Results	40
1.7 Conclusion	44
1.8 Chapter Acknowledgments	47
Chapter 2 Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates	48
2.1 Introduction	48
2.2 The U.S. Macroeconomy in 2008	54

2.2.1	Narrative of 2008	54
2.2.2	Alternative Measures of Consumption Expenditures	59
2.2.3	Forecasts of Consumption in 2008	61
2.3	Macro Counterfactuals	65
2.3.1	Two-Good, Two-Agent New Keynesian Model with Hand-to-Mouth Consumers and Durable Goods	65
2.3.2	Calibration	70
2.3.3	Macro Counterfactuals	72
2.4	The Micro MPC Estimates	78
2.4.1	Estimation Strategy	78
2.4.2	Results	81
2.5	Conclusion	88
2.6	Chapter Acknowledgments	88
Chapter 3	Household Inflation and Aggregate Inflation	90
3.1	Introduction	90
3.2	Measuring Household-Level Inflation	94
3.2.1	Nielsen Homescan	95
3.2.2	Consumer Expenditure Survey	100
3.3	How Consumers React to Higher Household Inflation	101
3.3.1	Simple Model	103
3.3.2	Household consumption response in the data	105
3.4	Inflation Dispersion and Inflation	110
3.5	Menu-cost model	115
3.5.1	Households	116
3.5.2	Firms	117
3.5.3	Simulation	118
3.6	Conclusion	119
3.7	Chapter Acknowledgements	121
	Bibliography	122
Appendix: Chapter 1	131
A1	Mathematical Appendix	131
A1.1	Proof of Proposition 1	131
A1.2	Derivation of Wage-Phillips Curve	134
A1.3	A Note on Aggregation	137
A2	Data Appendix	139
A2.1	Cross-walk between CPI and CEX	139
A3	Alternate Calibrations	140
A3.1	Demand Parameter Estimation	141
A3.2	Results	142
A4	Additional Figures and Tables	142

Appendix: Chapter 2	154
B1 Model	154
B1.1 Households	154
B1.2 Wages	155
B1.3 Production of capital goods	157
B1.4 Production of final goods	158
B1.5 Prices	159
B1.6 Government	160
B1.7 Durable Goods Production	161
B1.8 Market Clearing	162
B1.9 Functional Forms	162
B2 Nominal Counterfactuals	162
B3 Data Appendix	162
B3.1 Details for Figure 2.1	162
B4 Decomposing the Difference Between OLS and DID Imputation	164
B5 Additional Regression Table	169
B6 DID Decomposition for Non-durable and New Vehicle Expenditure	169
Appendix: Chapter 3	172
C1 Household Model	172
C1.1 Permanent Unanticipated Shock	173
C1.2 Permanent Anticipated Shock	175
C1.3 Transitory Shock	176
C1.4 Stochastic Shock Process	177
C2 Asymmetric Responses to Household Inflation Shocks	178
C2.1 Price Index	179
C2.2 Nominal Consumption	181
C2.3 Real Consumption	183
C3 Response of Consumption to Inflation Shock by Income Group	184
C4 Simple model of household inflation rates	184
C5 Inflation Rates and Volatility by Income Group	190
C6 Effective versus Common Prices	191
C6.1 Robustness check using effective prices	193
C6.2 Distribution of Relative Price Indexes over time	194
C6.3 Household inflation rates using store-specific prices	196

LIST OF FIGURES

Figure 1.1.	Expenditure Ratio Based on Engel Curve	10
Figure 1.2.	Aggregate Expenditure on Necessities and Luxuries	14
Figure 1.3.	Time Series of Relative Necessity Prices	19
Figure 1.4.	Relative Necessity Prices and Unemployment	20
Figure 1.5.	Relative supply and relative demand	26
Figure 1.6.	IRFs: Response to Monetary Policy Shock	31
Figure 1.7.	IRF Robustness Checks: Response to Monetary Policy Shock	32
Figure 1.8.	Model and Data: Necessity Shares by Income Group	41
Figure 1.9.	Model v. Data: Historical Simulation	46
Figure 2.1.	Expenditures on New Motor Vehicles: Actual vs. Counterfactual	50
Figure 2.2.	2008 Tax Rebates	55
Figure 2.3.	Aggregate Disposable Income and Consumption	56
Figure 2.4.	Real Consumption Expenditures by Type of Product	57
Figure 2.5.	Log Price Indices for Consumption	58
Figure 2.6.	Federal Funds Rate	59
Figure 2.7.	Comparison of PCE to Retail and CEX	60
Figure 2.8.	New Motor Vehicle Sales to Consumers	61
Figure 2.9.	Contemporary Real Consumption Forecasts	63
Figure 2.10.	Forecasts from Our Four Models	64
Figure 2.11.	Counterfactual Real Consumption Expenditures: Baseline Model	73
Figure 2.12.	Counterfactual Real Consumption Expenditures: Less Elastic Durable Supply Model	76
Figure 2.13.	Decomposing the OLS and DID Imputation Coefficients	85

Figure 3.1.	Nielsen Individual Inflation v. CPI-U	99
Figure 3.2.	Distribution of Individual Inflation Rates: CEX Laspeyres Index	101
Figure 3.3.	Path of Household Retail Price-Index Following Household Inflation Shock	107
Figure 3.4.	Response of Nominal Household Retail Consumption to one-unit Household Inflation Shock	108
Figure 3.5.	Response of Real Household Retail Consumption to one-unit Household Inflation Shock	109
Figure 3.6.	Inflation and SD of City Inflation Rates	115
Figure A1.	Model and Data: Necessity Shares by Income Group	143
Figure A2.	Model v. Data: Necessity Share	143
Figure A3.	Model v. Data: Relative Necessity Prices	144
Figure A4.	Necessity Share: Non-durables	145
Figure A5.	Necessity and Luxury Expenditures Normalized by Sector Level Prices ..	146
Figure A6.	Income Level CPI's during the Great Recession and Slow Recovery	147
Figure A7.	Relative Necessity Prices including energy	148
Figure A8.	Relative Necessity Prices 1990-2000 Consumption Patterns	149
Figure A9.	Relative Necessity Prices 2000-2010 Consumption Patterns	150
Figure A10.	Relative Necessity Prices 2010-2020 Consumption Patterns	151
Figure A11.	Great Recession: Change in Necessity Share by Income Quintile	152
Figure A12.	Additional IRF Robustness Checks	153
Figure B1.	Counterfactual Nominal Consumption Expenditures: Baseline Model	163
Figure B2.	Counterfactual Nominal Consumption Expenditures: Less Elastic Durable Supply Model	164
Figure B3.	Expenditures on New Motor Vehicles: Alternative Counterfactuals	165
Figure B4.	OLS: Weights and CATT	168

Figure B5.	Non-Durable: Decomposing the OLS and DID Imputation Coefficients ..	169
Figure B6.	New Vehicles: Decomposing the OLS and DID Imputation Coefficients ..	170
Figure C1.	Response to Price Index Shock	178
Figure C2.	Path of Household Retail Price-Index Following Positive Household Inflation Shock	179
Figure C3.	Path of Household Retail Price-Index Following Negative Household Inflation Shock	180
Figure C4.	Response of Nominal Household Retail Consumption to a Positive one-unit Household Inflation Shock	181
Figure C5.	Response of Nominal Household Retail Consumption to a Negative one-unit Household Inflation Shock	182
Figure C6.	Response of Real Household Retail Consumption to a Positive one-unit Household Inflation Shock	183
Figure C7.	Response of Real Household Retail Consumption to a Negative one-unit Household Inflation Shock	184
Figure C8.	Preference Intensity by Income Group	189
Figure C9.	Distribution of Relative Price Indexes: 2017Q1	195
Figure C10.	Distribution of changes in RP_{ht} over time	196

LIST OF TABLES

Table 1.1.	Top luxury and necessity products	11
Table 1.2.	Descriptive statistics for luxuries and necessities	12
Table 1.3.	Relationship Unemployment and Relative Necessity Shares	21
Table 1.4.	Relationship Unemployment and Relative Necessity Prices	23
Table 1.5.	Baseline Calibration	40
Table 1.6.	Welfare Difference Low v. High Income Households	43
Table 2.1.	Baseline Calibration of the Model	70
Table 2.2.	General Equilibrium Marginal Propensity to Consume: Baseline Model . . .	74
Table 2.3.	General Equilibrium Marginal Propensity to Consume: Model with less elastic Durable Supply	75
Table 2.4.	General Equilibrium Marginal Propensity to Consume: Model without Durable Goods	78
Table 2.5.	Contemporaneous Household Expenditure Response to Rebate	82
Table 2.6.	First Stage: Rebate Amount Conditional on Rebate Receipt	83
Table 2.7.	Contemporaneous Household New Vehicle Expenditure Response to Rebate	86
Table 2.8.	Contemporaneous Household Non-Durable Expenditure Response to Rebate	87
Table 3.1.	Distribution of Spending	95
Table 3.2.	Summary Statistics	102
Table 3.3.	Response of Household Spending to Household Inflation Shock: CEX	110
Table 3.4.	Household Inflation Dispersion and Aggregate Inflation	112
Table 3.5.	Regional Household Inflation Dispersion and Regional Inflation	113
Table 3.6.	Robustness Check with Unemployment Rate	113
Table 3.7.	Module level Household Inflation Dispersion and Module Inflation	114
Table 3.8.	Model Simulation Results	119

Table A.1.	Almost Ideal Demand System Parameter Estimation	142
Table A.2.	Relationship Unemployment and Relative Real Expenditure	144
Table B1.	Expenditure Response to Rebate: Full Sample (monthly frequency)	171
Table C2.	Response of Spending to Household Inflation Shock by Income Group: CEX	185
Table C3.	High Preference Intensity \implies Volatile Inflation	189
Table C4.	Household Inflation Volatility and Income	190
Table C5.	Average Inflation Rates of Income Group Compared to National Rate	191
Table C6.	Household Inflation Dispersion and Aggregate Inflation with Effective Prices	193
Table C7.	Household relative price distribution AR(1)	195
Table C8.	Household Inflation Rates with Store Specific Prices v. Other Inflation Measures	199

ACKNOWLEDGEMENTS

I would like to acknowledge Professor Valerie Ramey for her immense support as my adviser, Johannes Wieland for his guidance especially with the quantitative models in my dissertation, Munseob Lee for introducing me to heterogeneous household price indices, and Marc Muendler and Joey Engelberg for their guidance during the dissertation writing and job market.

Chapter 1, in full, is currently being prepared for submission for publication of the material. The dissertation author was the sole investigator and author of this paper.

Chapter 2, in full, is currently being prepared for submission for publication of the material. This project was co-authored with Valerie Ramey and Johannes Wieland. The dissertation author was a primary investigator and author of this paper.

Chapter 3 is calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

VITA

- 2016 Bachelor of Sciences, Brigham Young University
- 2016-2022 Teaching Assistant, Department of Economics
University of California San Diego
- 2018-2019 Research Assistant, Dr. Munseob Lee, University of California San Diego
- 2019-2020 Research Assistant, Dr. Valerie Ramey & Dr. Johannes Wieland, University of
California San Diego
- 2022 Doctor of Philosophy, University of California San Diego

PUBLICATIONS

“County-level racial prejudice and the black-white gap in infant health outcomes” (with Joseph Price), *Social Science & Medicine* 181:191-198, 2017.

FIELDS OF STUDY

Major Field: Economics

ABSTRACT OF THE DISSERTATION

Essays on Household Consumption and Relative Prices

by

Jacob Douglas Orchard

Doctor of Philosophy in Economics

University of California San Diego, 2022

Professor Valerie Ramey, Chair

Chapter one examines the cyclical behavior of low-income versus high-income household price indices and documents two new facts: (1) during recessions prices rise more for products purchased relatively more by low-income households (necessities); (2) the aggregate share of spending devoted to necessities is counter-cyclical. I present a mechanism where adverse macroeconomic shocks cause households to shift expenditure away from luxuries toward necessities, which leads to higher relative prices for necessities. I embed this mechanism into a quantitative model which explains around half of the cyclical variation in necessity prices and shares. The results suggest that low-income households are hit twice by recessions: once by the recession itself and again as their price index increases relative to other households.

Chapter two, presents evidence that the high estimated MPCs from the leading household studies result in implausible macroeconomic counterfactuals. Using the 2008 tax rebate as a case study, we calibrate a standard medium-scale New Keynesian model with the estimated micro MPCs to construct counterfactual macroeconomic consumption paths in the absence of a rebate. The counterfactual paths imply that consumption expenditures would have plummeted in spring and summer 2008 and then recovered when Lehman Brothers failed in September 2008. We use narratives and forecasts to argue that these paths are implausible. We go on to show that reasonable modifications of the model result in general equilibrium forces that dampen rather than amplify micro MPCs. We also show that estimators of the average treatment effect yield smaller micro MPC estimates than the standard two-way fixed effects estimators. The combination of smaller micro MPCs and dampening general equilibrium forces implies general equilibrium consumption multipliers that are below 0.2.

In **chapter three**, I construct novel measures of household-level inflation and show that an increase in a household's personal inflation rate leads to a persistent increase in their price index. Households respond to a personal inflation shock by decreasing nominal consumption, which means their real consumption falls more than one-for-one. I also find a statistically robust relationship between inflation dispersion (the variance of household inflation rates) and the level of absolute aggregate inflation.

Chapter 1

Cyclical Demand Shifts and Cost of Living Inequality

1.1 Introduction

Since the 2008 financial crisis, a flurry of research has shown that recessionary shocks have heterogeneous effects on households and can exacerbate inequality.¹ Much of the past literature has focused on the cyclical behavior of nominal consumption and income inequality and has overlooked cost-of-living differences across households, which is the denominator of real inequality. This paper shows that failing to include differential changes in the cost-of-living can dramatically understate the true distributional consequences of recessions.

This study asserts that higher consumer-price inflation for low-income households is a feature of recessions. I present a novel mechanism, “Cyclical Demand Shifts,” where contractionary shocks lead households to cut back on luxuries (e.g., vacations and pet services), but households continue to buy necessities (e.g., groceries). This shift in relative demand increases the relative price of necessities, which disproportionately affects poorer households since a larger share of their consumption basket is devoted to necessities. The mechanism implies that poor households are hit twice by recessions: once by the recession itself and again when the price of their basket increases relative to other households.

¹See Heathcote et al. (2020), Feiveson et al. (2020), Krueger et al. (2016), Meyer and Sullivan (2013), Hoynes et al. (2012)

This paper makes three main contributions. First, I show empirically that while consumption falls during recessions, it does not fall equally for all products; specifically, consumption falls more for luxury products than necessities. Second, I show that the relative price of necessities is counter-cyclical. Third, I present a theoretical framework that incorporates the “Cyclical Demand Shift” mechanism into a standard business cycle model. This model can explain a significant percentage of the cyclical behavior of relative necessity prices and consumption and estimates sizable increases in the relative cost-of-living for low-income households during recessions. Krueger et al. (2016) find that during the Great Recession, nominal consumption growth fell by 0.3 percent more for households in the lowest wealth quintile compared to the highest. A back-of-the-envelope calculation incorporating this paper’s cost-of-living inequality estimates suggests that the actual difference in the fall of *real* consumption is almost four times as high at 1.15 percent.²

In order to study differences in household-level price indices across time, I create 118 product sectors in the Consumer Expenditure Survey (CEX) that represent the same type of spending from 1991 to 2020. I then sort households into five different income quintiles. Next, I construct a measure of the relative importance of a product in a low-income household’s consumption basket by dividing the pooled average of the product’s nominal expenditure share for households in the first income quintile by the average expenditure share for that product of households in the highest income quintile (expenditure ratio).³ I define necessities as products purchased more by low-income households (expenditure ratio greater than one) and luxuries as products purchased more by high-income households. Next, I match these 118 product sectors with publicly available price data from the CPI.

Based on this categorization, I investigate how aggregate consumption shifts between luxuries and necessities over the business cycle. The aggregate expenditure share devoted to necessities increased during all three of the recessions in my sample (2001, the Great Recession,

²Krueger et al. (2016) classify households based on wealth levels, where this paper sorts households based on income.

³An expenditure ratio greater than one implies that the product’s Engel curve is downward sloping.

and the Covid-19 Recession). In fact, during both the 2001 recession and the Great Recession, all of the fall in real PCE expenditures can be accounted for by large declines in luxury expenditure, while nominal expenditures on necessities remain roughly constant at pre-recession levels. I formally test the relationship between aggregate spending on necessities and luxuries and economic slack in a panel regression using all 118 product sectors. I find that a one percent increase in the unemployment rate is associated with a 0.9-2 percent increase in the aggregate share of spending on necessities. This relationship continues to hold even when controlling for whether products are durables, services, or in the energy/transportation sector.⁴

Next, I examine the cyclical behavior of prices for necessity products. Because I have price data for a subset of products from 1967-2021, I can observe the cyclical behavior of necessity and luxury prices over seven different recessionary periods. I construct composite price indices for necessities and luxuries. I find that the price index for necessities relative to luxuries has increased during five out of the last seven recessions.⁵ Separately, in a panel regression using all 118 products, I find that a one percent increase in the unemployment rate is associated with a 0.7-1.5 percent increase in the relative price of necessity products. This relationship is also robust to including controls for whether products are services, durables, or energy.

Having documented that both necessity relative prices and aggregate shares increase during recessions, I formally introduce a static model that can rationalize these facts. The critical components of this model are non-homothetic preferences at the aggregate level and a concave production possibilities frontier. The non-homothetic preferences lead to cyclical demand shifts between necessities and luxuries that track the evolution of aggregate consumption expenditure. The concave production possibilities frontier leads to higher relative costs for the expanding sector. These components are sufficient for an aggregate decrease in expenditure to lead to a relative expansion in the necessity sector and higher relative necessity prices.

⁴This relationship is not simply mechanically related to higher necessity prices, as a necessity product's relative real expenditure (nominal aggregate expenditure divided by the product-specific price index) is also positively related to unemployment.

⁵Six out of the last seven when including the volatile energy/transportation sector.

Is aggregate demand non-homothetic? While the cross-sectional data show that low-income and high-income households buy different bundles, this does not necessarily imply that aggregate preferences are non-homothetic; i.e. in response to an exogenous shock that changes aggregate consumption, does the aggregate consumption bundle change?⁶ I test this assumption along with the model's primary conclusion, an increase in necessity prices following a decrease in aggregate expenditure, using Monetary Policy news shocks (Gürkaynak et al. 2004). Since the mechanism operates through changes in expenditure, I first show that 24 months after a 25 basis point contractionary monetary policy shock, aggregate expenditure falls by approximately 2 percent. Next, I show that the same contractionary shock leads the aggregate share of spending devoted to necessity products to increase by 5 percent and relative necessity prices increase by around 2.5 percent. Results are similar when conditioning on whether the product is a durable good or a service, sectors that typically have high-interest rate elasticities or sticky prices. These results show that an exogenous shock that lowers aggregate expenditure also leads to higher relative necessity prices and consumption.

Next, I present a quantitative New Keynesian model that incorporates non-homothetic preferences and can be calibrated to the US economy. Household preferences are represented by the Almost Ideal Demand System (AIDS) (Deaton and Muellbauer 1980). While these preferences have been used in the trade literature, to my knowledge this is the first paper to incorporate these preferences into a New Keynesian style model. The AIDS inherits well-behaved aggregation properties from the Generalized Linear class of demand systems (Muellbauer 1975), which allows me to solve for aggregate necessity shares and relative necessity prices using a representative agent framework. I calibrate the model to match the United State's aggregate expenditure and necessity share in 2005-2006, right before the Great Recession.

The quantitative model can explain a significant fraction of the cyclical variation in

⁶This question is also related to the relationship between income and expenditure elasticities. I define products as necessities/luxuries based on income elasticity and then test the aggregate expenditure elasticity of these products. The relationship between household income and aggregate expenditure elasticities is partially responsible for cyclical price index disparities across income groups.

relative necessity prices and shares. In a validation exercise, I introduce a series of shocks to the model so that expenditure in the model exactly matches the cyclical component of Personal Consumption Expenditures (PCE) from 1994-2021, which results in a model-produced time-series of necessity prices and shares. The model-produced time-series are highly correlated and of the same scale as their data counterparts: the model's necessity price series has a 44 percent correlation with cyclical necessity prices in the data, and the necessity share series has a 55 percent correlation.

With the model in hand, I examine the welfare consequences of the Great Recession when households have different price indices.⁷ Using the non-homothetic price index implied by the AIDS, I estimate that the price index for low-income households increased by 0.85 percentage points relative to the price index of high-income households during the Great Recession (2007Q3-2009Q2). This large relative increase in cost-of-living can have considerable welfare consequences. I perform a test of the expenditure equivalent welfare loss due to the Great Recession, and I find that the Great Recession was 22 percent more costly for households in the bottom income-quintile compared to households in the top quintile.

Taken together, the results suggest that the difference in cost-of-living between low- and high-income households varies systematically over the business cycle: increasing during recessions and subsiding during expansions. This cost-of-living channel is yet another reason why recessions are particularly costly for low-income households.

This paper is most closely related to a small but fast-growing literature examining changes in the cost-of-living across household groups. Early research by Amble and Stewart (1994), Garner et al. (1996), Hobijn and Lagakos (2005), and McGranahan and Paulson (2005) found only limited differences in inflation rates across demographic groups.⁸ However, more recent work has leveraged detailed product categories as well as barcode level data to document

⁷Since the model abstracts from differences in employment loss or ability to borrow during the recession, these results are due only to differences in relative prices

⁸An exception in this early-period is work by Crawford and Oldfield (2002) who found that few households in Britain have inflation close to the official Retail Price Index

substantial differences in inflation-rates across households (Kaplan and Schulhofer-Wohl 2017, Jaravel 2019, Cavallo 2020, Gürer and Weichenrieder 2020, Argente and Lee 2021, Orchard 2021, Lauper and Mangiante 2021) This literature has focused on either trends in inflation rate disparities (Jaravel 2019, Gürer and Weichenrieder 2020) or particular events such as the Great Recession (Argente and Lee 2021), the 1994 Mexican Devaluation (Cravino and Levchenko 2017), and the Covid-19 Pandemic (Cavallo 2020, Jaravel and O’Connell 2020). In contrast, this paper shows empirically and theoretically that inflation inequality increases following any shock that affects aggregate consumption expenditure.⁹

This paper also contributes to the literature on endogenous demand shifts. For example, Jaimovich et al. (2019) show that households switched from high- to low-quality products during the great recession and this shift in demand led to lower labor demand since low-quality products use less labor in production. Over a longer horizon, Boppart (2014) and Comin et al. (2021), show that non-homothetic demand can explain the shift from agriculture to manufacturing and services in advanced economies. Comin et al. (2020) shows how long-term shifts can contribute to labor-market polarization. Work by Bils and Klenow (1998) uses product expenditure elasticities to test competing business cycle models. This paper shows that over the short term, shifts in demand can lead to higher prices in the expanding sector, which can have heterogeneous effects on income-level cost of living.

The remainder of the paper proceeds as follows: Section 2 details the data I use in the analysis, Section 3 presents the twin motivating facts (counter-cyclical necessity prices and aggregate shares), Section 4 formally presents the cyclical demand shift mechanism, Section 5 tests the conclusions of the mechanism empirically via monetary policy news shocks, Section 6 presents the quantitative model, and Section 7 concludes.

⁹Inflation inequality may be a confusing term since price inflation traditionally has been defined as a general increase in the prices of goods and services in an economy or a decrease in the purchasing power of a particular currency. In the emerging literature on changes in the cost-of-living across income groups, “Inflation Inequality” is generally defined as differences in the change of the cost of achieving a particular level of utility across household groups (Jaravel 2021).

1.2 Data

This project's primary data sources are the Consumer Expenditure Survey (CEX) and publicly available product-level Consumer Price Index (CPI) series, both from the Bureau of Labor Statistics (BLS). The BLS uses the CEX and micro-level price data to construct the CPI-U. In doing so, they aggregate micro-price data into 243 different item strata and construct weights using the CEX (U.S. BLS, 2020). However, price time series for the 243 item strata are not publicly available. Instead, the BLS publishes CPI price series for a variety of more aggregated products, which I use in the analysis.

I create a cross-walk by hand between the publicly available item-level CPI categories and CEX MTBI micro-data. In this cross-walk, I create CEX products from base level UCC codes¹⁰ that were consistent across the 1991-2020 survey waves.¹¹ While some categories do not exist in earlier years (e.g., internet expenditures were not recorded prior to 1995 in the CEX), the categories are created so that comparison between years is possible and represent the same breadth of spending in each year. Next, I match these CEX categories to CPI item-level price data. Where this was not possible (for example, CPI has separate price series for premium and regular gasoline), I created broader CEX products to match with the CPI or use a broader CPI category (e.g., gasoline). The result is 120 distinct products that represent the same types of spending from 1991-2020 (118 excluding rent and owners equivalent rent). Taken together, these product categories represent approximately 97.5 percent of all consumption spending in the CEX.¹²

The CPI price series for these categories is not available across the entire sample period, as there was an expansion in published categories in 1967, 1977, 1987, and 1997. For this analysis, I use either a balanced sample of products with continuous price information over

¹⁰A UCC code is the most disaggregated expenditure category in the CEX.

¹¹While the CEX survey was fielded in earlier years, the more detailed MTBI files are only available starting with the 1990 survey. Most product categories in this analysis start in the 1991 Quarter 2 survey.

¹²Further details on this cross-walk are in section A2.1 of the appendix.

some period (for example, 1987-2019) or an unbalanced sample. Results are similar using either method.

I exclude rent and owners-equivalent-rent since most high-income households are homeowners while low-income households generally rent their homes. While the BLS constructs an imputed owners' equivalent rent series, homeowners do not actually pay this price. When rent prices change, homeowners can still consume at their initial endowment point and are shielded from increases in home prices. While studying the impacts of owning versus renting on real income and wealth inequality is an interesting area of research, it is not the focus of this article.

I divide households into five different income groups following Aguiar and Bils (2015). Namely, I keep only households that participate in all four CE interviews and are complete income reporters. I also include only urban households and households whose household head is between 25 and 64. This leaves me with 76,448 distinct households from 1991-2019.

I divide households into five different income groups based on their pre-tax income. In addition to pre-tax income reported in the CEX, I add in income from alimony, gifts, gambling winnings, inheritance, and any other payments from persons outside the household; similarly, I subtract from income the alimony, child support, etc. paid by the household. Next, I regress this income measure on dummies of the household size, age, and the number of income earners in the household. Then, I group households into groups based on their income percentile in the quarter they report their income (their fourth CEX interview). Similar to Aguiar and Bils (2015), the top income group are households in the 80-95 percentile of income (this lessens the degree to which changes in top-coding and outliers can change the composition of the top group). The bottom income group is households in the 5-20 percentile of income. Groups 2, 3, and 4 are households in the 20-40 percentile, 40-60 percentile, and 60-80 percentile, respectfully.

Households are interviewed four times three months apart and are asked about their spending in each of the previous three months in small categories (UCCs). These interview times do not necessarily correspond to calendar quarters. For example, a household interviewed in May would be asked about their April, March, and February spending. In principle, I should be

able to use the CEX data to create monthly expenditure variables for each household or quarterly expenditure based on each household's reported expenditure in that quarter. However, there is widespread expenditure smoothing across months within an interview (Coibion, Gorodnichenko, Kueng, Silva 2017). This means that reported expenditure in UCC u for a household interviewed in May would be relatively smooth from February to April, but would have a much larger change when compared to January spending (which would come from the previous survey). For this reason, I base household spending at time t based on the quarter or month they were interviewed rather than the quarter or month for which they report their spending (Coibion et al. 2017). In the main analysis, the measure of aggregate spending share in a category j in month t is smoothed across the three preceding months to capture all households in the interview wave.

I create quarterly expenditure shares for the 118 product groups for each household by dividing expenditure in category j by total consumption expenditure. Total consumption expenditure is defined as quarterly household expenditure minus savings in pension plans, life insurance, health insurance rebates, and cash contributions to those outside the household.

I create income group expenditure shares as the weighted average of household expenditure shares for all households in the income group. I use the household survey weights computed by the BLS. Note that this is different from how the BLS creates expenditure shares for the CPI, since they also base their shares on the contribution of the household to total spending, which puts more weight on higher spending households. Since this paper is focused on non-homotheticities in consumption shares, weighting based on expenditure is problematic since it would give more weight to households at the upper end of an income group (say those nearer to the 20th percentile vs. those nearer the 5th percentile). This could also be a problem when some households report more of their expenditure than others (see Aguiar and Bils (2015) for under-reporting in the CEX).

I pool the quarterly expenditure shares across quarters to create a single expenditure share for each income group and product. I define R_j , as the ratio of the share of consumer spending in the lowest income quintile to the share of spending in the highest quintile:

$$R_j = \frac{\sum_t \frac{1}{N_{t,Q1}} \sum_{h \in Q1} S_{jth}}{\sum_t \frac{1}{N_{t,Q5}} \sum_{h \in Q5} S_{jth}}. \quad (1.2.1)$$

R_j is equal to one if, on average, poor and rich households spend the same percentage of their expenditure on product j . I define products as necessity goods if poor households have a higher expenditure share on these goods relative to rich households ($R_j > 1$) and luxury goods as products with $R_j < 1$.




Figure 1.1. Expenditure Ratio Based on Engel Curve

Note: Panel (a) shows a product j with a downward sloping Engel curve (Necessity). Panel (b) shows a luxury product. Panel (c) shows a product with a hump shaped Engel curve. In this example, it is a necessity since the average expenditure share for j is higher for the lowest income group Q_1 than the highest Q_5 .

Figure 1.1 shows how this approach is similar to comparing the level of the share based Engel curve at the top and bottom of the income distribution. If the Engel curve is linear, then the “necessity” rank of the good using this method would be the same as the rank derived from the slope of the Engel curve (where a slope of zero would correspond to an expenditure share ratio of 1). If the underlying Engel curve is non-linear (as suggested by Atkin, Faber, Fally, and Gonzalez-Navarro (2020)), then this method ranks goods by their importance in the consumption basket of low-income versus high-income households.

Table 1 Panel A shows the top 10 luxury goods. The consumption category that has the highest comparative expenditure by those in the top income group is “Club memberships for shopping clubs, fraternal, or other organizations”, which has an expenditure ratio, R_j , of .31.

Table 1.1. Top luxury and necessity products

Panel A: Top Luxury Goods		
CPI Category	Expenditure Ratio	Percent Agg. Spending
Club memberships for shopping clubs, fraternal, or other organizations	0.31	0.34
Other Lodging away from home including hotels, and motels	0.33	0.80
Pet services	0.33	0.09
Day care and preschool	0.34	0.75
Fees for lessons or instruction	0.36	0.59
Other intercity transportation	0.36	0.2
Airline Fares	0.37	0.82
Alcohol Away from Home	0.40	0.44
Other Furniture	0.40	0.19
Elementary and high school tuition and fees	0.40	0.38
Panel B: Top Necessity Goods		
CPI Category	Expenditure Ratio	Percent Agg. Spending
Cigarettes	3.28	0.84
Electricity	1.68	3.11
Tobacco products other than cigarettes	1.63	0.07
Food at Home	1.51	12.04
Intracity transportation	1.49	0.20
Water and sewerage maintenance	1.45	0.8
Prescription drugs	1.44	0.6
Used Cars and Trucks	1.41	4.4
Telephone services	1.40	2.9
Gasoline (all types)	1.38	4.71

Source: Consumer expenditure survey and author's own calculations.

Note: Expenditure ratio is defined as the average expenditure share of households in the bottom income group divided by the average expenditure share of households in the top income group. Percent Agg. Spending is computed on households in sample.

This means that on average, households in the highest income group spend 3.3 times as much of their budget on this category compared to households in the lowest income group. Other top luxury goods include Airline flights, Daycare, Hotels, Private Lessons, and alcoholic beverages away from home.

Panel B shows the top 10 necessity goods. These include tobacco products, food at home, electricity, and intracity transportation (e.g., bus or subway). Table 1.2 shows that luxuries tend

to be more concentrated in services and durable goods, while necessities are more concentrated in energy and transportation.

Table 1.2. Descriptive statistics for luxuries and necessities

Descriptive Stats	Necessity	Luxury
Number	31	87
Number Durables	3	33
Number Services	17	33
Number Energy	5	4
Average Percent Aggregate Expenditure	1.3%	0.4%
Percent Expenditure Durables	11%	31%
Percent Expenditure Services	44%	54%
Percent Expenditure Energy	22%	4%

Note: These 118 products exclude the two housing products: rent and owners equivalent rent. Energy: denotes that the product is part of the energy or transportation sectors.

1.3 Two Facts

In this section, I use the combined CEX-CPI data to examine the consumption and pricing behavior of luxuries and necessities. To this end, I begin by creating composite necessity and luxury products so that the reader can visualize the relationship between relative prices/shares and the business cycle. I also perform panel regressions and show a strong positive correlation between the unemployment rate and the relative aggregate share/price of necessities.

1.3.1 Fact 1: Relative Spending on Necessities is Counter-Cyclical

Visual Evidence

First, I show that aggregate spending on necessities rises relative to luxuries during recessions. Using aggregate expenditures in the CEX on each of the 118 categories, I construct

the aggregate necessity share as:

$$s_{N,t} = \frac{\sum_{j \in \text{Necessity}} x_{jt}}{X_t}. \quad (1.3.1)$$


where, x_j is the total aggregate expenditure in the CEX on necessity sector j and X is total non-housing expenditure in the CEX. Panel A of figure 1.2 shows how the aggregate necessity share changes over time. The necessity share increased during the early 1990s, fell during the dot-com boom and increased during the mild recession of 2001. Then there was a drastic increase in the necessity share starting in 2007, the beginning of the great recession, which peaked between 2013 and 2014, which is around the same time that real per-capita GDP recovered from its 2007 peak. The necessity share then falls during the expansion of the mid-2010s and then rises again during the Covid-19 recession. Figure A4, in the appendix, shows that these same patterns are still present when we restrict the sample to only non-durables.

Not only does the aggregate spending share of necessities rise in recessions, almost all of the fall in consumption spending during recessions can be attributed to *falls in luxury expenditure* rather than falls in necessity expenditure. Panel B of figure 1.2 shows imputed aggregate expenditure on luxuries and necessities by multiplying equation (1.3.1) by real personal consumption expenditures (PCE). The vast majority of the fall in consumption during the 2001 recession and the Great Recession can be attributed to a decline in luxury spending, while necessity expenditure either remains at the same level as before or even *increases!*¹³ This fact remains when deflating luxury and necessity expenditure by each sectors relative prices (see figure A5).

The increase in the aggregate necessity share during the Great Recession was precipitated

¹³The larger increase in necessity rather than luxury expenditures from 1991-2020 could seem at odds with the rise in aggregate income/spending over this period, as well as papers in the structural change literature such as (Comin et al. 2021), which document the change from Agriculture to Manufacturing and then to service expenditure. I should note two things about the patterns I find: (1) the long-term increase in necessity expenditure relative to luxury expenditure is moderated considerably when expenditure is deflated by sector level prices (see figure A5); (2) in this period of the U.S. Economic history there is a shift from manufacturing towards service expenditure (Schettkat and Yocarini 2006), both of which are more likely be classified as luxuries in my categorization.

Panel A: Necessity Share of Aggregate Expenditures



Panel B: Necessity and Luxury Imputed Expenditure

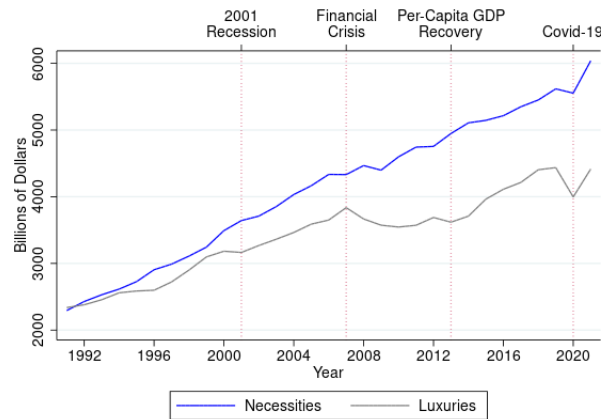


Figure 1.2. Aggregate Expenditure on Necessities and Luxuries

Source: Consumer Expenditure Survey and Author's own calculations. Excludes housing.

Note: Shaded lines indicate bootstrapped 90-percent confidence interval.

by all income groups. Figure A11 in the appendix shows the percentage point increase in the average necessity share for each income quintile during the Great Recession, 2007Q3-2009Q2, and the subsequent slow recovery (2009Q2-2012Q4). All income groups increased their share of necessity consumption expenditure by at least 2.5 percentage points during this period. The increase does vary by income group; for example, the lowest income quintile had the lowest increase in necessity share, especially during the official NBER recession, which may indicate a lack of an ability to substitute towards more luxuries (Argente and Lee 2021). It is important to note that while the shift in necessity expenditure varied by income group, the income group

ranking of necessity shares does not change. The lowest income group had the highest necessity share of expenditure during the Great Recession (around 72 percent), and the highest income quintile had the lowest (around 52 percent).

Regression Evidence

The visual evidence in the previous subsection shows that generally, relative necessity shares increase during recessions. Now I formally test the relationship between relative necessity shares and aggregate economic activity using a simple regression:

$$x_{j,t} = \beta_0 + \beta_1 Unemployment_t \times R_j + \beta_2 Unemployment_t \times Z_j + \delta_t + \gamma_j + \varepsilon_{j,t}. \quad (1.3.2)$$

Here, the dependent variable, $x_{j,t}$ is the log-relative price of products in sector j at time t or the log-aggregate share (presented in the next subsection). The dependent variable is regressed on the interaction of the unemployment rate with R_j the relative expenditure ratio, which is increasing for necessities. I also include time δ_t and sector γ_j fixed effects (which absorb the level effect in the interaction). Finally, Z_j is an indicator for whether the product sector is a service, durable, or in the energy/transportation sector.

The regression results have several advantages over the visual evidence. For example, I no longer have to rely on a binary definition of the necessity product since R_j is a continuous variable. Also, in the regression, I can control for a variety of confounding factors that may be correlated with a product's income elasticity and cyclical. For example, services have stickier prices than goods (Nakamura and Steinsson 2008) and high-income households also buy more services. Also, durable purchases are particularly sensitive to interest rates (McKay and Wieland 2021, Barsky et al. 2007) and could be another confounding factor.

Table 1.3 shows the correlation between the log-aggregate share of necessities and the unemployment rate. Panel A replaces R_j with a binary definition of necessity, while panel B shows the results of the regression in equation (1.3.2). Column 1 shows the baseline results.

In column 2, to determine that the results are not dependent on some arbitrary classification of spending into 118 categories, I weigh each observation by the sector's share in aggregate spending. In columns 3-5, I add in controls of the interaction of the unemployment rate with various aspects of the product j that may confound the results, including whether the product is directly related to oil prices (energy and transportation), whether the product is durable, or if the product is a service. Results are highly statistically significant and around the same size in all specifications. Overall, I find that a one percentage point increase in the unemployment rate is associated with an .9 – 2 percent increase in relative share of aggregate spending on necessities. This relationship is not simply the result of higher prices for necessities when unemployment is high (see next subsection); in the appendix Table A.2 I show a strong positive relationship between relative necessity sectoral-real expenditure and the unemployment rate.

1.3.2 Fact 2: Counter-cyclical Necessity Prices

Visual Evidence

Next, I show a visualization of the relative prices of necessities and luxuries over the business cycle. I create a geometric-price index for a representative necessity (luxury) good as:

$$P_t^K = \prod_j \left(\frac{p_{j,t}}{p_{j,b}} \right)^{\omega_j}, \quad (1.3.3)$$

where $K = \{N, L\}$ for necessity and luxury respectively, and ω_j is the pooled aggregate share of product j in total necessity or luxury spending from 1991-2020.¹⁴ Note that b refers to the prices in some base period, which I define as the first period in the sample. I then construct the relative necessity price as the ratio of the price of the composite necessity over the composite luxury:

¹⁴In the appendix, I show that my results are robust to pooling the aggregate share and income-share data over a smaller time period.

$$RP_t^N = \frac{P_t^N}{P_t^L}. \quad (1.3.4)$$

Figure 1.3 shows the results of this visualization. I have price data for some products since 1967, while for others, the publicly available series is much shorter. I construct multiple different versions of equation (1.3.4) corresponding to more inclusive balanced samples of products. For example, the series in blue comes from a balanced sample of 17 products with continuous price data from 1967-2020, while the series in red contains a much higher number of products (98) over a shorter period (1997-2020). For visualization purposes, I remove the volatile energy and transportation sectors from this graph (appendix figure A7 shows the results with energy and transportation).

Panel A shows the unfiltered series with NBER recession dates shown in gray, while panel B removes the trend component to produce a cyclical series following Hamilton (2018). There are some large differences between the balanced samples in the unfiltered series, but each of the filtered series closely track each other. Two patterns are apparent: (1) there is a large increase in the relative price of necessities during and around NBER recessions. For example, the relative price of necessities increased by more than 5 percent immediately following the Great Recession relative to trend. Relative necessity prices have increased in five of the last seven recessions. The second pattern (2) is that there is an increase over time in the relative price of necessities, which might be picking up the innovation in luxuries mechanism explained in Jaravel (2019). Both patterns are robust to varying the definition of necessities and luxuries. For example, suppose I define products as necessities or luxuries based on the consumption pattern of a particular decade (say 2010-2020) rather than pooling data from 1991-2020 together. In that case, the cyclical pattern of relative necessity prices and recessions holds, but the trend of increasing necessity prices does not (see figures A8, A9, and A10 in the appendix).

Figure 1.4 shows a bin-scatter plot comparing the level of slack in the economy (measured

by the unemployment rate) and the cyclical component of relative necessity prices. The left panel uses the long time series (1967-2021) with 17 balanced product sectors, while the right panel uses the shorter time series (1997-2021) with 98 balanced product sectors. In both cases there is a strong positive relationship between relative necessity prices and the unemployment rate. However, the relationship is tighter in the more recent period.





Figure 1.3. Time Series of Relative Necessity Prices

Source: Bureau of Labor Statistics and Author's own calculations. Excludes rent, owners equivalent rent, and energy. Data filtered following Hamilton (2018). Shaded area indicates NBER Recessions.

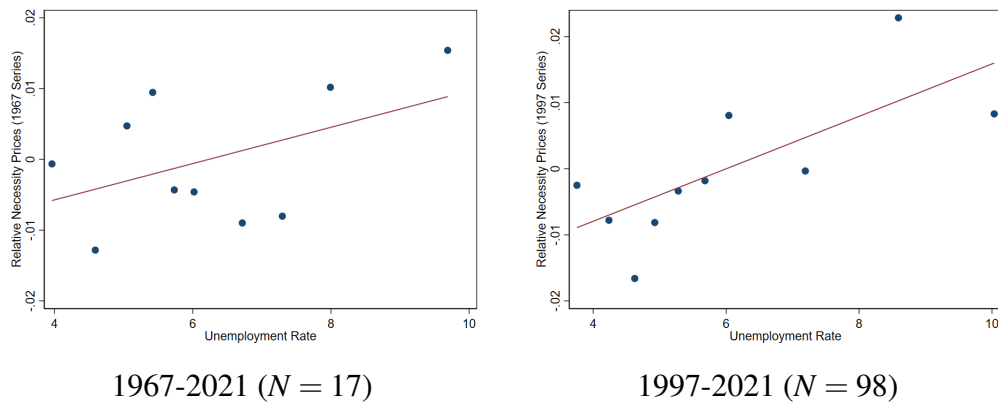


Figure 1.4. Relative Necessity Prices and Unemployment

Source: Bureau of Labor Statistics and Author's own calculations.

Note: In each plot, the y-axis is the (binned) residuals of the relative necessity price following the filtering method in Hamilton (2018), while the x-axis is the unemployment rate. The red line represents a bivariate regression line between these two variables. The left plot uses a balanced panel of products from 1967-2021, while the right panel uses the larger balanced panel under a shorter time horizon (1997-2021). Excludes rent, owners-equivalent rent, and energy.

Table 1.3. Relationship Unemployment and Relative Necessity Shares

Panel A: Binary necessity good					
	(1)	(2)	Log-Share (3)	(4)	(5)
Right hand side variables:					
UR × Necessity	0.019*** (0.006)	0.018*** (0.006)	0.014** (0.006)	0.009* (0.005)	0.020*** (0.007)
UR × Energy			0.023* (0.012)		
UR × Durable				-0.046*** (0.013)	
UR × Service					0.013 (0.011)
Sector FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
Weighted	No	Yes	Yes	Yes	Yes
Observations	42,700	42,700	42,700	42,700	42,700
Panel B: Scale by expenditure ratio					
	(1)	(2)	Log-Share (3)	(4)	(5)
Right hand side variables:					
UR × Exp. Ratio	0.026*** (0.006)	0.023*** (0.005)	0.019*** (0.004)	0.014*** (0.004)	0.027*** (0.006)
UR × Energy			0.022* (0.012)		
UR × Durable				-0.045*** (0.013)	
UR × Service					0.017 (0.011)
Sector FE	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes
Weighted	No	Yes	Yes	Yes	Yes
Observations	42,700	42,700	42,700	42,700	42,700

Notes: The unit of observation is the sector-month. Exp. ratio is the ratio of expenditure shares of poor over rich households for the sector. Standard errors, in parentheses, are clustered at the time level and are robust to auto-correlation. Significance at the 1, 5, and 10 percent levels indicated by ***, **, and *. Share is defined as the aggregate expenditure on sector j divided by total aggregate expenditure.

Regression Evidence

I repeat the regression exercise from the previous subsection, but I use the log-price of each sector as the dependent variable in equation (1.3.2). Results from these regressions are shown in table 1.4. Panel A replaces R_j with a binary definition of necessity, while panel B shows the results of the regression in equation (1.3.2). Column 1 shows the baseline results. In column 2, to determine that the results are not dependent on some arbitrary classification of spending into 118 categories, I weigh each observation by the sector's share in aggregate spending. Column 3 shows the results with a balanced sample. In columns 4-6, I add in controls of the interaction of the unemployment rate with various aspects of the product j that may confound the results, including whether the product is directly related to oil prices (energy and transportation), whether the product is durable, or if the product is a service. Results are highly statistically significant and around the same size in all specifications. Overall, I find that a one percentage point increase in the unemployment rate is associated with an 1.3 – 1.8 percent increase in relative prices for necessity products.

Table 1.4. Relationship Unemployment and Relative Necessity Prices

Panel A: Binary necessity good						
	Log-Relative Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Right hand side variables:						
UR × Necessity	0.015 (0.009)	0.015*** (0.004)	0.010** (0.004)	0.007** (0.003)	0.014*** (0.003)	0.013** (0.006)
UR × Energy				0.036*** (0.008)		
UR × Durable					-0.006 (0.017)	
UR × Service						-0.008 (0.013)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Weighted	No	Yes	Yes	Yes	Yes	Yes
Balanced Sample	No	No	Yes	No	No	No
Observations	49,963	49,963	24,480	49,963	49,963	49,963
Panel B: Scale by expenditure ratio						
	Log-Relative Price					
	(1)	(2)	(3)	(4)	(5)	(6)
Right hand side variables:						
UR × Exp. Ratio	0.018** (0.008)	0.020*** (0.004)	0.012*** (0.004)	0.013*** (0.003)	0.019*** (0.004)	0.018*** (0.006)
UR × Energy				0.035*** (0.008)		
UR × Durable					-0.005 (0.018)	
UR × Service						-0.006 (0.013)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Weighted	No	Yes	Yes	Yes	Yes	Yes
Balanced Sample	No	No	Yes	No	No	No
Observations	49,963	49,963	24,480	49,963	49,963	49,963

Notes: Standard errors, in parentheses, are clustered at the time level and are robust to auto-correlation. Significance at the 1, 5, and 10 percent levels indicated by ***, **, and *. The balanced sample are 59 sectors with continuous price data from 1987-2021.

To summarize, I find a statistical and economic significant correlation between relative necessity prices and shares with the unemployment rate. This result is not driven by differences in the service, energy, or durability composition of necessities. In the next section, I present a mechanism that can explain these two facts.

1.4 A Static Model of Relative Supply and Demand

In this section, I formalize the intuition behind the cyclical demand shift mechanism. I present a static model with a necessity and a luxury sector represented by perfectly competitive firms with concave production over labor. Households have non-homothetic preferences over these two sectors. This model is presented in partial equilibrium, and I abstract from the household labor market and savings decisions. Instead, the level of household expenditure, X , is exogenous. I show that a decline in the expenditure level, X , leads to higher equilibrium consumption shares and prices for the necessity sector.

1.4.1 Firms

There are two sectors $\{N, L\}$. Each sector is competitive and is represented by a firm with a homogeneous production function over labor:

$$Y_i = F(H_i). \quad (1.4.1)$$

I assume that $F(\cdot)$ is positive and homogeneous of degree $k \in (0, 1)$, implying that the firm has concave production over labor. Firms can hire labor at an exogenous fixed wage rate w . Profit maximization implies that the ratio of the wage and the sector price is equal to the marginal productivity of labor:

$$\frac{w}{p_i} = F_H(H_i). \quad (1.4.2)$$

Lemma 1 (see mathematical appendix), shows that the Marginal Rate of Transformation (MRT) between the two sectors is increasing (i.e. the production possibilities frontier (PPF) between the two sectors is concave). Since markets are competitive, this is akin to saying that:

$$\frac{p_i}{p_j} = \frac{F_{j,H}(H_j)}{F_{i,H}(H_i)} = \frac{F_{j,H}(F_j^{-1}(Y_j))}{F_{i,H}(F_i^{-1}(Y_i))} \quad (1.4.3)$$

is sloping upward in $\left(\frac{Y_i}{Y_j}, \frac{p_i}{p_j}\right)$ space over some range Y . Intuitively, in the short term firms, can only expand by changing their labor input. If one sector expands relative to the other, they must expand by increasing their relative share of labor, which increases their relative marginal cost. An example of this type of production function pair would be $F_i(H_i) = A_i H_i^\alpha$ where $\alpha \in (0, 1)$ and is common across sectors. If both sectors have linear production over labor, then the relative marginal cost curve would be flat. An increasing marginal product of labor would lead to a downward-sloping curve.¹⁵

1.4.2 Households and Intratemporal Substitution

The representative household is given an exogenous endowment of expenditure, X . They have non-homothetic preferences over consumption in the necessity and luxury sectors $U(c_N, c_L)$ such that for prices p_N, p_L and nominal expenditure X over some interval around X , the ordinary demand of the luxury good $C^L(\cdot)$ increases in relation to that of the necessity good with an increase in X :

$$\frac{\partial C^L(X, p_N, p_L)}{\partial X C^N(X, p_N, p_L)} > 0. \quad (1.4.4)$$

¹⁵If sectors each have production over labor, but not of the same curvature (i.e. it violates the assumption of production being homogeneous of degree $k \in (0, 1)$ for each sector) then the relative supply curve is not necessarily upward sloping across the domain. For example, suppose both sectors decrease production, but one sector j decreases production more. Sector j will shrink relative to the other sector, but the actual change in relative marginal costs will depend on the size of the decrease in average production versus the relative decrease in production in sector j .

Since we only have two goods, this implies that when X increases, the share spent on the necessity good s_N decreases.

Figure 1.5 shows a representation of how the relative marginal cost curves (relative supply) and relative demand could look in $(s_N, \frac{p^N}{p^L})$ space. The relative supply curve slopes upward due to homogeneous production of degree $k \in (0, 1)$ in each sector. The relative demand curve can slope upward or downward (as pictured, the downward sloping relative demand implies that the goods are gross substitutes). If there is a decrease in expenditure X , then relative demand for necessities will rise, and the relative demand curve will shift to the right. Equilibrium necessity expenditure share and the relative price will both increase (as pictured, this is a move from point A to B).

The intuition behind figure 1.5 is stated formally in the following proposition (the proof is included in the mathematical appendix).¹⁶




Figure 1.5. Relative supply and relative demand

¹⁶In the proposition, the representative household is assumed to have non-homothetic consumption preferences. However, this is not always the same assumption as the micro-level households having non-homothetic consumption preferences. I discuss this issue in more detail in the mathematical appendix.

Proposition 1 *In a two-sector competitive economy with a representative household that has preferences satisfying equation (1.4.4), production function in each sector $F_i(H_i) : [0, \infty) \rightarrow [0, \infty)$ both homogeneous of degree $k \in (0, 1)$ and standard market clearing conditions, then an decrease/increase in household expenditure will lead to an increase/decrease in the relative price of necessities.*

1.5 Empirical Strategy

The empirical approach centers around: (1) testing how the aggregate relative demand curve shifts in response to a macroeconomic shock and (2) measuring the slope of the relative marginal cost curve. These two questions are directly related to the assumption that the representative consumer has non-homothetic preferences and that the relative supply curve is upward sloping. In order to address both of these questions, I need a macroeconomic shock that will shift *only* the relative demand curve and leave the relative supply curve unchanged. This is important as any shock that directly affects the slope/position of the relative supply curve will obscure efforts to test its slope.

I use monetary policy shocks to test the non-homotheticity of aggregate demand and the corresponding effect on relative prices. Interest rate shocks are typically treated as demand rather than supply shocks, as they directly affect households expenditure and savings, but not relative costs across sectors.¹⁷ This ignores the potential cost channel of monetary policy (Barth III and Ramey 2001), as well as changes in household preferences for sector products that are correlated with the necessity/luxury classification. I partially address these latter concerns in the robustness section.

Since central banks respond to macroeconomic events, making interest rate changes endogenous, there is a large literature using monetary policy news as an external shock on interest rates (Gürkaynak et al. 2004, Swanson 2021, Bauer and Swanson 2020). As a proxy

¹⁷In the textbook New Keynesian model, the interest rate appears only in the household side of the model and operates through the Euler Equation (Galí 2015)

for a monetary policy shock, I use the estimated monetary policy news shock from Swanson (2021). This news shock is computed by looking at the change in a variety of asset prices in a 30-minute window around each FOMC meeting from July 1991-June 2019. I use the first principle component of changes in this vector of asset prices, which corresponds to a change in the interest rate (rather than changes to forward guidance or Quantitative Easing).

In order to test the differential response of interest changes on necessity and luxury product shares and prices, I estimate a local projection of the dependent variable (x_j) on the interaction between the monetary policy shock and the product's expenditure ratio (Jordà 2005):

$$x_{j,t+h} = \beta_0 + \sum_{k=0}^K \beta^k x_{j,t-k} + \gamma^h i_t \times R_j + \delta_t + \psi_j + \varepsilon_{j,t+h} \quad (1.5.1)$$

In the above equation, the dependent variable ($x_{j,t+h}$) is either the log-aggregate share of product j at time $t+h$ or the log-price. The coefficient of interest γ^h (the coefficient of the interaction of the monetary policy shock i_t and expenditure ratio R_j) is the differential response of sector shares/prices based on expenditure ratio, which corresponds to the Blinder-Oaxaca extension to the local projection framework discussed in Cloyne et al. (2020). I include a year of lags of the dependent variable, $\sum_{k=0}^K \beta^k x_{j,t-k}$, so $K = 12$. I also include time fixed effects, δ_t , which absorb the direct effect of monetary policy on shares/prices, as well as any other macroeconomic events occurring at time t . Finally, I include product fixed effects, ψ_j , which control for the average level of share/prices for product j . I compute these local projections on monthly share and price series for the panel of products in the data. In the appendix, I consider alternate specifications: including lags of the interaction of the shock i_t and the expenditure ratio R_j (Ramey 2021a), limiting the shock data to pre-2008 to avoid the zero lower bound period, and including sector-specific time trends. Results for these alternate specifications are shown in figure A12.

If aggregate demand responds non-homothetically to monetary policy shocks, then I would expect γ^h to be positive when the dependent variable is log-share. A positive coefficient

means that products bought more by poor households (the expenditure ratio R_j is higher) increase in price following a contractionary monetary policy shock compared to other products (which the model in the previous section predicts). Furthermore, an upward sloping relative supply curve implies that γ^h in the price regression should have the same sign as γ^h in the demand regression. If γ^h is positive when the dependent variable is log-share, this implies that demand shifts towards necessities (away from luxuries) after a contractionary monetary policy shock and an upward sloping relative supply curve require γ^h to also be positive in the price-regression.

In the model presented in the proceeding section, a fall in expenditure causes households to shift their demand to necessities due to non-homothetic preferences. Accordingly, I test directly how the monetary news shocks affect aggregate expenditure using a simple local projection of Log-real personal consumption expenditure (PCE) on the monetary policy shock (Jordà 2005). I follow Ramey (2016) and include lags of the monetary instrument and lags of the dependent variable. I also include lags of the price level (CPI), one-year treasury yield, the unemployment rate (Leahy and Thapar 2019).

All regressions use standard errors that are clustered at the time level and are robust to serial correlation.¹⁸ Results are scaled so that a one-unit monetary shock corresponds to a 25-basis point increase in the one-year treasury bill. Finally, regressions are weighted by the pooled aggregate share of sector j in consumer spending.

1.5.1 Results

Figure 1.6 shows the impulse response functions estimated following equation (1.5.1). Panel (a) shows the response of the One-Year Treasury yield to the monetary policy news proxy. This result was scaled so that on impact, the one-year Treasury yield increases by 25 basis points. Panel (b) shows the response of log-real consumption; consumption falls by approximately 2 percent two to three years following the monetary shock. Panel (c) shows that aggregate

¹⁸Standard errors are similar when using heteroskedasticity-consistent robust standard errors that are not robust to auto-correlation (Herbst and Johansson 2021, Montiel Olea and Plagborg-Møller 2021).

expenditure shifts towards necessity products following a contractionary monetary shock. The IRF peaks at around 0.05 following the shock, which means that products with an expenditure ratio of 1-point higher than average increase their aggregate share by approximately 5-percent relative to other products. Finally, panel (d) shows how the relative price of necessity goods increases following the monetary contraction. A product with expenditure ratio 1 point higher than average increases in price by around 2-percent, compared to other products, one to three years following the shock. The empirical results provide evidence for the mechanism presented in the static model. Following shocks that lower aggregate expenditure, aggregate spending shifts towards necessities raising their relative prices.

1.5.2 Robustness

The main identifying assumption is that monetary shocks affect product prices differently only to the extent that they shift demand through non-homothetic preferences. However, demand for durables can be more sensitive to interest rate changes than non-durables (McKay and Wieland 2021, Barsky et al. 2007), services tend to have stickier prices (Nakamura and Steinsson 2008), and the central bank can react to oil shocks directly. As a robustness check, I perform a similar local projection to equation (1.5.1), but I include an interaction between the monetary policy shock and dummies for whether the product is a durable, a service, or in the energy or transportation sector (energy). Estimates of γ^h , the differential response of necessities, with these additional controls are shown in Figure 1.7. Results are similar to the baseline for both shares and prices, with the exception of the price response when the energy interaction is included. Here the magnitude of the necessity relative price increase is smaller, but it follows a similar path as the baseline set of local projections.

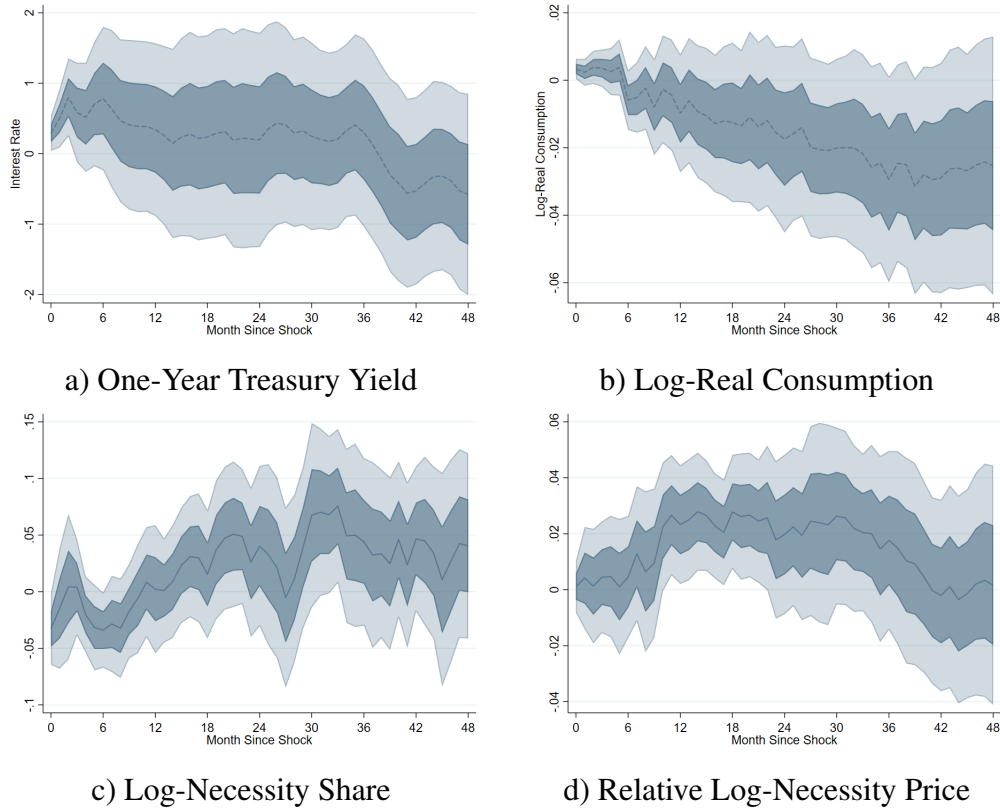


Figure 1.6. IRFs: Response to Monetary Policy Shock

Note: Data from 1991-2019. Estimated coefficients, γ^h from Local Projections in equation (1.5.1). The unit of observation is the month in panels a) and b), and the sector-month in c) and d). Robust standard errors are shown by one- and two- standard error confidence bands indicated by the dark and light shaded areas respectively. Standard errors are robust to auto-correlation and are clustered at the monthly level for panels c) and d). Sectors weighted by their share in pooled aggregate expenditure. Monetary Policy shock normalized to 25-basis point increase in 1-year treasury in month $t = 0$. Figure d) uses a balanced sample of 60 sectors with price data available for the entire period.

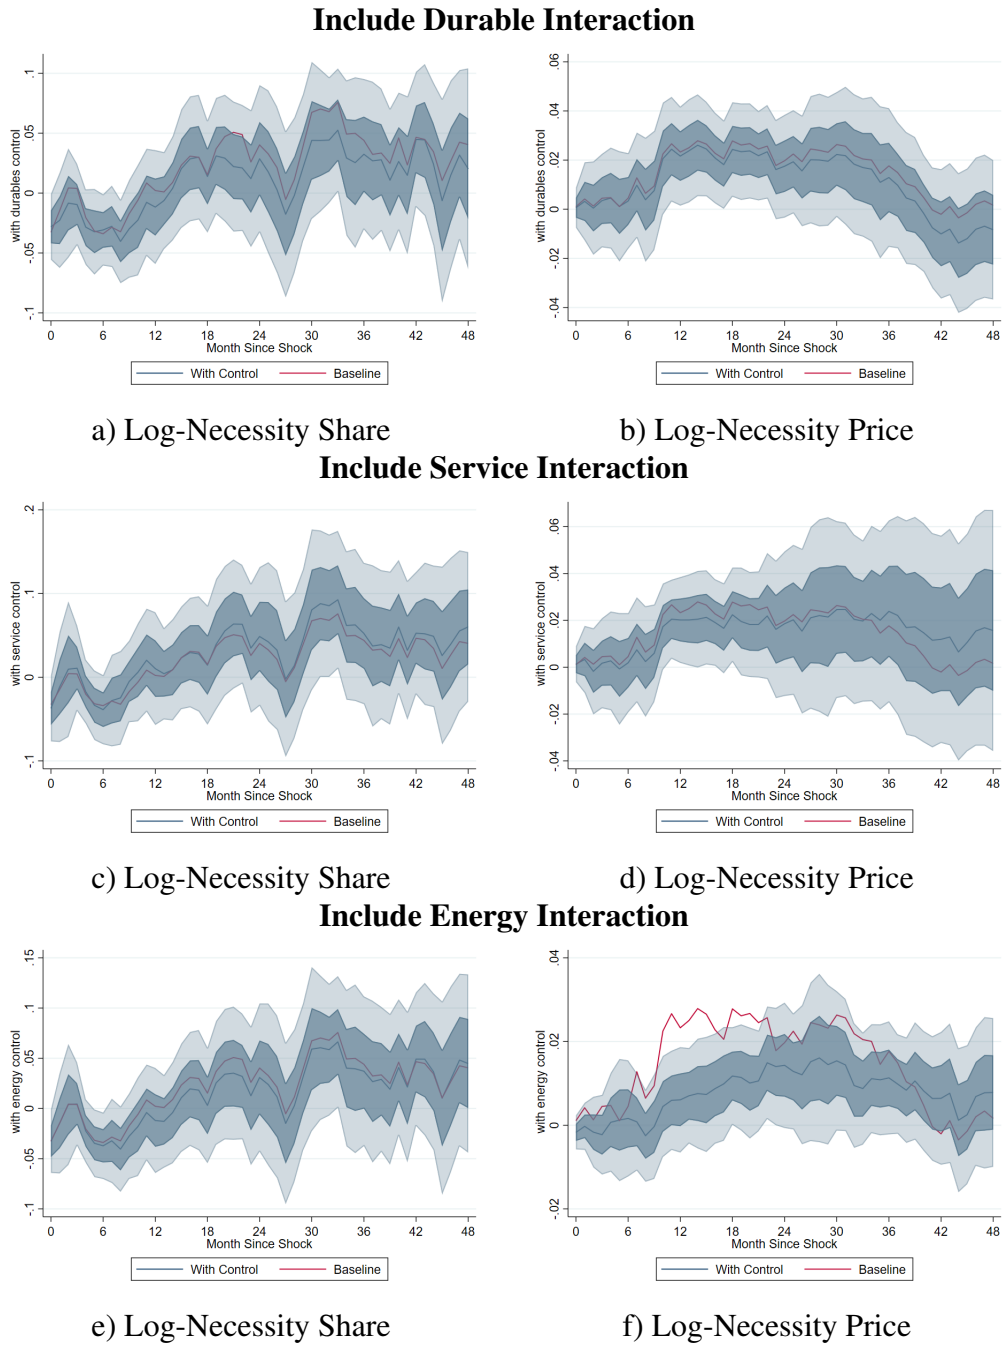


Figure 1.7. IRF Robustness Checks: Response to Monetary Policy Shock

Note: Data from 1991-2019. Estimated coefficients from Local Projections explained in section 5. The unit of observation is the sector-month. Robust standard errors clustered at the monthly level are shown by one- and two-standard error confidence bands indicated by the dark and light shaded areas respectively. Sectors weighted by their share in pooled aggregate expenditure. Monetary Policy shock normalized to 25-basis point increase in 1-year treasury in month $t = 0$. When the dependent variable is log-price a balanced sample is used of 60 sectors with price data available for the entire period.

1.6 New Keynesian Model with Non-homothetic consumption preferences

I have already formally presented the cyclical demand shift mechanism and shown that this mechanism is qualitatively consistent with the empirical results. This section shows that the theoretical results also quantitatively match the cyclical behavior of necessity prices and aggregate shares in the data. I include non-homothetic consumption preferences in a two-sector New Keynesian model with sticky wages and calibrate this model to the U.S. economy in 2005-2006. I then use the model to examine the welfare consequences of the cost-of-living channel of recessions for low- and high-income households.

1.6.1 Households

Intratemoral Consumption Choice: The Almost Ideal Demand System

Household preferences follow the Almost Ideal Demand System (AIDS) first introduced by Deaton and Muellbauer (1980). I choose AIDS for two reasons: (1) the model relies on aggregate demand shifts, and since AIDS is a form of PIG-Log preferences, they are within the Generalized Linear Class of preferences and can be aggregated (Muellbauer 1975). Aggregation is a clear advantage over other types of non-homothetic demand systems, such as the Non-homothetic CES system presented in Comin et al. (2021). AIDS aggregation properties allow me to estimate aggregate parameters using micro-data since the parameters for the representative and micro-level households are the same. The second reason, (2) is that the Almost Ideal Demand System was originally designed to be extremely flexible; in fact, it is a first-order approximation to any demand system (Deaton and Muellbauer 1980).¹⁹

¹⁹A disadvantage is that the AIDS is not generally regular. There are levels of expenditure and prices for which AIDS is not a valid utility function. However, this is not an issue for the calibration and expenditure levels that I study.

The functional form for the household level indirect utility function is:

$$V(X^h, \mathbf{p}) = \left(\frac{X}{a(\mathbf{p})} \right)^{1/b(\mathbf{p})} \quad (1.6.1)$$

where $a(\mathbf{p})$ and $b(\mathbf{p})$ are price aggregators over a vector of sector level prices \mathbf{p} defined by:

$$\log(a(\mathbf{p})) = a_0 + \sum_k a_k \log(p_k) + \frac{1}{2} \sum_j \sum_k \gamma_{jk} \log(p_j) \log(p_k) \quad (1.6.2)$$

$$\log(b(\mathbf{p})) = \sum_j \beta_j \log(p_j) \quad (1.6.3)$$

where γ_{jk} are cross-price semi-elasticities and β_j are expenditure semi-elasticities. Parameters have the following restrictions: $\sum_{j=1}^N a_j = 1$, $\sum_{j=1}^N \beta_j = \sum_{j=1}^N \gamma_{jk} = 0$ and $\gamma_{ij} = \gamma_{ji} \forall i, j$.

The indirect utility function equation (1.6.1) has a corresponding cost function:²⁰

$$\log c(u_h^0, \mathbf{p}) = \log(a(\mathbf{p})) + (b(\mathbf{p})) \log(u_h). \quad (1.6.4)$$

The cost function shows that households must pay some cost for subsistence level consumption $\log(a(\mathbf{p}))$, where $a(\mathbf{p})$ is a homothetic translog price aggregator. The second aggregator, $b(\mathbf{p})$ introduces non-homotheticities into the cost-function. A household's cost to reach a higher level of utility (expenditure) increases with $b(\mathbf{p})$. This allows me to construct the theoretically consistent non-homothetic price index for a household with fixed utility u_h :

$$\log P(\mathbf{p}^1, \mathbf{p}^0, u_h^0) = \log \left(\frac{a(\mathbf{p}^1)}{a(\mathbf{p}^0)} \right) + \log \left(u_h^{b(\mathbf{p}^1) - b(\mathbf{p}^0)} \right) \quad (1.6.5)$$

The greater the household's utility (expenditure) x^h , the higher the welfare gain from reductions in $b(\mathbf{p})$. Similarly, households with a low-expenditure level have changes in the cost of living closer to changes in the subsistence price index $a(\mathbf{p})$.

²⁰This functional form differs from the cost function in Deaton and Muellbauer (1980) due to a slight change in the definition of $b(\mathbf{p})$. If written out entirely, the two cost functions are identical

Roy's identity applied to equation (1.6.1) yields the following Marshallian demand share for products in sector j :

$$s_j = a_j + \sum_k \gamma_{jk} \log(p_k) + \beta_j \left(\frac{x^h}{a(\mathbf{p})} \right). \quad (1.6.6)$$

A household's share of expenditure on a particular product j is dependent on prices and real expenditure level. The demand share increases with real expenditure if $\beta_j > 0$ (luxuries). The household's expenditure elasticity for good j is $1 + \frac{\beta_j}{s_j}$, while the cross price elasticity is $\delta_{jk} + \frac{\gamma_{jk} - \beta_j(\alpha_j + \sum_k \gamma_{jk} \log(p_k))}{s_j}$ where δ_{jk} is the Kronecker delta term.

Household intratemporal aggregate demand can be represented completely by a representative household. However, unlike homothetic preferences, the representative consumer does not have an expenditure level equal to the aggregate household. In the non-homothetic case, the representative consumer's expenditure level must increase with the level of expenditure inequality in the economy. A less equal distribution of expenditure means that high-expenditure households command a larger portion of aggregate spending, which means that the aggregate share spent on luxuries is higher than in an otherwise equivalent economy with lower expenditure inequality. A collection of households with PIG-Log preferences can be represented by a household with income $X^r = X^{mean} \exp\left(\sum \frac{x^h}{X^{mean}} \ln\left(\frac{x^h}{X^{mean}}\right)\right)$ where the term on the right $\left(\sum \frac{x^h}{X^{mean}} \ln\left(\frac{x^h}{X^{mean}}\right)\right)$ is the Theil index of the expenditure distribution, which increases with expenditure inequality Muellbauer (1975), Deaton and Muellbauer (1980).

Intertemporal Consumption Choice and Labor Supply

Each household chooses consumption expenditures to maximize their sum of discounted indirect utility over time.

$$\mathbb{E}_0 \sum_{t=0} \beta^t \left[F\left(V(X_t^h, \mathbf{p}_t)\right) - g(H_t^h) \right]. \quad (1.6.7)$$

where $g(\cdot)$ is the disutility of labor and H is hours worked. $F(\cdot)$ is taken to be the isoelastic utility function:

$$F(y) = \frac{y^{1-\eta} - 1}{1 - \eta}.$$

One feature of Isoelastic preferences, is the the elasticity of intertemporal substitution is generally constant. However, that is not the case in this model. Following Browning (2005), I define the elasticity of intertemporal substitution as:

$$EIS = -\frac{v_x(X_t, \mathbf{p}_t)}{X_t v_{xx}(X_t, \mathbf{p}_t)},$$

where $v(X_t, \mathbf{p}_t) = F(V(X_t^h, \mathbf{p}_t))$. So in this model the elasticity of intertemporal substitution is $-\frac{b(\mathbf{p}_t)}{1-\eta-b(\mathbf{p}_t)}$, which varies with the level of relative prices in the economy (Crossley and Low 2011, Attanasio and Weber 1995). When relative prices for luxuries are higher, this increases the concavity of the indirect utility function making further increases in utility more difficult, which raises the elasticity of intertemporal substitution.

One important thing to note, is the while the elasticity of intertemporal substitution is dependent on relative prices, it does not depend on the households income or expenditure level. Household's disutility of labor also does not depend on household expenditure or income (in this model). So, household intertemporal and labor supply decisions can also be represented by a representative household.²¹ In practice, I solve for equilibrium prices and aggregate shares using the representative household. I can then back out household level price indices given aggregate

²¹While there has been extensive work showing that households intertemporal responses vary based on income level (see Kaplan, Moll, Violante (2018) for an example), heterogeneous intertemporal responses is not the key feature of this paper. Some macroeconomic policies such as the 2020 and 2021 stimulus checks could have first-order effects on relative prices, as only low to moderate-income individuals were given checks. If low-income household expenditure increases sufficiently after such a policy then the Theil Index could rise enough to partially offset aggregate increases in expenditure.

prices. This approach has the advantage of being able to study welfare effects with heterogeneous consumption bundles using the large toolbox of solution methods for representative agent models.

The representative household works for wages W_t and can invest in a one-period nominally riskless bond B_t that pays one monetary unit in the next period at price Q_t . The resulting household budget constraint and the no-Ponzi scheme condition are shown below:

$$X_t + Z_t Q_t B_t \leq B_{t-1} + W_t H_t + D_t \quad (1.6.8)$$

$$\lim_{T \rightarrow \infty} \mathbb{E}_t (\Lambda_{t,T} B_T) \geq 0.$$

In the above expression, D_t is a dividend from firm profits and $\Lambda_{t,T} = \beta^{T-t} \frac{V_{X,T}}{V_{X,t}}$ where β is the discount factor. Z_t is an interest rate wedge shock that is distributed *i.i.d* and acts to dampen or increase a household's per-period expenditure.

The household's optimization problem and budget constraint yield the following Euler Equation:

$$Q = \beta \mathbb{E} \left[\frac{a(\mathbf{p})b(\mathbf{p})}{a(\mathbf{p}')b(\mathbf{p}')} \frac{\left(\frac{X'}{a(\mathbf{p}')}\right)^{\frac{1-\eta}{b(\mathbf{p}')} - 1}}{\left(\frac{X}{a(\mathbf{p})}\right)^{\frac{1-\eta}{b(\mathbf{p})} - 1}} \frac{1}{Z} \right]. \quad (1.6.9)$$

I assume that the disutility of labor takes the familiar form (with ϕ the inverse of the Frisch elasticity of labor supply):

$$g(H_t) = \varphi \frac{H_t^{1+\phi}}{1+\phi}. \quad (1.6.10)$$

However, households do not decide how much labor to provide. Rather, they allow a labor union to bundle and sell their labor, which introduces sticky wages and nominal rigidity (see Erceg et al. (2000), Auclert et al. (2018), Auclert et al. (2020), Broer et al. (2020), Ramey (2020)). The

mathematical appendix shows that the Wage-Phillips curve is:

$$(1 + \pi_t^w) \pi_t^w = \beta \mathbb{E}_t [(1 + \pi_{t+1}^w) \pi_{t+1}^w] + \left(\frac{\varepsilon_w}{\psi_w} \right) \left(\phi H_t^\phi - \left(\frac{\varepsilon_w - 1}{\varepsilon_w} \right) \frac{W_t}{a(\mathbf{p}_t) b(\mathbf{p}_t)} \left(\frac{X_t}{a(\mathbf{p}_t)} \right)^{\left(\frac{1-\eta}{b(\mathbf{p}_t)} - 1 \right)} \right) \quad (1.6.11)$$

1.6.2 Firms

There is a necessity and a luxury sector. Each sector has flexible prices and perfect competition. Firms have concave production over labor; they can scale up labor in the short run, but other factors of production are constrained. The production function for the representative firm in sector i is:

$$Y_t(i) = A_{it} H_t(i)^{1-\alpha} \quad \alpha \in (0, 1). \quad (1.6.12)$$

Firms sell their good for price $p_t(i)$ in a competitive market. Firms take prices and wages as given. Firm optimization implies that:

$$p_t(i) = \frac{W_t}{(1-\alpha) A_{it} H_t(i)^\alpha}. \quad (1.6.13)$$

This yields a relative supply curve, that is upward sloping:

$$\frac{p_t(i)}{p_t(j)} = \frac{A_{jt} H_t(j)^\alpha}{A_{it} H_t(i)^\alpha}. \quad (1.6.14)$$

The elasticity of marginal cost to an increase in output, which governs the slope of the relative supply curve, is $\frac{\alpha}{1-\alpha}$.

1.6.3 Equilibrium

An equilibrium for this model is defined as series of prices $\{W_t, \mathbf{p}_t\}$ and quantities $\{Y_{N,t}, Y_{L,t}, H_t, \mathbf{h}_{j,t}, X_t, D_t, \mathbf{s}_{j,t}\}$ such that households optimize intertemporally and intratemporally given prices, the union chooses labor to maximize household utility, firms maximize profits given prices, and markets clear.²²

1.6.4 Calibration

The two most important parameters for the model are $\beta_L = -\beta_N$ the degree of non-homotheticity, and α , which is one minus the labor share. The first is important since it governs the degree to which representative household spending shifts between sectors over the course of the business cycle. For example, a value of $\beta_L = -\beta_N = 0$ would imply that the household has homothetic preferences, and macroeconomic shocks would not affect the relative demand for necessities or luxuries. The second, α , controls the price response of the expanding sector.

In the baseline calibration, I choose β_L so that the steady-state necessity share for low- and high-income households in the model match that of low- and high-income households in the data. In an alternate calibration, I estimate β_L and the other (AIDs) parameters directly from the microdata; the results of this alternate calibration are in the appendix.

There are a variety of estimates of α , the capital share, in the literature. These can range from as low as 0.16, the implied value based on the estimated elasticity of marginal cost to quantity produced from Feenstra and Weinstein (2017), to as high as 0.37 estimated directly in Fernald (2014). For the baseline specification, I choose α as the midpoint of these extreme values ($\alpha = 0.26$). Alternate calibrations with other values of α are included in the appendix.

The remaining parameters I either take from the literature, or from targeting the steady-state expenditure and necessity share of the representative agent to match representative expen-

²²There is also a central bank that uses a Taylor rule to set interest rates:

$$-\log(Q_t) = i_t = F(\pi_t^w) \quad (1.6.15)$$

where $F(\cdot)$ is increasing in wage inflation.

Table 1.5. Baseline Calibration

Parameter	Desc.	Value	Source
α	Capital Share	0.26	(Midpoint Fernald (2014) and Feenstra and Weinstein (2017))
β	Discount Rate	.99	
$1/\eta$	EIS at Steady State	0.5	
ϕ	Inverse Frisch Elasticity	1	
ψ_w	Wage Adjustment Penalty	20.7	(Wage Phillips Slope 0.29 Galí and Gambetti (2019))
ε_w	Substitutability of labor	6	(Colciago 2011)
β_L	Degree of non-homotheticity	0.29	(Target High- and Low- income steady state necessity shares)
γ_{LN}	Cross-price semi-elasticity	0.95	(Feenstra and Weinstein 2017)
α_N		2.9	(Target necessity share 0.53)

diture and aggregate necessity shares in the period immediately preceding the great recession (2005-2006).²³ I target the calibration, so that in steady-state necessity and luxury prices are equal (which means that the Elasticity of Intertemporal Substitution is equal to $1/\eta$). Table 1.5 shows the chosen calibration.

1.6.5 Results

How well can the calibrated model explain the distribution of household consumption and historical changes in necessity shares and prices? I start by comparing the steady-state necessity shares in the model with those in the data. While I targeted the aggregate steady-state share of necessities and those from the top and bottom income groups, the other income groups' necessity share was not targeted. Figure 1.8 shows the model implied necessity shares for the five different income groups alongside their actual values in the data (2005-2006). In the data, low-income households spend around 70 percent of their budget on necessities compared to around 50 percent for high-income households, which by design, the model matches exactly. The model also matches the necessity shares for the non-targeted income groups within 2 percentage

²³Representative expenditure in the data is average expenditure multiplied by the calculated Theil index.

points.




Figure 1.8. Model and Data: Necessity Shares by Income Group

Note: Data from 2005-06. Model income-group shares at steady state. Author targeted calibration so model necessity shares for the top and bottom income quintiles would match empirical necessity shares. Necessity shares for the middle income quintiles are untargeted.

Historical Simulation

How well does the model predict necessity prices and shares over time? As a validation exercise, I shock the model with a series of i.i.d. interest wedge shocks so that the expenditure series in the model exactly matches the filtered real personal consumption series from the BEA. I then compare the necessity share and price series in the simulated model with their filtered counterparts in the data. Figure 1.9 shows the results of this simulation. The data series of prices and shares excludes the volatile energy and transportation sectors.

The top panel shows the path of both model and data expenditure from 1994-2019. The second panel shows the untargeted model necessity share series compared to the data.²⁴ Similar

²⁴The expenditure share series begins in 1991, but filtering necessitates dropping the first few years of data.

to the data, the model necessity share series falls during the late 1990s, rises around the 2001 recession, falls again during the housing boom, increases drastically during the Great Recession, falls again in the subsequent recovery and then rises during the Covid-19 Pandemic. The time series in the model and data are highly correlated (0.55), and a simple regression of the data series on the model series yields a coefficient of 0.6.

The bottom panel compares relative necessity prices in the data with the cyclical component of the composite necessity price in the data. I use a balanced sample of products with continuous price data from 1987-2021 (this is the red series in figure 1.3). The data and the model series match each other quite closely, however the model overstates the fall in necessity prices during the Dot-Com boom (late 1990s) and the rise in necessity prices during the Covid-19 recession. A simple regression of the data series on the model series yields a coefficient of 0.54.²⁵ I conclude that the model is highly effective at predicting the cyclical path of relative necessity shares and prices.

Welfare Implications

What are the welfare implications of this model? In this model, the expenditure inequality of households is fixed at the steady-state level. However, households price indices can diverge since low-expenditure households spend more of their budget in the necessity sector. How much can this matter? Table 1.6 shows the difference in the non-homothetic price index (equation (1.6.5)) between households with expenditure matching expenditure in the bottom income quintile with expenditure in the top quintile. During the great recession, the price index of poor households increased by 0.85 percent more than rich households. This result closely matches the difference in the change in core inflation in the data over this same period (0.86 see figure A6). Failing to incorporate changes in the price index could lead to large underestimates of the change in consumption inequality over the Great Recession. For example, Krueger et al. (2016) use the PSID and find that household consumption in the first wealth

²⁵Correlation coefficient is 0.44.

quintile fell by approximately 0.3 percent more than consumption in the highest quintile from 2006-2010. A back of the envelope calculation suggests that the change in real consumption is $\Delta \ln\left(\frac{c}{p}\right) = \Delta \ln(c) - \Delta \ln(p) = 0.0115$ or 1.15 percent, which is nearly a fourfold increase compared to Krueger et al. (2016).

While the model predicts that this price index gap will eventually close (as the model returns to steady-state), the price index of the lowest income quintile remains elevated during the slow recovery (GDP per-capita did not return to pre-great recession levels until 2013Q1). The average difference in the cost of living from the beginning of the great recession until GDP per capita recovered is 0.5 percentage points.

Table 1.6. Welfare Difference Low v. High Income Households

Difference in Price Index		
Time Period	End Period	Average
Great Recession (2007Q3-2009Q2)	0.85	0.42
Recession to Recovery (2007Q3-2012Q4)	0.12	0.51
Expenditure Equivalent Welfare Loss		
	Low Income	High Income
Expenditure Equivalent Welfare Ratio	0.59 %	0.48 %
		1.22

Note: Price Index difference is defined as the percentage point difference in the change of the cost-of-living for Q1 v. Q5 households as calculated in the model. Expenditure equivalent welfare is the present discounted value of *all* future expenditure the household would be willing to forgo in exchange for avoiding shocks lead

Next, I calculate the expenditure equivalent welfare loss of the Great Recession for a household in the lowest income group and the highest. This measure is the present discounted value of all future expenditure streams that the household would relinquish in order to avoid the Great Recession:

$$\mathbb{E}^{No\ Recession} \left[\sum_{t=0}^{\infty} \beta^t (V((1 - \xi)X_{ht}, \mathbf{p}_t) - g(H_t)) \right] = \mathbb{E}^{Recession} \left[\sum_{t=0}^{\infty} \beta^t (V(X_{ht}, \mathbf{p}_t) - g(H_t)) \right] \quad (1.6.16)$$

where ξ is the share of all future expenditure the household would relinquish so that the present discounted value of all future utility streams is equal in the counterfactual world where the Great Recession never happens. Table 1.6 shows that low-income households would be willing to give up 0.59% of all future expenditure, while high-income households would relinquish only 0.48%, a difference of approximately 22%. A similar model where the level of non-homotheticity (β_L) is set to 0 results in no difference in welfare loss between low- and high-income households.

1.7 Conclusion

In this project, I present new evidence on the cyclical behavior of necessity and luxury prices. I create a new dataset combining dis-aggregated CPI price indices with micro-level CEX data, and I find that the prices and aggregate shares of products bought relatively more by low-income households are counter-cyclical. I show that these facts likely come via demand shifts by testing how aggregate necessity prices and shares respond to monetary policy shocks. I show that a model with non-homothetic preferences and an upward sloping relative supply curve can jointly reconcile these empirical facts. The calibrated model can explain around half of the cyclical variation of necessity prices and shares. I find that recessions can be more costly for low-income households as their price index increases relative to the price-index of other households.

It is important to note that this project studies changes in sector-level prices rather than prices within a sector; e.g. furniture is a category made up of many different micro-products with their own quality and prices. This project also ignores product entry and exit, which could also impact income-level cost-of-living (Feenstra 1994). To the extent that cyclical demand shifts occur within product categories, causing price increases for low-quality products or changes in product variety (at the business cycle frequency) is a topic for future research.²⁶

This study also has ramifications for the measurement of aggregate changes in the Cost of

²⁶Jaimovich et al. (2019) show that household engage in quality downgrading within sectors during the Great Recession.

Living. For example, in the measurement of the Consumer Price Index (CPI), the BLS uses the Consumer Expenditure Survey to weigh product sectors so they are representative of spending by the average household. However, these weights are only updated with a lag (typically 36 months). Since my study shows that *aggregate* spending shifts to necessities during recessions, that means that the CPI underweights necessities in recessions and overweights them during expansions. This implies that measurement of inflation via the CPI is potentially biased downward during both recessions (when necessity prices are rising more rapidly) and during expansions (when luxury prices are rising more rapidly).

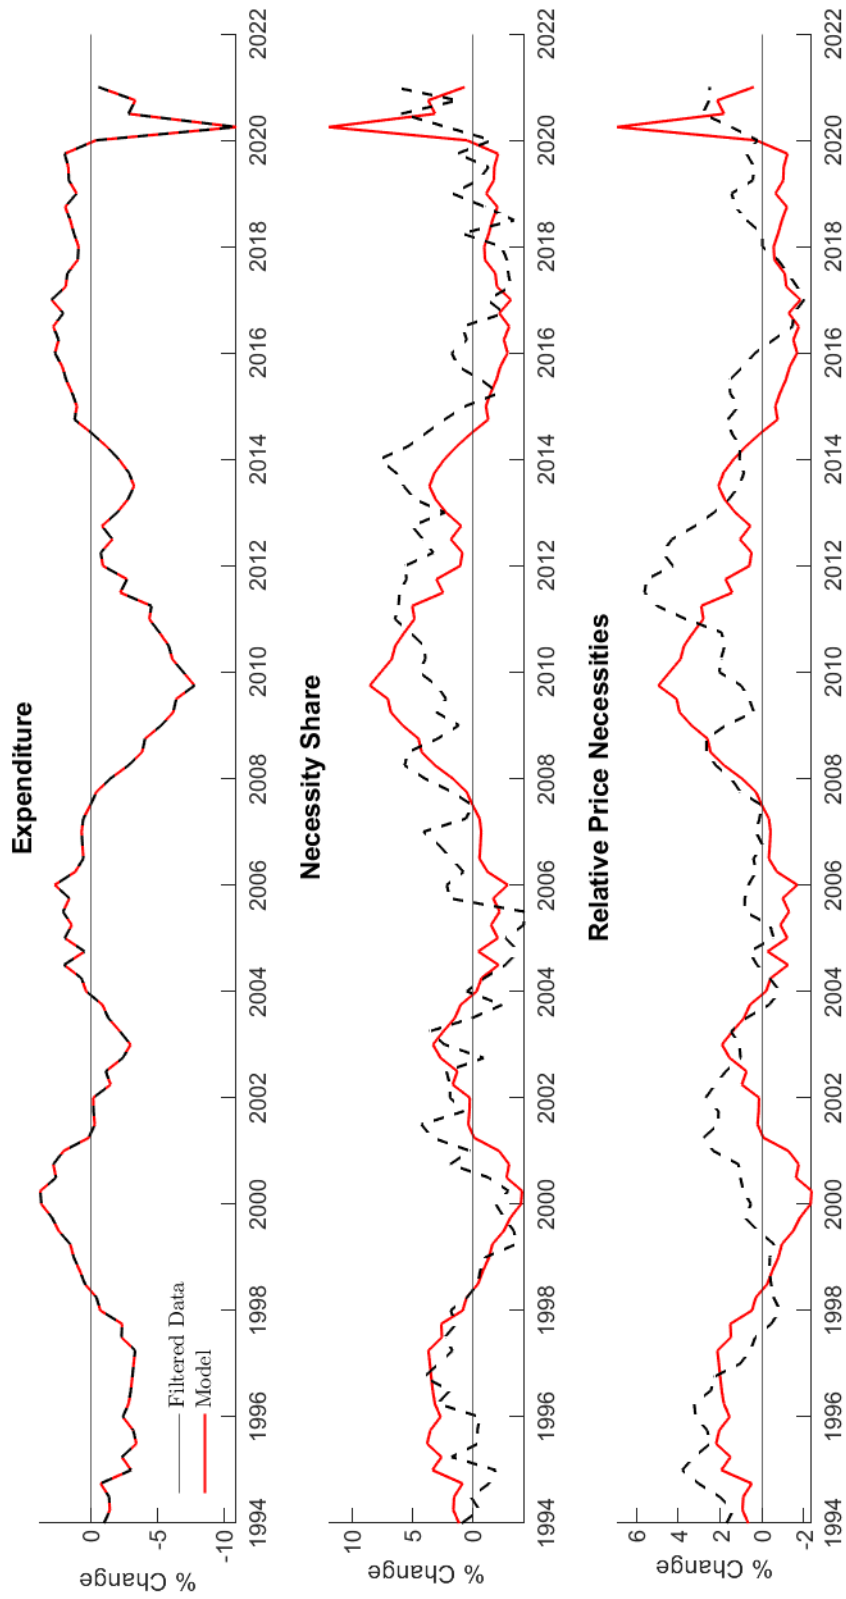


Figure 1.9. Model v. Data: Historical Simulation
 Author targeted shock to match expenditure data (top panel). Necessity share and relative necessity price are untargeted. Data is filtered following Hamilton (2018) and excludes energy.

1.8 Chapter Acknowledgments

Chapter 1, in full, is currently being prepared for submission for publication of the material. The dissertation author was the sole investigator and author of this paper.

Chapter 2

Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates

2.1 Introduction

Numerous studies in the last twenty years have used panel data from households to estimate the marginal propensity to consume out of anticipated, temporary changes in income. Some of the leading studies in this area estimate the effects of the temporary tax rebates of 2001 and 2008. For example, the Shapiro and Slemrod (2003; 2009), Johnson et al. (2006), Sahm et al. (2012), Parker et al. (2013), and Broda and Parker (2014) analyses are exemplars in the use of natural experiments to obtain estimates of this key micro parameter of interest to macroeconomists. Moreover, in some of the best examples of entrepreneurial data collection, these authors added special questions to existing household surveys in order to match the household behavior to the timing of its receipt of the rebate. Shapiro and co-authors found smaller marginal propensities to consume (MPCs), around 30 percent, but Parker and co-authors found some very high estimates. For example, Parker et al. (2013) found a marginal propensity to spend out the temporary tax rebate of 50 to 90 percent on total consumption within three months of receiving the 2008 tax rebate (p. 2531, Table 3).

Estimates from these studies have motivated the thriving literature on heterogeneous agent models in which some households live hand to mouth because of myopia or financial market imperfections. The estimates have been used to calibrate a wide variety of macro New

Keynesian heterogeneous agent models and to argue that temporary tax rebates can have large aggregate multipliers. For example, Kaplan and Violante (2014), Kaplan et al. (2018b), and Auclert et al. (2019) calibrate their heterogeneous agent models to match an MPC of 25 percent on the nondurables component of consumption expenditures. Government policy in recent years has been guided by the high MPC estimates.

In this paper, we present evidence that the high estimated MPCs from the leading household studies result in implausible macroeconomic counterfactuals. Using the 2008 tax rebate as a case study, we calibrate a standard medium-scale New Keynesian model with the estimated MPCs to construct counterfactual macroeconomic consumption paths in the absence of a rebate. The counterfactual paths imply that consumption expenditures would have plummeted in spring and summer 2008 and then would have recovered when Lehman Brothers failed in September 2008. We use narratives and forecasts to argue that these paths are implausible.

In their early analyses of the aggregate effects of the tax rebates of 2008, Feldstein (2008) and Taylor (2009) found little evidence of a response in aggregate consumer expenditures and suggested that consumers mostly saved the rebate. However, their aggregate analyses were soon overshadowed by the impressive household-level analysis conducted by Parker et al. (2013) and Broda and Parker (2014), which estimated very high propensities to consume out of the rebates.

Sahm et al. (2012) also estimated micro MPCs out of the 2008 rebate from rich survey data, but found lower MPCs than the other household-level studies. They conducted an interesting counterfactual analysis using the Parker et al. (2013) estimates. In particular, they used the Parker et al. (2013) estimate of the marginal propensity to spend the 2008 rebate on new vehicles to calculate the implied fraction of actual motor vehicle sales that were induced by the rebate. They noted that this estimate was "surprisingly high" given that there were no dramatic shifts in motor vehicle sales around that time.¹ They pointed out, however, that their exercise does not allow for any partial or general equilibrium effects.

¹See p. 242 and Table 14 of Sahm et al. (2012). Sahm et al. (2010) compare their own micro MPC estimates to total aggregate consumption in a similar exercise.




Figure 2.1. Expenditures on New Motor Vehicles: Actual vs. Counterfactual

Note. Based on Sahm, Shapiro, and Slemrod calculations applied to revised data.

The literature has overlooked Sahm et al. (2012) important calculation, perhaps because it appears in a table at the end of the paper. To demonstrate the striking implied counterfactual path, we update the numbers from their table, calculate the counterfactual path, and graph it relative to the actual path. Figure 2.1 shows actual expenditures on motor vehicles as the green solid line, along with the implied counterfactual spending estimate depicted by the red dashed line. This counterfactual is created as the difference between the actual spending and the estimated induced spending from the rebate.

The graph shows that had there been no tax rebates, expenditures on motor vehicles would have declined by over 85 percent from \$17.3 billion in March 2008 to only \$2.6 billion in June 2008 and then would have rebounded sharply in late summer, averaging \$14.4 billion per month in August and September 2008. This counterfactual strains credulity, especially since the lowest actual level of motor vehicle expenditures during the Great Recession was \$11.7 billion in April 2009.²

In this paper, we extend the logic of the Sahm et al. (2012) exercise to a dynamic general

²The appendix contains details of the calculation. It also shows that when we allow consumers to smooth the spending over more months, the counterfactual remains implausible.

equilibrium setting to study the implications of estimated micro MPCs for the counterfactual path of consumption in 2008 with no rebates. Our method proceeds as follows. We first construct a medium-scale two-good, two-agent New Keynesian (TG-TANK) model in which some households are life-cycle permanent income households and others are “hand-to-mouth” households who consume all their current income each period. We calibrate the fraction of hand-to-mouth households in the economy to match the MPC estimates from the household-level data. In this model, aggregate consumption rises due to both the direct micro effect of the rebate on consumption at the household level and the induced macroeconomic effect on income through Keynesian multipliers. We call the sum of these two effects on aggregate consumption per dollar of rebate the *general equilibrium marginal propensity to consume* out of the rebate, or GE-MPC for short. We then use the model to simulate the macroeconomic effects of a path of rebates that matches the timing and size of the actual 2008 rebate, which was announced in February and distributed mostly from April through July 2008. To create the counterfactual path of aggregate consumption in 2008 with no tax rebate, we multiply actual aggregate NIPA consumption by the ratio of the model-simulated consumption path to the model steady state.

The counterfactual paths created from our baseline simulations with average household MPCs above 0.2 imply that the path of aggregate consumption in the U.S. economy would have been V-shaped from April 2008 through August 2008 had there been no rebates. Specifically, the counterfactual implies that consumption would have collapsed from May through July 2008 and recovered in August and September 2008, when Lehman Brothers failed.

Our argument that the counterfactual path of consumption is implausible rests on three pillars: (i) a credible macroeconomic model that produces dynamic general equilibrium responses of aggregate consumption to rebates; (ii) the absence of other factors that would have led to a collapse of consumption in summer 2008; and (iii) aggregate monthly consumption data that accurately capture the spending effects of the rebates. For the first pillar, we use a standard New Keynesian model that features the type of general equilibrium amplification that is widely used in the literature and policy models. We allow more lags in the response to spending to the rebate

than estimated in order to mute the V-shape, yet the implied paths are still implausible. For the second pillar, we demonstrate that other events, such as the dramatic peaking of gasoline and other energy prices in July 2008 or the bankruptcy of Lehman Brothers in September 2008, were unlikely to have induced a V-shape of consumption absent rebates. Our evidence is based on both professional forecasts at the time and our own time series forecasts using a variety of alternative assumptions. Neither the professional forecasts nor any of the variations on our forecasting model predict a V-shape in consumption in late spring and summer of 2008. For the third pillar, we present evidence that monthly NIPA consumption does not mismeasure the consumption response during that period. To explore the possibility that aggregate consumption rose more than is reflected in the monthly NIPA numbers, we study how alternative measures of consumption, such as unit sales of automobiles, retail sales, and our own aggregates constructed from the Consumer Expenditure Survey (CEX), behaved during this period. We find no evidence of a burst in aggregated consumption from any of those sources that would be consistent with a high MPC.

Our claim about counterfactual aggregate consumption paths begs the question: how does one reconcile the high estimated micro MPCs from the literature with the implausible general equilibrium counterfactuals? One possibility is that general equilibrium forces, rather than magnifying the micro MPCs, actually dampen them. A second possibility is an upward bias in the existing household MPC estimates. We explore each of these explanations and conclude that both are key to explaining the implausible counterfactuals.

To assess the impact of dampening general equilibrium forces, we recalibrate our New Keynesian model, which has a perfectly elastic supply of durable goods, to one with a supply elasticity of five. We find that this dampening goes far toward eliminating implausible counterfactuals. However, this dampening means that even high micro MPCs do not result in sizeable Keynesian general equilibrium multipliers. Relative to the analysis in Wolf (2021), we find significant crowding out of tax rebates because they were spent on motor vehicles which have more elastic demand than nondurable goods.

With regard to a possible upward bias in the existing household MPC estimates, we re-examine the Parker et al. (2013) estimates from the CEX in light of the recent econometric papers highlighting potential problems with event studies. Those papers, such as Sun and Abraham (2020), Borusyak and Jaravel (2017), Borusyak et al. (2022) and others, have raised questions about event study estimates based on standard OLS two-way fixed effects estimators. These estimators implicitly adopt the assumption that the treatment effect is homogeneous in the population. To maximize efficiency these estimators then assign large weights to certain treatment effects and small or negative weights to others. When treatment effects are heterogeneous, this weighting scheme can result in estimates of the aggregate treatment effect that are very different from the Average Effect of Treatment on the Treated (ATT). We apply Borusyak et al. (2022) new method for computing an average MPC among treated households in the CEX data and find smaller estimates of the MPC than the original Parker et al. (2013) paper does.

The combination of dampening general equilibrium forces and more modest micro MPC estimates yields macroeconomic counterfactuals that we consider plausible. However, they also imply that the effect of the rebate on consumption expenditures in general equilibrium was modest. With our preferred micro MPC of 0.3, we find that the general equilibrium increase in total consumer spending was only 16 cents per dollar of the total rebate.

Section 2.2 begins with a narrative of details of the 2008 tax rebate and the behavior of other key variables in 2008. It then presents alternative measures of consumption expenditures that support the patterns indicated by the NIPA data. Finally, it presents contemporaneous forecasts as well as our forecasts for consumption in 2008 before the rebate was passed. Section 2.3 presents the counterfactual experiments. It begins by presenting a medium-scale New Keynesian model with two goods and two types of agents. It then calibrates the model and uses it to perform DSGE simulations of the effects of rebates. It uses the simulated impulse responses to infer what actual consumption would have been had there been no rebate. It then modifies the model to incorporate more dampening effects in general equilibrium to produce alternative counterfactual paths. Section 2.4 re-examines the micro MPC estimates. It begins with a brief discussion

of potential issues with past micro MPC estimates and then applies Borusyak et al. (2022) to re-estimate micro MPCs from 2008. Section 2.5 summarizes and concludes.

2.2 The U.S. Macroeconomy in 2008

This section sets the stage for thinking about the plausibility of counterfactual paths by reviewing the tax rebates and the behavior of other key macroeconomic aggregates in 2008. The first subsection reviews the nature and timing of the tax rebates and then shows the behavior of disposable income, consumption, inflation, oil prices, and monetary policy. The second subsection provides alternative measures of consumption expenditures that support the patterns displayed in the standard NIPA measures. The third subsection shows two types of forecasts of consumption in 2008. The first type is professional forecasts of aggregate consumption, based on information before the rebates were passed. The second is our own set of time series forecasts of consumption during the Summer 2008.

2.2.1 Narrative of 2008

In early January 2008, numerous forecasters and policymakers began to discuss the possibility of a recession in 2008. The employment report released on January 4, 2008 showed a jump in the unemployment rate from 4.7 percent to 5 percent in December; this jump followed an earlier rise from a low of 4.4 percent in May 2007. After release of the report, Goldman Sachs forecasted that the U.S. was either in a recession or would enter one shortly, but predicted that it would be a mild downturn. That forecast assumed that the federal funds rate target would be cut from 4.25 to 2.5 by the end of the year, with the first 50 basis point cut at the next FOMC meeting on January 30th.

In fact, the Federal Reserve enacted an inter-meeting cut in the funds rate of 75 basis points on January 23rd, and another 50 basis points at the January 30th FOMC meeting. The Greenbook forecasts prepared for that meeting did not predict declines in GDP or consumption expenditures in any quarter during 2008, but the New York Federal Reserve Bank's Blackbook




Figure 2.2. 2008 Tax Rebates

Notes. Data from Shapiro and Slemrod (2009). Rebates are nominal. Vertical red dashed line indicates May 2008.

was more pessimistic, predicting an annualized decline in real GDP of -0.8 percent in the first quarter of 2008 with a recovery starting in the second quarter.

The Congress and Administration also recognized that the economy was slowing. They began to discuss tax rebates in January and quickly enacted them in February 2008. Both houses of Congress passed the legislation in the first week of February and President Bush signed it on February 13th. As a result, \$100 billion in rebates were distributed from April through July 2008 to approximately 85 percent of households. The \$100 billion in rebates was large, totaling eleven percent of January disposable income (measured on a monthly basis). The amount of the rebate depended on tax status and dependents and was phased out at higher income levels. Among households receiving a check, the average amount was \$1,000. The timing of distribution was randomized according to the last two digits of the Social Security number. The actual time path of the rebates is shown in Figure 2.2. The graph shows that almost half of the total amount was distributed in May alone, with most of the remaining rebates distributed in June and July.

Figure 2.3 shows the behavior of nominal and real NIPA disposable personal income and consumption from mid-2007 through mid-2009. The vertical red dashed line indicates May

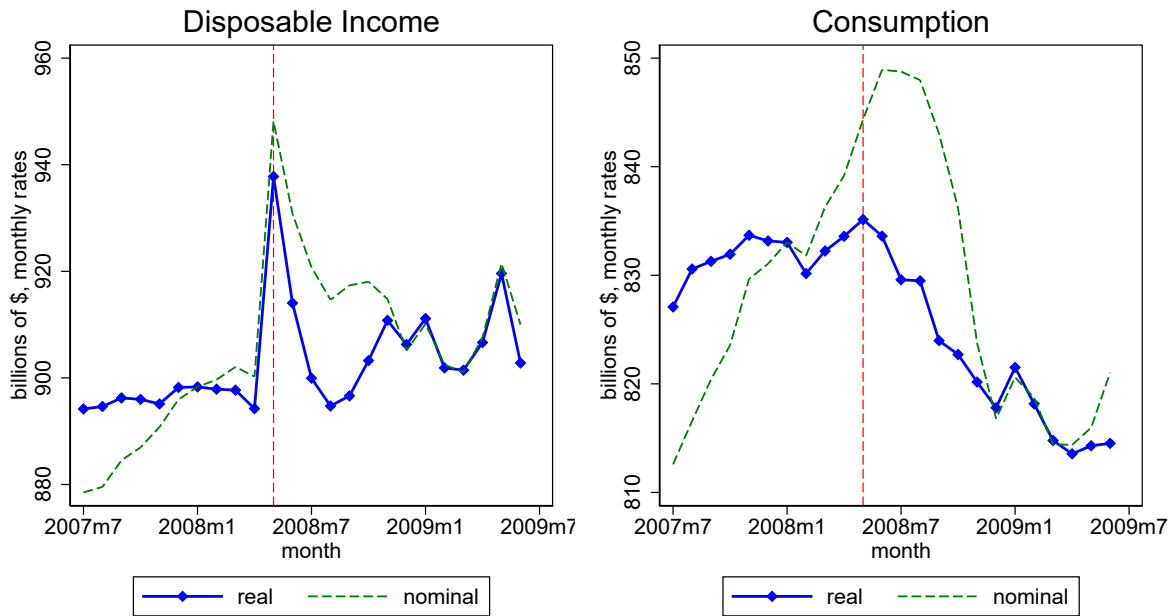


Figure 2.3. Aggregate Disposable Income and Consumption

Source. BEA data. Vertical red dashed line indicates May 2008.

2008 when almost half of the rebate checks were distributed. We normalize real income and consumption to be equal to nominal values in January 2008 for better illustration. Also, note that the range of the y-axis of the graph for disposable income is \$80 billion, which is twice the range of the graph for consumption, which has less variability.

The effect of the 2008 tax rebate on disposable income is clearly evident in the spikes in both the nominal and real disposable income series, shown in the left panel. For both disposable income and consumption, however, the nominal and real paths look quite different from each other because of the behavior of inflation. After falling in February, real consumption rises to a peak in May 2008 before falling through the end of 2008. The sharpest decline is between August 2008 and September 2008, and was likely due to the shock wave caused by the fall of Lehman Brothers in mid-September. Nominal consumption shows a prominent hump in Summer 2008, but real consumption displays only a small bump.

Figure 2.4 shows real consumption expenditures disaggregated by type: nondurable

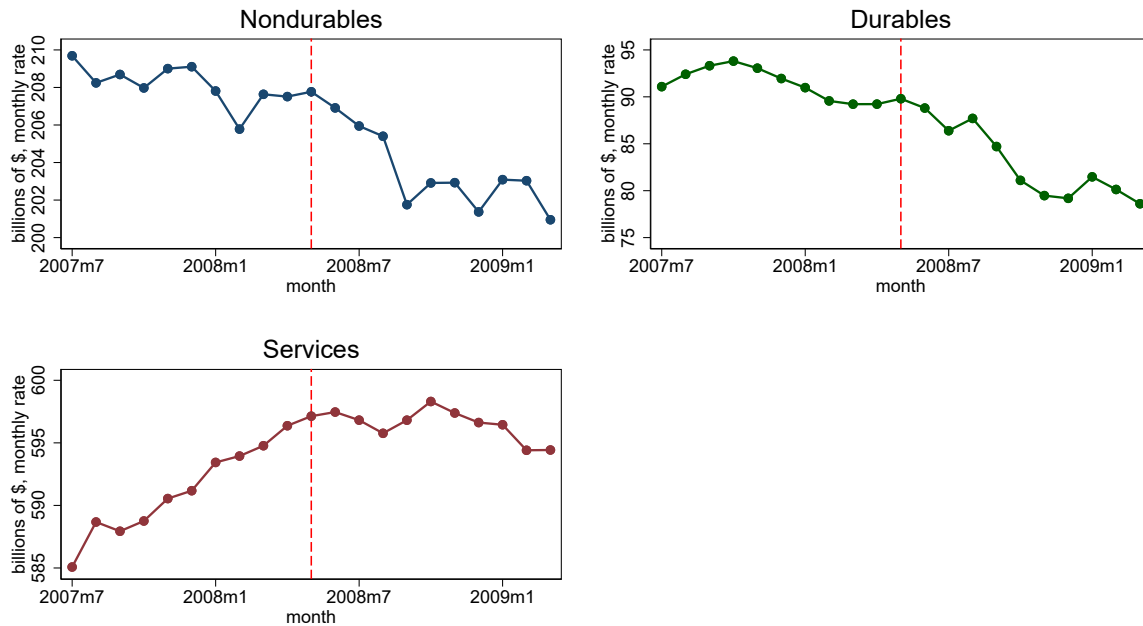


Figure 2.4. Real Consumption Expenditures by Type of Product

Source. BEA data. Vertical red dashed line indicates May 2008.

goods, durable goods, and services. In general, consumption of goods (both nondurable and durable) decline over this period whereas consumption of services rises. In none of the three aggregates is there much evidence of a big boost to spending in May through July 2008.

We now turn to the behavior of other key factors that might have influenced consumption expenditures. The first is the behavior of consumer prices. Figure 2.5 shows the price indices for total consumption expenditures and consumption expenditures excluding food and energy, transformed to logarithms so that the slope of the path indicates the inflation rate. Consider first the behavior of the price deflator for total consumption. The rate of inflation for total consumption accelerated after April, resulting in July prices that were 1.6 percent above April prices. Price levels then reached a plateau and fell after the failure of Lehman Brothers in September, so that by the end of the year the level of prices was slightly lower than at the start of the year.

In contrast, the price index for consumption excluding the volatile food and energy

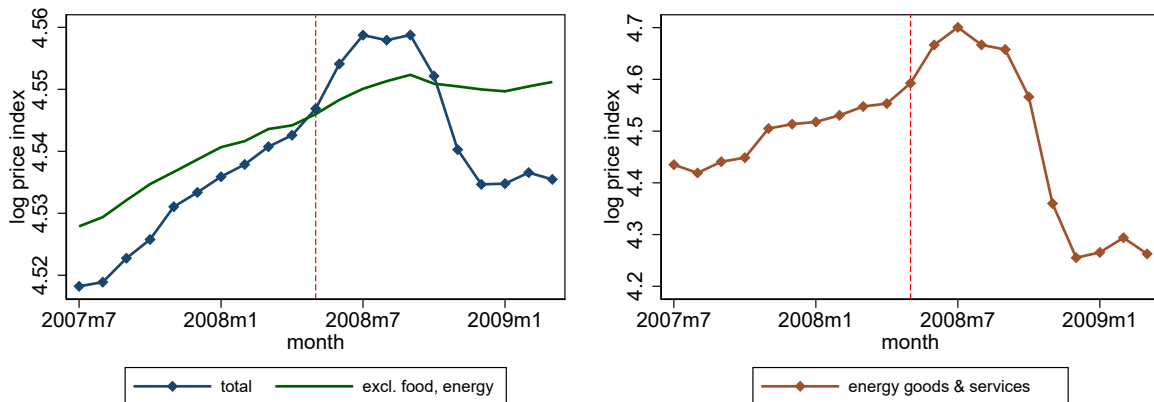


Figure 2.5. Log Price Indices for Consumption

Source. BEA data.

components showed a more modest rate of inflation, averaging 3.4 percent annualized for January through the peak in September 2008. This price level then declined slightly after the collapse of Lehman Brothers.

A key source of volatility of consumer prices in 2008 was the behavior of crude oil prices (not shown). The price for West Texas Intermediate rose from \$98 per barrel in January 2008 to a peak of \$140 per barrel in July 2008. By the end of 2008, it had fallen to only \$33 per barrel.

Turning to interest rates, Figure 2.6 shows the behavior of the nominal and ex post real federal funds rate, constructed from the monthly rate of inflation of PCE. The nominal series shows cuts every month from mid-2007 to May 2008, a leveling off from May through August, and then cuts until the zero lower bound was reached. The combination of the cuts and the higher rates of inflation result in negative real interest rates in May through July.³

To summarize, these graphs reveal several key aspects of 2008. First, the rebate was large relative to aggregate disposable income. Second, most of the rise in nominal consumption in the first half of 2008 was due to inflation. Real consumption expenditures show a bounce from February to the peak in May 2008, the month with the largest rebate payments, but the magnitude

³If we instead use the PCE price deflator excluding food and energy, the real interest rates are still negative in May through July, but are between 0 and minus 1 percent.

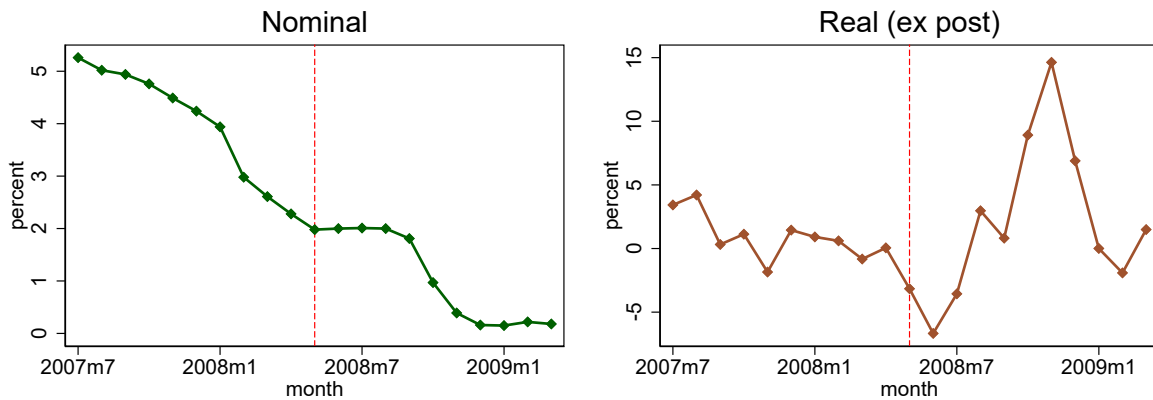


Figure 2.6. Federal Funds Rate

Source: FRED, based on Federal Reserve Board of Governors. The ex post real interest rate is constructed using the annualized monthly rate of inflation in PCE.

is modest. Third, consumer expenditure prices were volatile during 2008. Oil prices and the PCE deflator hit a peak in July and then fell. Fourth, the Fed paused the downward trajectory of the funds rate near the end of May; however the ex post real rate turned negative in Summer 2008 because of the behavior of inflation.

2.2.2 Alternative Measures of Consumption Expenditures

In this section, we show alternative measures of consumption expenditures. The motivation is twofold. First, because the monthly NIPA consumption data are based on combining and smoothing various data sources, we want to provide supplemental evidence that the patterns we showed in consumption expenditures in the last section are not due to smoothing procedures. Second, since the micro estimates suggest that a large portion of the rebate was spent on motor vehicles, it is useful to look at the behavior of aggregate spending on motor vehicles.

We first compare the NIPA measures of personal consumption expenditures (PCE) on goods to two other series: the Census series on retail sales of goods and our own constructed series based on the CEX data that is the basis for the micro estimates. As described by Wilcox (1992), government statisticians use retail sales as an input to monthly consumption, but then

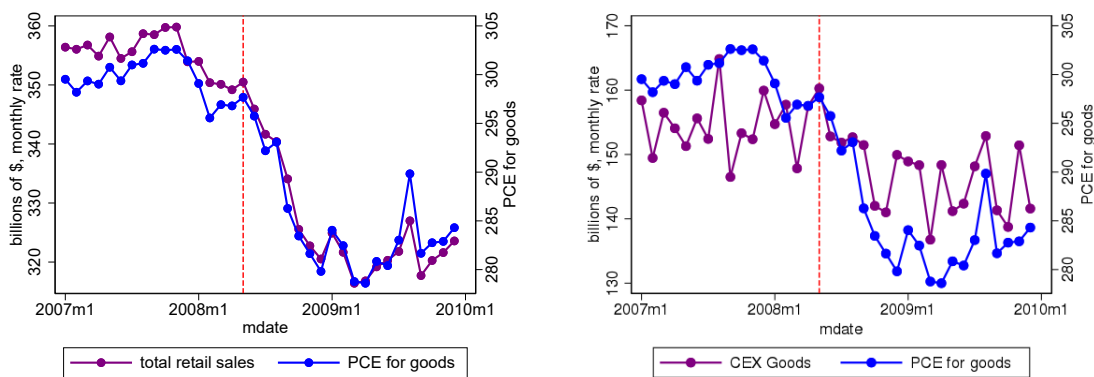


Figure 2.7. Comparison of PCE to Retail and CEX

Source. PCE (Personal Consumption Expenditures) from BEA; Retail sales from Census; Authors' aggregation from CEX. Vertical red dashed line indicates May 2008.

make a number of adjustments. To make sure those adjustments are not smoothing out jumps in consumption due to the rebate, we examine the key underlying series as well as our constructed alternative from the CEX. For all series, we use the PCE goods deflator to create real spending series.

Figure 2.7 shows the comparisons from 2007 through 2009. Consider first the left side graph, which compares PCE on goods to retail sales. The movements in the two series match up very well over the two years. Both show a slight blip up in May 2008, with the retail series showing a more muted blip. Thus, it is unlikely that BEA smoothing of retail sales would account for the consumption pattern.

The right-hand side graph compares PCE on goods to our aggregates of household spending on goods using CEX micro data. The CEX aggregate is much noisier than either the PCE data or the retail sales data. The CEX series falls from February to March, recovers in April, and then declines in May and June. These movements look similar to those in other months, suggesting more noise than information. We conclude that the PCE data is not smoothing out a large jump in consumption when the rebates are distributed.

Finally, we consider detailed data on new motor vehicle expenditures since expenditures on motor vehicles and parts constitute the main part of the high MPC estimated by Parker et al.




Figure 2.8. New Motor Vehicle Sales to Consumers
Source. BEA.

(2013). Another advantage of focusing on motor vehicles is the very high quality of the data. Figure 2.8 shows sales of new motor vehicles to consumers, measured as units on the left-hand side and as billions of dollars on the right-hand side.

Both the unit measure and the dollar measure of sales follow a downward trend from 2007 to early 2009. The unit sales measure shows a small blip in May 2008. This small blip contrasts with the huge spike that occurs later in August 2009 in response to the cash-for-clunkers program. As Sahm et al. (2012) accounting exercise demonstrates, the high MPC estimated by Parker et al. (2013) implies that the bulk of all sales of new motor vehicles in spring and summer 2008 were induced by rebate.

2.2.3 Forecasts of Consumption in 2008

In this section, we present both contemporary forecasts by professional forecasters and our own forecasts that incorporate some of the negative events that occurred in 2008.

Contemporary Forecasts

As discussed in the narrative section above, the employment report released on January 4, 2008 led policymakers and forecasters to raise their probabilities of recession. We begin by

discussing the Goldman Sachs forecast released on January 9, 2008, since they were among the first to predict that the U.S. was already in recession. The Goldman Sachs forecast also contained the following key predictions.⁴ First, the Fed would cut the federal funds rate target from 4.25 to 2.5 by the end of the year, with the first 50 basis point cut at the next FOMC meeting on January 30th. Second, housing prices would decrease 20 to 25 percent below their peak. Third, Congress and the President would pass a temporary tax break as part of a fiscal stimulus plan later in the year.

They forecasted no change in real consumption expenditures (PCE) in 2008Q1, a decrease of 0.125 percent (not annualized) in each of 2008Q2 and 2008Q3, and a 0.25 percent increase in 2008Q4. Thus, they forecasted actual declines in real consumption expenditures, but they were tiny in magnitude. Similarly, contemporary forecasts from the Federal Reserve Board Staff (Greenbooks) and the Survey of Professional Forecasters also did not predict large drops of consumption in summer 2008. Most forecasters predicted an increase in real consumption and even the most pessimistic forecaster from the Survey of Professional forecasters only predicted a small decrease in consumption in summer 2008. We show all these forecasts alongside actual values in figure 2.9.⁵

Our 2008 Consumption Forecasts

In the last section, we showed that even the more pessimistic forecasts did not predict a significant U-shape or V-shape of real consumption between the second and third quarters of 2008. However,, the forecasts in January 2008 did not foresee the rapid run-up in oil prices in spring and summer or the failure of Lehman Brothers in September, both of which could have affected consumption. Thus, we construct our own forecasts that factor those negative events in to create more pessimistic forecasts to compare to our counterfactuals.

⁴This summary is based on contemporaneous news accounts, such as the CNN Money article "Recession may already be here," January 10, 2008.

⁵In each case, we select the last survey prior to the passage of the Economic Stimulus Act of 2008 since afterward forecasters would include the rebate response as part of their forecast. The January Greenbook actually does incorporate the tax rebates in their consumption forecasts, however, they predict that the rebates will be

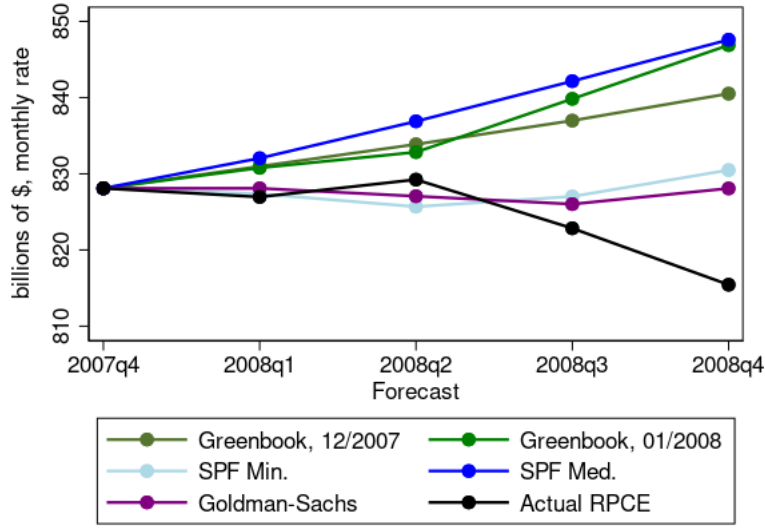


Figure 2.9. Contemporary Real Consumption Forecasts

Source. BEA data, Federal Reserve Board, Survey of Professional Forecasters. All forecasts normalized to monthly real consumption in 2007Q4.

Our forecasting model is a simple monthly frequency time series model with the following endogenous variables: log real consumption, log real disposable income, log consumption deflator, and the Gilchrist and Zakrajšek (2012) excess bond premium. We also include a dummy variable for recession, log real oil prices, and a dummy variable for the Lehman Brothers bankruptcy in September 2008. We explored the addition of a number of other variables, such as consumer confidence but they did not noticeably change the forecasts and/or were not statistically significant. We use six lags of each variable, except for the Lehman Brothers dummy variable where we use current and two lags. We include current values of the recession dummy, oil prices, and the excess bond premium in the equations for the endogenous variables. We estimate the model on data from 1984m1 - 2019m12 and forecast dynamically starting in January 2008. We start the estimation period in 1984 because the effects of oil prices on consumption expenditures changed significantly post-1984 (e.g. Edelstein and Kilian (2009)).

We produce four forecasts by varying our assumptions on the exogeneity of oil prices received in the second half of 2008, not in the second quarter when most of them were received.

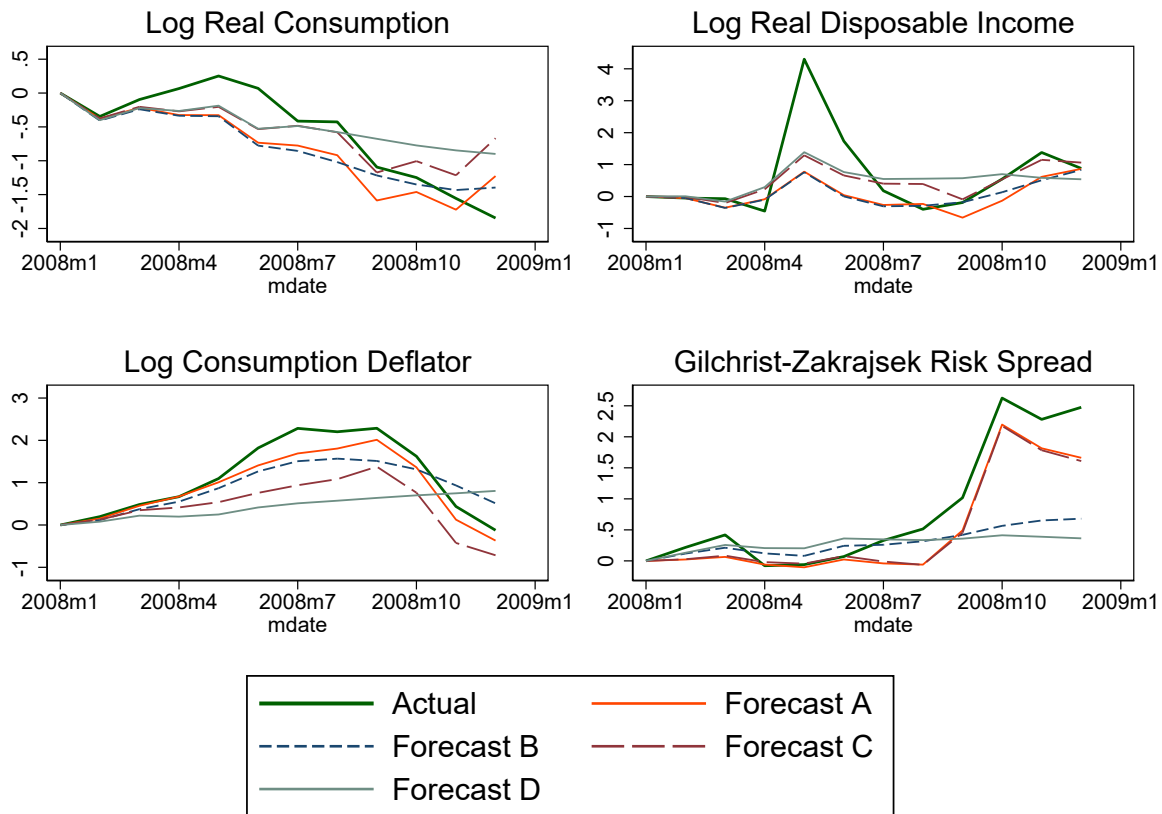


Figure 2.10. Forecasts from Our Four Models

Forecasts are based on information through January 2008, with exception of models in which oil prices are exogenous and Lehman Brothers dummies are included. Real oil prices are assumed to be exogenous in Models A and B; Lehman Brothers bankruptcy dummy variables are included in Models A and C.

and whether Lehman Brothers went bankrupt. The most pessimistic forecasts are those in which oil prices are assumed to follow their actual path exogenously and in which the Lehman Brothers bankruptcy dummy variable is included in the forecasting equation. We keep the current and lagged recession dummy variable in all forecasts; if we omit them, the forecasts are substantially more optimistic.

Figure 2.10 shows the forecasts for the four endogenous variables in each of the four models. The most pessimistic forecast (Forecast A) assumes both exogenous oil prices and that Lehman Brothers went bankrupt in September 2008. The reason that allowing oil prices to respond exogenously leads to a more pessimistic forecast is that the alternative model in which

oil prices respond *endogenously* does not predict a rise in spring and summer 2008, but instead predicts a gentle drift down until they plummet after the bankruptcy of Lehman Brothers in September 2008. None of the forecasts hints at a V-shape path of consumption in 2008.

2.3 Macro Counterfactuals

In this section, we begin by constructing a medium-scale New Keynesian (NK) model that allows us to generate counterfactual paths of consumption expenditures that include general equilibrium feedbacks. We then simulate the response of consumer expenditures to rebates and apply the results to actual consumer expenditures to create counterfactual paths had there been no rebates.

2.3.1 Two-Good, Two-Agent New Keynesian Model with Hand-to-Mouth Consumers and Durable Goods

We construct a two-good, two-agent New Keynesian (TG-TANK) model, which features both nondurable and durable goods and both optimizing and hand-to-mouth agents. Most elements of our model are standard for a medium-scale New Keynesian model. In particular, it builds on the model analyzed by Ramey (2021b), which is an extension of Galí et al. (2007) fiscal NK model. The main addition to the model is a durable consumption good, which we interpret as motor vehicles. This part of the model builds on McKay and Wieland (2021) recent analysis of durable goods expenditures.

We begin by describing the household's problem in more detail, since it is less standard than the other parts of the model. We then briefly summarize the other features, and refer interested readers to the appendix for more details.

Optimizing Households

A measure $1 - \gamma$ of *ex-ante* identical households maximizes utility subject to their budget

constraints. The utility function for these optimizing households is:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{(C_t^o)^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} + \psi \frac{(D_t^o)^{1-\frac{1}{\sigma^d}}}{1-\frac{1}{\sigma^d}} - \nu \frac{(H_t^o)^{1+\phi}}{1+\phi} \right]$$

where C_t^o is nondurable consumption, D_t^o is the durable stock, and H_t^o is hours worked. The household budget constraint is

$$A_t^o = \frac{R_{t-1}}{\Pi_t} A_{t-1}^o - C_t^o + W_t H_t^o - X_t^o - \eta D_t^o - T_t^o + \text{Profits}_t^k + \text{Profits}_t^s$$

where R_t is the gross nominal interest rate, Π_t is the gross inflation rate measured in nondurable goods prices, A_t^o are holdings of the nominal bond, W_t is the real wage, X_t^o is durable expenditure denominated in nondurable goods, ηD_t^o are operating expenditures for the durable good (e.g., gasoline), T_t^o are net taxes (i.e. taxes less transfers), Profits_t^k are profits of the capital good producing firms, and Profits_t^s are profits of the sticky-price firms, which produce nondurable goods.

Durables follow a standard accumulation equation

$$D_t^o = (1 - \delta^d) D_{t-1}^o + \frac{X_t^o}{p_t^d}$$

where δ^d is the depreciation rate of household durables and p_t^d is the relative price of durable goods.

Optimizing households pick an optimal plan $\{C_t^o, D_t^o, A_t^o, X_t^o\}_{t=0}^{\infty}$ to maximize utility. Labor supply is not chosen by the household, but instead by a union as discussed below. The

first order conditions for the household problem are:

$$\begin{aligned}\lambda_t &= (C_t^o)^{-\frac{1}{\sigma}} \\ \lambda_t &= \beta \frac{R_t}{\Pi_{t+1}} \lambda_{t+1} \\ p_t^d \lambda_t &= \mu_t \\ \mu_t &= -v \lambda_t + \beta(1 - \delta^d) \mu_{t+1} + \psi(D_t^o)^{-\frac{1}{\sigma^d}}\end{aligned}$$

where λ is the Lagrange multiplier on the household budget constraint and μ is the Lagrange multiplier on the durable accumulation equation.

Hand-to-Mouth Households

In order for lump-sum transfers to have general equilibrium effects, we require non-Ricardian households. We adopt Galí et al. (2007) assumption that a certain fraction γ of consumers neither borrow nor save and simply consume all of their current income,

$$\begin{aligned}A_t^m &= 0 \\ C_t^m + \eta D_t^m + X_t^m &= W_t H_t^m - T_t^m\end{aligned}$$

where variables superscripted by m denote the hand-to-mouth household.

We also assume that in steady state, hand-to-mouth households consume the same relative amount of durable and nondurable services,

$$\frac{C^m}{X^m} = \frac{C^o}{X^o}$$

Finally, we directly specify dynamic marginal propensities to consume for nondurable

and durable expenditures to match both the allocation across goods and any lagged effects implied by the micro MPC estimates,

$$\begin{aligned}
C_t^m - C^m + \eta(D_t^m - D^m) &= \sum_{l=0}^L mpc_l [W_{t-l} H_{t-l}^m - T_{t-l}^m - (WH^m - T^m)] \\
X_t^m - X^m &= \sum_{l=0}^L mpx_l [W_{t-l} H_{t-l}^m - T_{t-l}^m - (WH^m - T^m)] \\
1 &= \sum_{l=0}^L (mpc_l + mpx_l) \\
mpx_l &= \frac{\theta}{1 - \theta} mpc_l, \quad \forall l = 0, \dots, L
\end{aligned}$$

where mpc_l is the marginal propensity to spend on nondurable goods today out of income l periods ago, and mpx_l is the marginal propensity to spend on durable goods today out of income l periods ago.

Durable Goods Production

Durable goods are produced competitively using nondurables N_t as inputs,

$$\frac{X_{it}}{p_t^d} = N_{it} \left(\frac{X_t}{\bar{X}} \frac{1}{p_t^d} \right)^{-\zeta}$$

where $\frac{X_{it}}{p_t^d}$ is the real production of durable goods by firm i and ζ is a negative production externality. ζ could alternatively represent a fixed factor of production as in McKay and Wieland (2021). We model it as a production externality because this yields zero profits in durable production.

Real profits from the sale of durable goods are given by

$$\max_{N_{it}} (X_{it} - N_{it}) = \max_{N_{it}} \left[p_t^d N_{it} \left(\frac{X_t}{\bar{X}} \frac{1}{p_t^d} \right)^{-\zeta} - N_{it} \right]$$

Profit maximization yields an upward sloping supply curve,

$$p_t^d = \left(\frac{X_t}{\bar{X}} \right)^{\frac{\zeta}{1+\zeta}}$$

where \bar{X} is steady state durable expenditure, so the steady state relative durable price is normalized to 1. Since durable expenditure is denominated in units of nondurable consumption, the supply elasticity of real durable goods is given by $\frac{1}{\zeta}$.

Summary of the Model's Other Features

We summarize the other features of the model only briefly since they are standard. The market for nondurable goods features sticky prices and sticky wages and noncompetitive product and labor markets. Intermediate goods firms are monopolistically competitive and face a Calvo-style (1983) adjustment cost on prices. In labor markets, households mark up wages over the marginal rate of substitution and face Calvo-type (1983) adjustment costs. The result is that short-run employment fluctuations are driven more by labor demand than labor supply. Firms face an adjustment cost on capital investment. However, they can vary their utilization of capital, so capital services are more cyclical than the capital stock. The result is more elastic output supply since it mutes the diminishing returns to labor and prevents real marginal cost from increasing much when output rises. The monetary rule is inertial, with a coefficient of 1.5 on the inflation gap and 1 on the output gap. Lump-sum taxes respond to the deviation of government debt from its steady-state values but with a lag of one year. A more complete description with equations is provided in the appendix.

Table 2.1. Baseline Calibration of the Model

Parameter	Value	Description
β	0.997	Subjective discount factor
ψ	1.435	Weight on durable service flow
σ	0.5	IES for nondurable consumption
σ_d	varies	Utility curvature on durable service flow
η	0.018	Durable operating cost
ν	70.956	Weight on disutility of labor
ϕ	1	Inverse of the Frisch elasticity of labor supply
γ	varies	Fraction of Hand-to-Mouth consumers
θ	0.83	Hand-to-Mouth fraction of MPC spent on durables
δ_d	0.015	Depreciation of durable consumption goods
α	0.36	Exponent on private capital in production function
δ	0.005	Depreciation of private capital
κ	40	Investment adjustment cost parameter
δ_1	0.008	Parameter on linear term of capital utilization cost
δ_2	0.017	Parameter on quadratic term of capital utilization cost
ζ	0	Inverse supply elasticity of durable goods
μ_p, μ_W	1.2	Steady-state price markup, wage markup
θ_p, θ_W	0.917	Calvo parameters on price and wage adjustment
$\varepsilon_p, \varepsilon_W$	6.0	Elasticities of substitution between types of goods and types of labor
gy	0.175	Steady-state share of total govt spending to GDP
ϕ_b	0.1	Debt feedback coefficient in fiscal rule
ρ_r	0.947	Monetary policy interest rate smoothing
ϕ_π	1.5	Monetary policy response to inflation
ϕ_{gap}	0.083	Monetary policy response to the output gap

2.3.2 Calibration

The calibrated parameters with their descriptions are shown in Table 2.1. Note that the model is calibrated to a monthly frequency. In addition to the calibrations shown in the table, we calibrate the steady-state transfers by type of household so that hand-to-mouth and life-cycle permanent income households consume the same amount in the steady state. We also calibrate the steady-state ratio of government purchases to GDP to equal 0.175 to match the U.S. economy average. The durable goods parameters are chosen to match the share of motor vehicle spending in PCE and its depreciation rate in the fixed asset table. Operating costs are based on PCE expenditures on motor vehicle fuels, lubricants, and fluids. The appendix shows more details of

the model.

The timing of spending by hand-to-mouth households is important for constructing the counterfactual path of consumption. We assume that the hand-to-mouth households respond to a shock to their disposable income by spreading their spending over three months. Estimates from Broda and Parker (2014) using higher-frequency Nielsen data on nondurable expenditures suggest that two-thirds of expenditure occurs in the month of the rebate, and one-sixth each of the following two months. In our own investigation using CEX data, we find no evidence of additional expenditure after three months.⁶ Unfortunately, the CEX does not lend itself to estimate monthly expenditure patterns as most household report expenditures divided equally across the three months within an interview. One exception to this limitation is reported car expenditure, which more precisely identifies the month of purchase. Appendix table B1 shows that the car expenditure response occurs in the three months around the rebate. We conservatively choose an equal spread of expenditure since this minimizes the extent of V-shapes in our counterfactuals and is thus more consistent with larger MPCs.

A key distinction in both the estimates and in our model is the allocation of spending between nondurable goods and motor vehicles. We assume that hand-to-mouth households allocate 83% of their expenditure towards motor vehicles. This is based on our preferred estimated MPCs after implementing the Borusyak et al. (2022) method in the next section of this paper. The estimate for nondurable spending is 0.057 (table 2.8, panel B column 1) and for cars is 0.3 (table 2.7, panel B column 1).

We simulate several versions of the model, across a range of fractions of households who are hand to mouth. We set values for γ , and thus the overall quarterly MPC, equal to 0.3, 0.5, and 0.7. The lower value, 0.3, reflects our preferred estimate after implementing the Borusyak et al. (2022) method (table 2.5, panel B column 1). The other two values, 0.5 and 0.7, are the lower bound and mid-point for the MPC reported in Parker et al. (2013).

The supply and demand elasticities for durable goods are two important parameters

⁶See the appendix table

for the general equilibrium outcomes of the model. We set the durable good supply elasticity $\zeta^{-1} = \infty$, implying perfectly elastic supply of durable goods. We later allow for a less elastic supply of durable goods.

We calibrate the curvature of durable utility σ^d to match a motor vehicle demand elasticity of -0.87 estimated by McCarthy (1996). For example, when the fraction of hand-to-mouth households is $\gamma = 0.3$, targeting this value of the elasticity yields $\sigma^d = 0.25$.

2.3.3 Macro Counterfactuals

With the model constructed and calibrated, we now compute counterfactual paths of consumption that take into account the full dynamic general equilibrium effects. We start the economy in steady state in January 2008, and assume that households do not anticipate in advance the equilibrium path of prices resulting from the rebate until after the first rebate payments are made in April.⁷ We feed a path of rebate shocks into the model that matches the relative size and timing of the actual rebate shown in figure 2.2.

We use first-order perturbation methods to solve for the general equilibrium impulse responses of the variables to the path of rebates. We then construct macro-counterfactuals by subtracting the model-implied impulse response functions for consumer expenditures from the observed consumer expenditure data.⁸

Figure 2.11 plots counterfactuals based on both the micro MPCs, excluding any general equilibrium effects, and on the GE-MPCs, which incorporate full dynamic general equilibrium feedbacks. The figure shows the results for both total consumer expenditure (real PCE) and motor vehicle expenditure.⁹ The micro counterfactual graphs in the left column are the analogs to the Sahm et al. (2012) counterfactual for motor vehicles, except that we assume that expenditure is equally spread over three months rather than over two months and we assume a greater fraction

⁷Without this assumption, optimizing households would foresee the future rise in motor vehicle prices and would increase their purchases immediately.

⁸Because the model is linearized, the counterfactuals for the tax rebate would be identical if we also fed the model with other shocks that hit the economy at the time.

⁹Appendix Figure B1 shows the counterfactual for nominal PCE and motor vehicle expenditure.

of the rebate is spent on motor vehicles. The figures show prominent, and we have argued implausible, V-shapes for both total expenditure and motor vehicle expenditure, even for micro MPCs for total consumption expenditures as low as 0.3.

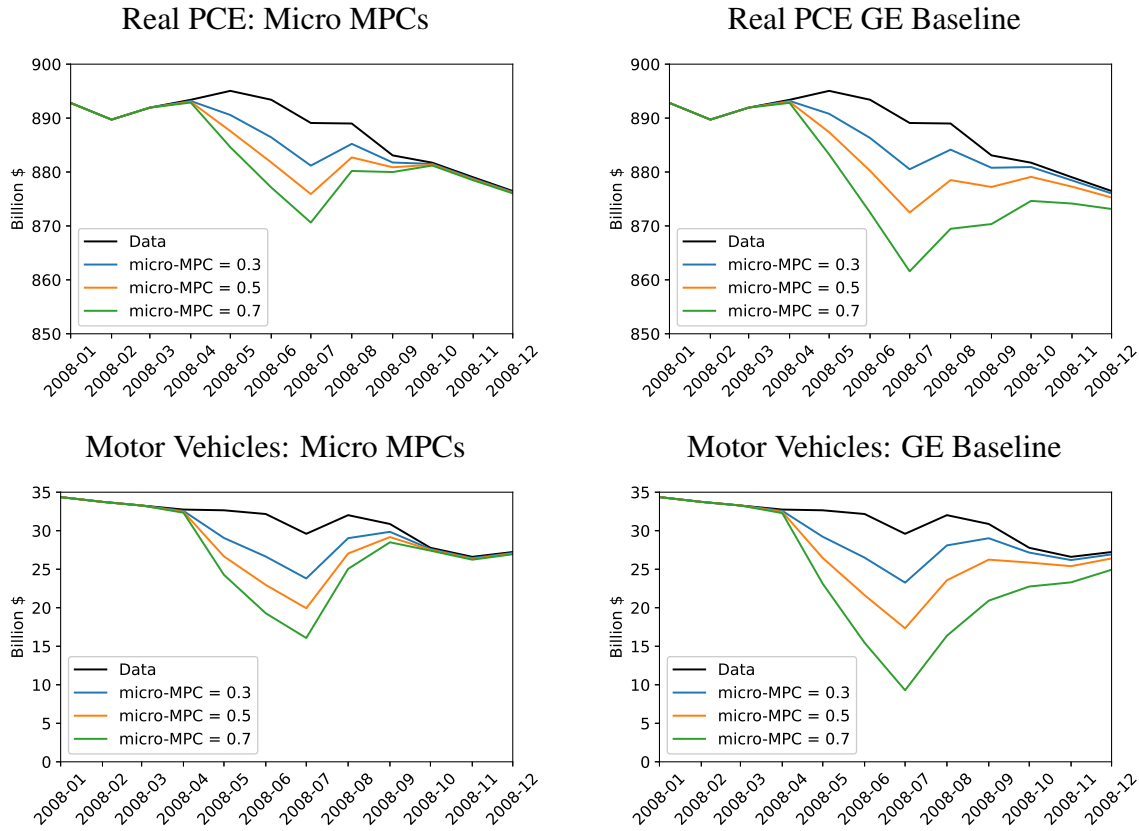


Figure 2.11. Counterfactual Real Consumption Expenditures: Baseline Model

Notes. Based on NK model simulations and actual data on rebates and consumption. The micro MPC value refers to the MPC for total consumption.

The graphs in the right column of Figure 2.11 plot the corresponding counterfactuals in general equilibrium. In this standard New Keynesian model, the general equilibrium forces amplify the effects, particularly as the micro MPCs become larger, so the counterfactual paths become even more V-shaped. For example, for a total micro MPC of 0.7, the implied counterfactual path of motor vehicles falls to \$5 billion in the general equilibrium experiment rather than \$13 billion in the experiment that neglects general equilibrium effects.

To quantify the total change in consumption following the rebate, we compute micro

MPCs and GE-MPCs over a twelve month period in response to the rebate shock.¹⁰ Table 2.2 shows the correspondence between the micro MPCs (which equal the fraction of hand-to-mouth households) and general equilibrium MPCs. When the micro MPC is 0.3, the amplification is modest so that the GE-MPC for total consumption is only 24 percent higher (0.37) than the micro MPC. In contrast, when the micro MPC is 0.7, the GE-MPC is double the micro MPC.

Table 2.2. General Equilibrium Marginal Propensity to Consume: Baseline Model

PCE		Motor vehicles		Nondurable goods	
micro	GE	micro	GE	micro	GE
0.3	0.371	0.249	0.307	0.051	0.065
0.5	0.765	0.415	0.634	0.085	0.131
0.7	1.406	0.581	1.169	0.119	0.237

How do we reconcile the high micro MPCs with these implausible counterfactuals? To answer this question, we now explore modifications of the standard New Keynesian model that dampen rather than amplify the micro MPCs. In the next section, we re-examine the micro MPC estimates.

There are a number of ways to introduce dampening forces in general equilibrium that might help solve the puzzle of the implausible counterfactual. Possibilities include less accommodative monetary policy or lower elasticity of aggregate output.¹¹ We instead choose the most straightforward way to do this in our two-good model, which is to make the supply of durable goods more elastic. Our baseline calibration assumes a perfectly elastic supply of durable goods, which mimics the results one would obtain in a one-good model.¹² We thus calibrate the elastic supply of durable goods more realistically, by changing the supply elasticity of durable goods from $\zeta^{-1} = \infty$ to $\zeta^{-1} = 5$ which is midway between the elasticities reported in House and Shapiro (2008) and Goolsbee (1998).

¹⁰GE-MPCs are computed in terms of real quantities.

¹¹The elasticity of aggregate output will be lower if prices and wages are more flexible, the labor supply elasticity is lower, or there is less scope for varying the utilization of capital.

¹²Recall that in our model durable goods are produced competitively using nondurables as inputs, so a perfectly elastic supply means that the two goods are perfect substitutes in production.

Figure 2.12 plots the corresponding counterfactuals for the revised model. The left column reports the same micro counterfactuals (which exclude general equilibrium effects) from the previous graph for comparison purposes and the right column reports the new general equilibrium counterfactuals based on less elastic durable goods supply. For total PCE we no longer see evidence of V-shapes in the general equilibrium counterfactual. This change occurs because the general equilibrium response of motor vehicle expenditure to a tax rebate is much less than implied by the micro MPCs. With our preferred micro MPC of 0.3, real motor vehicle spending in general equilibrium falls from \$33 billion in March 2008 to \$28 billion July 2008, rather than from \$33 billion to \$22 billion based on the micro-MPCs. For higher micro MPCs these differences are even larger.

Our preferred micro MPC estimate also shows a continuous decline of the counterfactual consumer expenditure path for both total expenditure and motor vehicles. In particular, this estimate implies that motor vehicles decline further as Lehman Brother fails in September 2008. In contrast, with a micro MPC of 0.5 or 0.7, motor vehicle expenditure in July 2008 is at or below the level of spending when Lehman Brothers fails.

Table 2.3. General Equilibrium Marginal Propensity to Consume: Model with less elastic Durable Supply

micro	PCE	Motor vehicles		Nondurable goods	
	GE	micro	GE	micro	GE
0.3	0.155	0.249	0.12	0.051	0.035
0.5	0.36	0.415	0.286	0.085	0.074
0.7	0.831	0.581	0.67	0.119	0.161

Table 2.3 shows the correspondence between the micro MPCs and the GE-MPCs. When the micro MPC is 0.3, the GE-MPC is only half as large, 0.155. In this case, the general equilibrium forces of the model dampen the effect of the rebate on consumer expenditure. For a micro MPC of 0.5, this dampening is smaller and the GE-MPC is 0.36. For a micro MPC of 0.7, general equilibrium amplifies the initial partial equilibrium spending response resulting in a

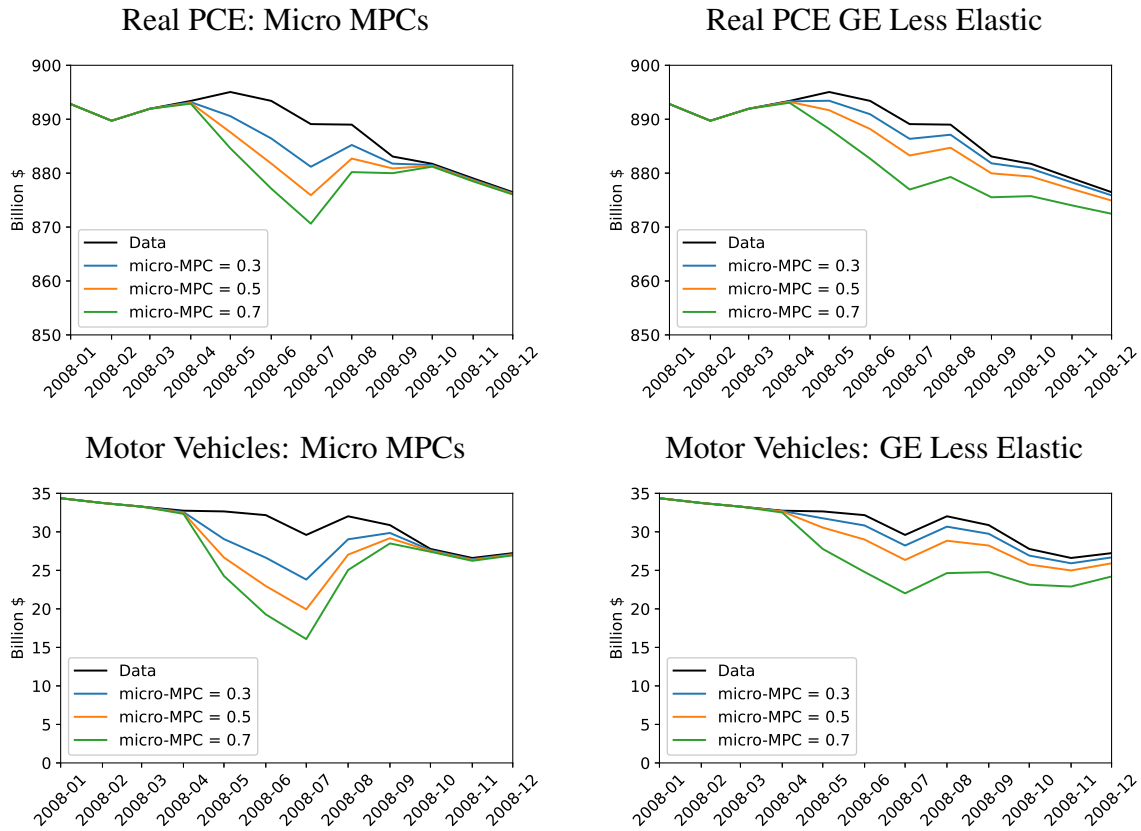


Figure 2.12. Counterfactual Real Consumption Expenditures: Less Elastic Durable Supply Model

Notes. Based on NK model simulations and actual data on rebates and consumption. The micro MPC value refers to the MPC for total consumption.

GE-MPC of 0.83. The general equilibrium spending response is non-linear in the micro MPC primarily because the Keynesian multiplier is also non-linear. For example, for a micro MPC of 0.3, the Keynesian multiplier is only 0.4; for a micro MPC of 0.7, the Keynesian multiplier is 2.3.¹³

The next four columns decompose the MPCs into durable expenditure (motor vehicles) and nondurable expenditure. By construction, the durable micro MPC accounts for 83% of the total expenditure micro MPC. The GE-MPCs show that the dampening in general equilibrium is concentrated in durable expenditure. For example, when the micro MPC on durables is 0.25,

¹³The simple Keynesian multiplier on rebates is $mpc/(1-mpc)$.

then the GE-MPC is less than half that magnitude. In contrast, the MPC on nondurables is only dampened by one third in general equilibrium.

The general equilibrium dampening of the consumption responses stems from the rise in relative durable goods prices. Optimizing households intertemporally substitute away from durable goods because their price is temporarily high; however, there is only a small amount of intratemporal substitution toward nondurable goods. Hand-to-mouth households also reduce their real expenditures on durable goods, but in their case, it is because their MPCs are fixed in nominal terms so the rise in relative prices of durable goods eats up part of their spending. Aggregate nondurable expenditure excluding operating costs is essentially invariant to changes in the relative price.¹⁴

These results have important implications for models with homogenous goods. While many models in the heterogenous agent literature are calibrated to match micro MPCs around 0.3, these models typically include only nondurable spending and therefore abstract from the stronger general equilibrium forces on durable expenditure.¹⁵ Table 2.4 shows that a model that abstracts from durable goods features amplification in general equilibrium across the range of micro MPCs we consider.¹⁶ For instance, a micro MPC of 0.3 translates into a GE-MPC of 0.41 in the model without durable goods. In our model with durable goods the GE-MPC is less than half as large as the micro MPC. This shows that it is not only important to match an overall micro MPC for consumer spending, but also its composition across nondurables and durables and their heterogeneous general equilibrium effects.

¹⁴Nondurable expenditure excluding operating costs is also less responsive to changes in the real interest rate than durable expenditure. However, this difference is less important in our simulation because the change in the real interest rate peaks at only 10 basis points when the micro MPC is 0.3. This small response of the real interest rate reflects short-lived rebate, that prices are sticky, and that the nominal interest rates are inertial.

¹⁵Notable exceptions include Berger and Vavra (2015), McKay and Wieland (2021), and McKay and Wieland (2022).

¹⁶In this model we set the weight on the utility of durables stock $\psi = 0$, durable operating cost $\eta = 1$, and fraction of MPC that is allocated to durables $\theta = 0$.

Table 2.4. General Equilibrium Marginal Propensity to Consume: Model without Durable Goods

PCE		Motor vehicles		Nondurable goods	
micro	GE	micro	GE	micro	GE
0.3	0.415	−∅	−∅	0.3	0.415
0.5	0.89	−∅	−∅	0.5	0.89
0.7	1.767	−∅	0.0	0.7	1.767

2.4 The Micro MPC Estimates

We now reconsider the micro MPC estimates. We first summarize the latest developments in the estimation of treatment effects in the type of model used by Parker et al. (2013) on the CEX data. We then replicate the Parker et al. (2013) results using our version of the data and their methods and then apply some of the recently-developed econometric methods to generate new estimates of the micro MPCs. Our new estimates imply lower micro MPCs.

2.4.1 Estimation Strategy

The most widely cited micro MPC estimates, which range from 0.5 to 0.9, come from Parker et al. (2013). In a case of entrepreneurial data collection, the authors worked with the U.S. Bureau of Labor Statistics to add a question about the 2008 Tax Rebate receipt to the monthly Consumer Expenditure Survey (CEX). Since the CEX is a rotating panel survey of household expenditure, this allowed the authors to analyse consumption expenditure alongside rebate receipt in an already established survey. Furthermore, since rebate checks were sent to households based on the last two-digits of their social security number, the timing of treatment (i.e. distribution of the rebate) was effectively random.

Parker et al. (2013) leverage the variation in treatment time (i.e., the month in which the household received the rebate) and in some cases the treatment size (i.e. the dollar value of the rebate check) to estimate the causal impact of receiving a rebate on household spending using a standard difference-in-differences (DID) event-study methodology. We will focus on their specifications that leverage only the treatment timing, since the recently-developed method

that we use does not allow for continuous treatment variables. For this specification, Parker et al. (2013) estimate the following regression,

$$C_{i,t+1} - C_{i,t} = \sum_s \beta_{0s} month_{s,i} + \beta_1' \mathbf{X}_{i,t} + \beta_2 I(ESP_{i,t+1}) + u_{i,t+1} \quad (2.4.1)$$

where t indexes the interview (performed once every three months), and i indexes individual households. The regression includes fixed effects for each month ($month_{s,i}$), household controls for age and change in household size $X_{i,t}$, and the main variable of interest, $I(ESP)$, which is a dummy variable equal to one if the household received a rebate, i.e., an Economic Stimulus Payment (ESP).

In the last few years, the literature on staggered event-studies and two-way fixed effect models has made significant progress, first by uncovering problems with standard OLS estimators, and second by developing new estimators appropriate for this context (see e.g., Borusyak and Jaravel 2017, De Chaisemartin and d'Haultfoeuille 2020, Sun and Abraham 2020, Borusyak et al. 2022). The problems arise in the weights that are used by standard methods. The standard OLS estimators implicitly adopt the assumption that the treatment effect β_2 is homogeneous in the population. To maximize efficiency in this context, OLS assigns a large weight (relative to population size) to certain treatment effects and a negative weight to other treatment effects (see e.g., De Chaisemartin and d'Haultfoeuille 2015, Sun and Abraham 2020, Borusyak et al. 2022). But this weighting scheme is inappropriate when treatment effects are heterogeneous and the object of interest is the average effect of treatment on the treated (ATT) in the population.¹⁷

Our goal is to estimate the average MPC in the population of households treated by the rebate. For this purpose we adopt the method in Borusyak et al. (2022). Their method consists of imputing a counterfactual spending path based on untreated and non-yet-treated households, and then aggregating the implied treatment effects among the treated population using equal weights.

¹⁷Misra and Surico (2014) were the first to note the heterogeneity across households in the responses to the rebates in 2001 and 2008 and used quantile regression methods to allow for heterogeneity.

The identifying assumptions are that there are no anticipation effects and that the untreated households are on parallel trends with the treated households.¹⁸

Both Borusyak and Jaravel (2017) and Borusyak et al. (2022) apply versions of the imputation estimator to Broda and Parker (2014) estimates of MPCs using the Nielsen data. In both cases, they find that the imputation method produces MPC estimates that are half those estimated by Broda and Parker (2014). Thus, our application to the CEX data used by Parker et al. (2013) complements the results of these two studies.

We estimate the following regression on the sample of untreated observations, which consists of all observations on households that never received a rebate and observations on households prior to their receiving a rebate, i.e. the "untreated." The estimating equation is:

$$Y_{i,t+1} \equiv C_{i,t+1} - C_{i,t} = \sum_s \beta_{0s} \text{month}_{s,i} + \beta_1' \mathbf{X}_{i,t} + \tilde{u}_{i,t+1}, \quad \forall (i,t+1) \in \{\text{Untreated}\}$$

Because these observations are untreated, this equation omits $ESP_{i,t+1}$ in contrast to equation (2.4.1). We use the estimated coefficients from this equation to "impute" the change in spending for all observations as if they had never received a rebate check as:

$$Y_{i,t+1}(0) = \sum_s \hat{\beta}_{0s} \text{month}_{s,i} + \hat{\beta}_1' \mathbf{X}_{i,t}, \quad \forall (i,t+1) \in \{\text{Full Sample}\}$$

where $Y_{i,t+1}(0)$ is the imputed change in expenditure of household i if it is never treated. We

¹⁸We adopt the weaker parallel trends rather than the random treatment assignment for the following reasons: (1) The actual rebate timing is not fully random because households received the rebate sooner if they filed taxes via EFT. (2) The reported rebate dates appear non-random as households are more likely to report receiving a rebate in the month before the interview compared to the previous two months. (3) We prefer to use the never-treated group as a control group because the OLS weighting problems are more severe when no never-treated group exists.

then create the difference between the actual change in expenditures and the imputed change as:

$$\tau_{i,t+1} = Y_{i,t+1} - Y_{i,t+1}(0), \quad \forall (i, t+1) \in \{\text{Treated in } t+1\}.$$

The average treatment effect of receiving the rebate on spending by households that received a rebate in the last interview period is then just:

$$\tau = \sum_{i,t+1 \in I(ESP_{i,t+1})=1} \omega_{i,t+1} \tau_{i,t+1},$$

where the weights $\omega_{i,t+1}$ are chosen so that τ is the sample average given the CEX survey weights (ATT).¹⁹

2.4.2 Results

We first report the results for the version of equation (2.4.1) that uses the change in total consumer expenditure as the dependent variable. These results are reported in Table 2.5. Panel A reports the estimates the treatment effects using standard OLS as in equation (2.4.1). Column (1) is a replication of Parker et al. (2011) (the detailed working paper version of (Parker et al. 2013)) estimates in Table 4, column 8. While the samples are not exactly identical,²⁰ the estimates—\$483.2 in our sample, \$494.5 in theirs—are extremely close. We construct an implied MPC by dividing this estimate by an estimate of the rebate received for each household. We obtain the rebate amount estimate by regressing the rebate amount on the rebate indicator and the other control variables in shown in equation (2.4.1). These results are tabulated in Table 2.6. The ratio yields an MPC of $\frac{483.2}{930.5} = 0.519$, very close to the value of 0.523 reported in Parker et al.

¹⁹We use Borusyak et al. (2022)'s **did_imputation** STATA command to construct point estimates and standard errors.

²⁰We were unable to create the exact same dataset as in (Parker et al. 2013) based on the replication instructions provided by (Johnson et al. 2006, Parker et al. 2011; 2013). But the difference appear to be very small in the vast majority of cases.

Table 2.5. Contemporaneous Household Expenditure Response to Rebate

Panel A: OLS				
	Full Sample		Rebate Only Sample	
	(1)	(2)	(3)	(4)
Rebate Indicator	483.2** (2099)	325.7* (1782)	779.2** (3102)	593.6** (2388)
Implied MPC	0.52	0.35	0.86	0.65
Extra Controls	No	Yes	No	Yes
Observations	17,229	17,229	10,343	10,343

Panel B: DID Imputation				
	Full Sample		Rebate Only Sample	
	(1)	(2)	(3)	(4)
Rebate Indicator	287.0 (2160)	116.2 (1914)	984.4 (6656)	-64.3 (5790)
Implied MPC	0.30	0.12	1.03	-0.07
Extra Controls	No	Yes	No	Yes
Observations	12,499	12,499	5,585	5,585

Notes: The dependent variable is the change in Expenditure from the previous interview. Standard errors, in parentheses, are clustered at the household level. Significance is indicated by: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions include interview (time) fixed effects, as well as household level controls for age, change in number of adults, and change in number of children. Extra controls refer to additional controls for household income decile and lagged total spending. Rebate sample includes only households that receive a rebate at some point during our sample period.

(2011), Table 4, column 16. Column (3) of Table 2.5 repeats the same analysis in the sub-sample of households that report receiving a rebate. Our implied MPC, $\frac{779.2}{905.5} = 0.861$ is again very close to the estimate of 0.866 reported in Parker et al. (2011), Table 4, panel B, column 12.

In columns (2) and (4) of Table 2.5, Panel A, we include additional controls for household income decile and lagged spending. These controls are not included in the original Parker et al. (2011) specifications, but we do find that they reduce the implied MPCs relative to the baseline specifications. This suggests that the rebate timing is not fully orthogonal to household characteristics. Nevertheless, the two-way fixed effects estimates for the MPC remain statistically

Table 2.6. First Stage: Rebate Amount Conditional on Rebate Receipt

Panel A: OLS				
	Full Sample		Rebate Only Sample	
	(1)	(2)	(3)	(4)
Rebate Indicator	930.5*** (10.2)	926.6*** (10.1)	905.5*** (12.9)	907.5*** (12.8)
Extra controls	No	Yes	No	Yes
Observations	17,229	17,229	10,343	10,343

Panel B: DID Imputation				
	Full Sample		Rebate Only Sample	
	(1)	(2)	(3)	(4)
Rebate Indicator	952.4*** (9.62)	952.4*** (9.62)	954.1*** (9.69)	954.1*** (9.69)
Extra controls	No	Yes	No	Yes
Observations	12,499	12,499	5,585	5,585

Notes: The dependent variable is the dollar value of Economic Stimulus Payments (ESP) received by the household. Standard errors, in parentheses, are clustered at the household level. Significance is indicated by: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions include interview (time) fixed effects, as well as household level controls for age, change in number of adults, and change in number of children. Extra controls refer to additional controls for household income decile and lagged total spending. Rebate sample includes only households that receive a rebate at some point during our sample period.

significant, and remain large in the rebate-only sub-sample.

In Panel B of Table 2.5 we instead apply the Borusyak et al. (2022) imputation estimator. Column (1) shows that average rebate spending is only \$287.0, compared to the OLS estimate of \$483.2 in column (5). The implied MPC in column (15) is 0.3. Note that while the point estimate drops by almost half, the standard errors are almost unchanged. This is also the case once we include extra household controls in column (2), which only further depress the estimate for the MPC.

In columns (3) and (4) of Panel B of Table 2.5 we restrict the estimation to the rebate-only sample, which results in much noisier and statistically insignificant estimates. The noisier

estimates reflect the fact that the imputed part of the dependent variable $Y_{i,t+1}(0)$ is calculated using a shrinking sample of not-yet-treated observations. Most households receive their rebate in May or June 2008, which means that a very small number of households are used to calculate the time fixed effects for the imputed dependent variable for the majority of the sample. The higher precision of the estimates in panel A columns (3) and (4) suggests that OLS heavily leverages comparisons with previously treated units. Borusyak et al. (2022) call these “forbidden comparisons” because they may result in negative weighting of treated observations, which in turn yields misleading estimates of the ATT when treatment effects are not homogeneous. For this reason, we believe the full sample estimates yield more reliable estimates of the aggregate MPC.

In Figure 2.13 we decompose the OLS coefficient in column (1) of Panel A and the imputation estimator in column (2) Panel B into weights (top left panel) and their treatment effects (top right panel). The headline coefficients in Table 2.5 are simply the weighted sum of the period treatment effects (see appendix B4 for details). The top right panel shows that the imputation estimator applies more weight to periods with more treated households, consistent with its interpretation as an average treatment effect. The top right panel shows that the imputation estimator and OLS estimator largely agree on the treatment effects from rebates reported in June through August, but they imply very different treatment effects for rebates reported in the months of September and October.

To understand why the period treatment effects are different in OLS and the imputation method, the bottom left-panel shows the decomposition of the period coefficients into contributions from currently treated households compared to not-yet treated households and households that received their rebate in the past. For September, almost all the difference in the treatment effect between the DID imputation estimator and OLS comes from the comparison with the previously treated group. Put another way, OLS sees that households that report receiving their rebate in June display substantial negative consumption growth in the following interview (September); OLS then uses the sizeable negative consumption growth for past-treated units as a

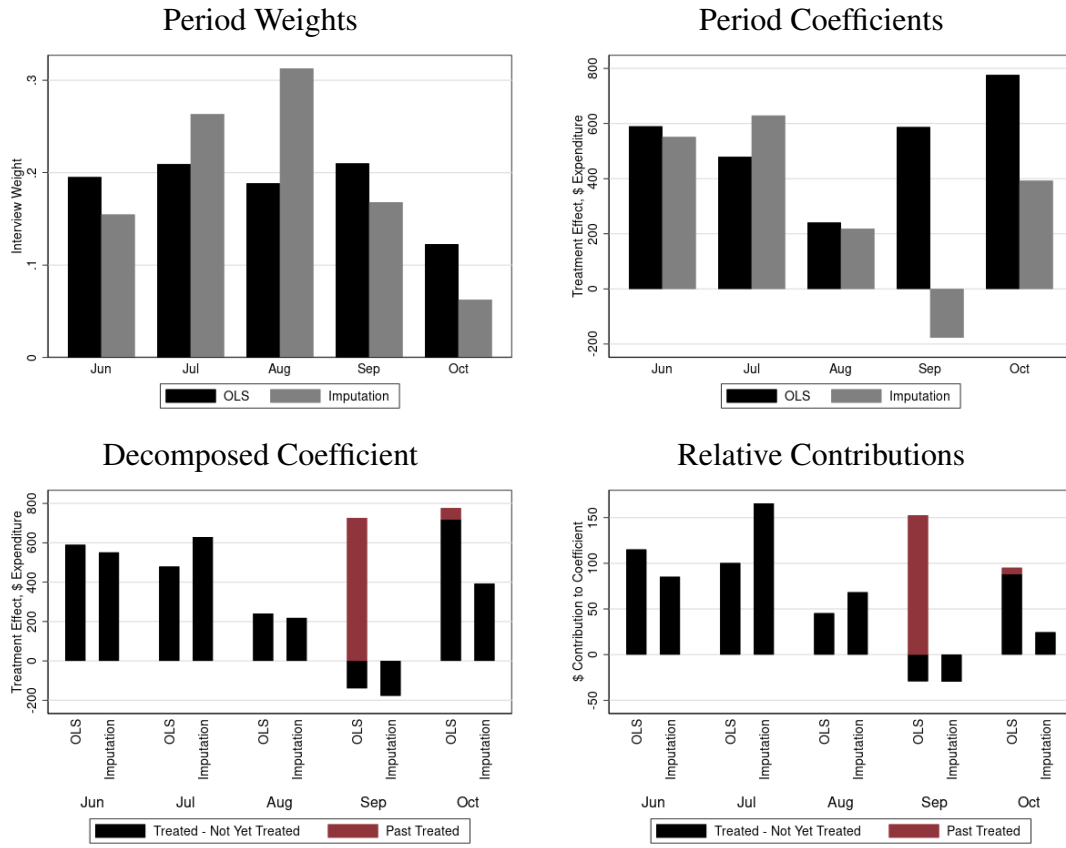


Figure 2.13. Decomposing the OLS and DID Imputation Coefficients

Notes. The dependent variable is the change in total expenditure. Based on estimations of equation 2.4.1 via OLS and the DID imputation method described in section 2.4. Periods after October, 2008, also receive positive weight, however, these weights are quite small and are not shown here.

counterfactual for the treated group.²¹ Borusyak et al. (2022) call these “forbidden comparisons” and remove them from their imputation estimator by dropping previously treated observations.

Previously treated households are unlikely to form a valid control group: expenditure growth in the September interview month is likely relatively low if the rebate did raise reported expenditures in the previous interview in June.²² For this reason we prefer the imputation

²¹Recall the the June interview captures expenditures from February through May, and the September interview captures expenditures from June through August.

²²Controls for lagged spending or lagged rebate do not solve the “forbidden comparison” problem: the comparison will remain invalid if treatment effects are heterogeneous across rebate cohorts. This is likely the case here because the cohorts differ by composition and time to spend the rebate: because EFT rebates were sent in May, the proportion of electronic filers among rebate recipients is highest in the June interview cohort and then decays to zero by September. Furthermore, the June and July interview cohorts had less time to spend the rebate as the earliest they

Table 2.7. Contemporaneous Household New Vehicle Expenditure Response to Rebate

Panel A: OLS				
	Full Sample		Rebate Only Sample	
	(1)	(2)	(3)	(4)
Rebate Indicator	301.2** (1287)	231.4* (1214)	310.8 (1922)	245.2 (1768)
Implied MPC	0.32	0.25	0.34	0.27
Extra Controls	No	Yes	No	Yes
Observations	17,229	17,229	10,343	10,343

Panel B: DID Imputation				
	Full Sample		Rebate Only Sample	
	(1)	(2)	(3)	(4)
Rebate Indicator	301.3** (1268)	235.8* (1212)	539.0* (3098)	173.7 (2992)
Implied MPC	0.32	0.25	0.56	0.18
Extra Controls	No	Yes	No	Yes
Observations	12,499	12,499	5,585	5,585

Notes: The dependent variable is the change in New Vehicle Expenditure from the previous interview. Standard errors, in parentheses, are clustered at the household level. Significance is indicated by: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions include interview (time) fixed effects, as well as household level controls for age, change in number of adults, and change in number of children. Extra controls refer to additional controls for household income decile and lagged total spending. Rebate sample includes only households that receive a rebate at some point during our sample period.

estimator. The bottom right panel shows the contribution the decomposed period treatment effect to the overall estimate in Table 2.5. It shows that comparison with the previously treated group in September accounts almost all of the difference between the imputation and the OLS estimator.

Table 2.7 displays the same analysis with new vehicle expenditure as outcome variables. We find that the MPCs for new vehicles in the full sample are quite similar using OLS and the imputation estimator. This suggests that the instability of the MPCs for total expenditure reflects other components of spending. Indeed, Table 2.8 shows that the MPC estimates for non-durable could have received it is in May.

Table 2.8. Contemporaneous Household Non-Durable Expenditure Response to Rebate

Panel A: OLS				
	Full Sample		Rebate Only Sample	
	(1)	(2)	(3)	(4)
Rebate Indicator	126.4*	116.2*	262.9***	241.5***
	(67.2)	(66.8)	(94.8)	(91.2)
Implied MPC	0.14	0.13	0.29	0.27
Extra Controls	No	Yes	No	Yes
Observations	17,229	17,229	10,343	10,343

Panel B: DID Imputation				
	Full Sample		Rebate Only Sample	
	(1)	(2)	(3)	(4)
Rebate Indicator	57.0	44.8	175.2	42.8
	(68.9)	(70.5)	(2125)	(2032)
Implied MPC	0.06	0.05	0.18	0.04
Extra Controls	No	Yes	No	Yes
Observations	12,499	12,499	5,585	5,585

Notes: The dependent variable is the change in Non-Durable Expenditure from the previous interview. Standard errors, in parentheses, are clustered at the household level. Significance is indicated by: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. All regressions include interview (time) fixed effects, as well as household level controls for age, change in number of adults, and change in number of children. Extra controls refer to additional controls for household income decile and lagged total spending. Rebate sample includes only households that receive a rebate at some point during our sample period.

expenditure are only half as large when estimated using imputation rather than OLS.

In short, we find that household MPC estimates are substantially smaller when we employ additional household controls or use an estimation method that is robust to “forbidden comparisons.” Our preferred estimates indicate an MPC for total consumer spending of 0.3 or below, with all of it accounted for by durable goods expenditures. Our finding that the MPC estimate declines by at least 40 percent is similar to Borusyak et al. (2022) finding of a decline of 50 percent using the same imputation estimator.

2.5 Conclusion

In this paper, we have argued that a standard New Keynesian model calibrated with the leading micro estimates of the marginal propensity to consume out of temporary stimulus payments implies counterfactual paths of consumption that are implausible. Using the 2008 tax rebate as a case study, we presented narrative and forecasting evidence that no events in late spring and summer 2008 should have caused aggregate consumption expenditures to plummet and then recover in August and September 2008. Using a two-good, two-agent New Keynesian model with standard amplification and high MPCs, we simulate the effect of the 2008 tax rebates and apply the simulated responses to actual aggregate consumption to create counterfactual paths of consumption had there been no rebate. The resulting counterfactual paths imply that consumption would have exhibited a sharp V-shape in late spring and summer 2008 if there had been no tax rebates. We argue that this counterfactual path is implausible.

We have reconciled the implausible counterfactual with the micro MPC estimates in two ways. First, we modified our two-good model, which features nondurable consumption goods and durable consumption goods (interpreted as motor vehicles), to allow more realistic supply elasticities of durable goods. This modification goes far to creating counterfactual consumption paths that are more plausible. Second, we re-estimated the micro MPCs in the CEX data using new methods that overcome problems with standard OLS estimates of treatment effects. The new method results in estimated MPCs that are noticeably lower than those in the literature. The combination of the modified model and lower micro MPC estimates results in counterfactual paths that are no longer implausible. However, they imply that the general equilibrium consumption multiplier on the 2008 tax rebates was below 0.2.

2.6 Chapter Acknowledgments

Chapter 2, in full, is currently being prepared for submission for publication of the material. This project was co-authored with Valerie Ramey and Johannes Wieland. The dissertation

author was a primary investigator and author of this paper.

Chapter 3

Household Inflation and Aggregate Inflation

3.1 Introduction

A critical decision for macroeconomic policy makers is how much to focus on controlling inflation versus other indicators such as unemployment or the output gap. Correct judgment on the optimal level of inflation “hawkishness” requires an understanding of the costs of inflation. Past research has focused on both the optimal level inflation for central banks to target (see Diercks 2017 for a summary), which ranges from -7.6 to 6 percent, as well as how much weight the central bank should place on stabilizing inflation at that level, usually as a coefficient in a Taylor rule.

This paper postulates an additional cost of increased inflation: an increase in the cross-sectional diffusion of household inflation rates, which I call inflation dispersion. The central idea behind inflation dispersion is that households do not have the same inflation rates because of (1) different preferences for goods (due to age, income, or other idiosyncrasies), or (2) different abilities to adjust to changes in prices (because of preference intensity, ability to travel to stores, accept lower quality products, etc). I document that the standard deviation of household level inflation rates each quarter is around 3-5 percentage points using both a large scanner data-set and the consumer expenditure survey. However, the dispersion of inflation rates is not constant; during periods of greater aggregate price change, inflation dispersion increases.

While a household's wage or pension and the interest rate on borrowing and savings may be indexed to the aggregate inflation rate, only the price paid for the household's basket of consumption goods is related to their personal inflation rate. This means that with inflation dispersion, households with high inflation will have a real welfare loss compared to households with low inflation. I show that an inflation shock leads to a persistent increase in a household's price index. Households respond to the personal inflation shock by (weakly) reducing their nominal consumption expenditures; this means that real consumption falls more than one-for-one. Poorer households seem to be the least able to smooth their consumption, which means that an increase in inflation dispersion disproportionately increases the volatility of consumption of poor households relative to wealthier households. Although households benefit from deflation shocks, risk averse households should prefer lower inflation dispersion and smoother consumption.

I relate inflation dispersion to aggregate inflation by showing that there is a robust relationship between the level of aggregate inflation and cross-sectional distribution of household level inflation rates. I estimate that a one percent increase in the absolute value of aggregate inflation increases the standard deviation of individual inflation rates by 0.38 percentage points (in my preferred specification). This relationship is robust to differences in price index calculation (Laspeyres, Paasche, Fisher, or Sato-Vartia), data, and time (present in city level CPI data back to 1915). I find some evidence that unexpected inflation is driving this relationship (higher expected inflation may even reduce inflation dispersion).

I develop a multi-sector menu cost model where households have heterogeneous preferences and I solve this model in partial-equilibrium. In this model, the elasticity of demand for a firm's product differs across sectors, which means that the profit function for firms is more curved in some sectors than in others, which endogenously leads to differences in price setting behavior. Households have heterogeneous preferences across sectors. When aggregate inflation increases, firms with high elasticities of substitution for their products are more likely to change their prices. Household's with a greater consumption share of products in these sectors experience a higher inflation rate than others and the distribution of household level inflation

rates widens. Simulations of my model lead to a relationship between inflation dispersion and the absolute value of aggregate inflation that fall within my empirical estimates.

My project contributes to several strands of literature. The first is a very large literature on the costs of inflation and the corresponding optimal inflation rate. While I will not endeavor to summarize this entire literature, past literature has shown that increases in inflation is related to increases in inflation volatility (Kim and Lin 2012). Doepke and Schneider (2006) show that unexpected inflation can lead to redistribution from lenders to borrowers. Past research has discussed the so called “shoe leather” cost of inflation where households spend real resources to protect themselves from inflation (Pakko 1998). Menu costs incurred by firms as they must employ more labor to determine optimal prices (Golosov and Lucas Jr 2007, Nakamura and Steinsson 2010) and many more. An increase in inflation dispersion is unlike many of the other costs of inflation, since it cannot be resolved by indexing to the aggregate inflation rate; i.e. loans, wages, rent, etc. can be indexed to the aggregate inflation rate, but not to a household’s personal price index.

The cost of inflation is a perennial question with a mature literature, but there have been some recent developments. In a New Keynesian model, the largest cost of inflation is an increase in price dispersion (Coibion et al. 2012). Higher inflation rates mean that firm prices are more likely to be far from their desired price, which leads costly misallocation of resources to firms with “artificially” low prices. However, Nakamura et al. (2018) do not find increases in price dispersion during the US great inflation of the 1970s. While Alvarez et al. (2018) do find evidence of price dispersion during the Argentine Hyperinflation, they do not find price dispersion when inflation is at more moderate levels. My project shows that while price dispersion may not be a consequence of higher inflation rates, higher inflation dispersion may be. My multi-sector menu-cost model does not produce a relationship between inflation and price dispersion, but is able to reproduce the relationship between the absolute value of aggregate inflation and inflation dispersion that I find in the data.

There is also an emerging literature on demographic group specific price indexes. Argentine

and Lee (2017) shows that during the great recession, inflation for the lowest income group was around half a percent greater than for the highest income group. Cravino et al. (2018), show that rich households spend a larger fraction of their income on “sticky” goods such as education compared to middle-income households who spend more money on goods like gasoline that change their prices often; this makes the price index of middle-income households more responsive to monetary policy shocks than the price index of richer households. Like my project, Kaplan and Schulhofer-Wohl (2017) also use the Nielsen Homescan data to compute household specific inflation rates; however, they restrict their inflation measure only to products defined at the barcode level that the household buys in between two periods. This results in their inflation measures only being representative of around a quarter of total spending in the Nielsen Homescan. My project makes many improvements on their household inflation measure (including increasing the relevance of the measure to be representative of 60-99 percent of Nielsen spending) and I also use the broader Consumer Expenditure Survey to calculate household level inflation rates. My project is also the first to examine how households react to personal inflation rates that are higher or lower than aggregate inflation.

This project also expands on recent work by Gelman et al. (2016) on the marginal propensity to consume (MPC) out of changes in gasoline prices. They find that the estimated (MPC) out of savings from lower gasoline prices is approximately one. My work looks at changes in the entire price index of the household rather than just gasoline prices. I do find evidence that households increase their consumption spending following a fall in their price index (which can be thought of as a persistent wealth shock). However, my estimates for increases in spending following a deflation shock are smaller (around 0.8).

The remainder of the paper proceeds as follows. Section 2 explains how I create the individual inflation measures using first the Nielsen Homescan and second the consumer expenditure survey. Section 3 presents a simple household level model showing that a household reacts to increases in its price index by reducing real consumption and then confirms this fact empirically. I also show that changes in a household’s personal price index are quite persistent. Section 4

shows that there is a robust relationship between inflation and inflation dispersion. Section 5 develops a menu-cost model that can explain this relationship. Section 6 concludes.

3.2 Measuring Household-Level Inflation

For this project, I construct novel measures of household-level inflation using two large and complementary datasets: the Nielsen Homescan Consumer Panel Data and the Consumer Expenditure Survey (CEX). The Nielsen Homescan, is a for-profit market research survey that tracks the retail purchases of approximately 178 thousand households from 2004-2017. The Consumer Expenditure Survey, administered by the Bureau of Labor Statistics, surveys households up to five times at three month intervals (only four of these surveys are available for public use) on all of the household's consumer expenditure; the consumption weights in the Consumer Price Index (CPI) are constructed using results from the CEX.

Each of these data have distinct advantages. The Nielsen Homescan is able to track households for a long period of time (the average household is surveyed for eight years) includes very detailed information on the products the households purchase (at the barcode level), and has a large number of households (40,000 from 2004-2005, and 60,000 from 2006-2017); however, the Nielsen Homescan only includes information on the household's retail purchases (about 30 % of the spending in the CEX (Kaplan and Schulhofer-Wohl 2017)) and excludes purchases of some major categories including housing, transportation, and medical care. In contrast, the CEX explicitly asks households the total sum of all of their consumption spending in the past three-months and includes more detailed information on all of the large categories that are used to construct the CPI. Additionally, the CEX is available for a longer time period than the Nielsen data: this project is using CEX data from 1996-2017, but it is possible to extend the sample back until 1980¹. However, the CEX only includes survey responses for about 10% of the number of households as the Nielsen Homescan (5,416 in 2017Q1) and households are in the survey for at

¹There may be problems with survey quality in earlier years (see NBER's discussion www.nber.org/data/ces_cbo.html)

Table 3.1. Distribution of Spending

	CPI-U	CEX	Consumption (Nielsen)	Household Inflation
Health and Beauty Aids	2.57	1.43	10.3	10.0
Food	8.60	8.91	63.4	64.3
Alcohol	0.95	1.03	3.5	3.7

Note: Raw unweighted shares. CPI-U and CEX shares in Kaplan and Schulhofer-Wohl (2017). The household inflation data is based on the common price and represents 82% of the spending in the full Nielsen Homescan.

most 4 quarters. By using both the Nielsen Homescan and the CEX, I am able to show that my results are robust to the main weaknesses of each dataset.

3.2.1 Nielsen Homescan

The Nielsen Homescan Consumer Panel Data tracks the retail purchases of 40,000-60,000 households from 2004-2016 and includes information on household demographics and income levels. The most novel feature of these data is that they track the products that households purchase at the barcode level, which makes these data uniquely situated to measure very detailed household level retail inflation rates , as well as the total retail consumption of these households.

Households in the panel are given financial incentives to record consumption purchases (similar to a credit card rewards program). To facilitate the survey, each household is given a barcode scanner so that they can easily record the individual products that they buy. Household's are also asked to record where they bought the product. If the product was bought at one of Nielsen's partner stores then the price is automatically recorded as the average price in that store for that product during the week of purchase, if the product is bought somewhere else then the consumer is asked to record the price. Nielsen argues that the homescan panel is representative of 30% of all consumption (Kaplan and Schulhofer-Wohl 2017).

While households are able to record gasoline purchases and other non-grocery products, the actual survey responses are heavily skewed toward grocery purchases. Table 3.1 shows an overview of the purchases in the Nielsen Homescan compared to the CPI-U and the Consumer Expenditure Survey. While food at home represents less than 9 percent of the basket for the

CPI-U and the CEX, it is more than 60 percent in the Nielsen Homescan.

Purchases in the Nielsen Homescan are recorded at the barcode (UPC) level. Each UPC is also associated with a hierarchy of classifications of increasing levels of aggregation: brand, module, group and department. For example, if a household were to purchase a particular type of toothpaste it would be associated with the Health and Beauty Aids department in the toothpaste module and the UPC would denote the specific flavor/ingredients. For my analysis I will use the product module as the definition of an individual product, however, my main result is qualitatively robust to using more aggregate or dis-aggregate product definitions.

Since households may enter or exit the survey in the middle of a quarter, I exclude each household's first and last quarter in the panel. I also exclude households with breaks in recorded transactions for more than a quarter and households with extreme changes in their price index ($> 300\%$) from quarter to quarter.

Using the Nielsen Homescan, I can track individual households recorded retail spending each quarter, the quantities they purchase of each product, and the prices they pay for each product². I use this information to compute Laspeyres, Paasche, Fisher, and Sato-Vartia style household specific chained price indexes. Household level inflation rates are computed as the annual percentage change in the chained price index between quarter t and $t-4$. I create each of these household level price indexes using both national prices (the weighted average of all prices paid for that product) and regional prices (the average in one of Nielsen's 52 metropolitan areas). Summary statistics for each of these 8 price indexes are shown in table 3.2.

Laspeyres and Paasche are the familiar undergraduate price indexes formed by weighting each product in a household's basket by its beginning (Laspeyres) or ending (Paasche) expenditure share (they implicitly assume the household has Cobb-Douglas utility over products). By nature, the Laspeyres index understates changes in the cost of living as it does not allow for substitution on the part of households. Similarly, the Paasche index overstates inflation. The

²There is measurement error in the Nielsen Homescan (Einav et al. 2010). However, as long as the measurement error is orthogonal to inflation rates and monetary policy shocks then it should not bias my results.

fisher index is the geometric average of the two. Lastly, the Sato-Vartia index assumes that individual households have CES utility over the N products in their basket:

$$P_{h,t} = \left(\sum_{k=1}^N \left(\frac{p_{k,t}}{\varphi_{k,h}} \right)^{1-\sigma_h} \right)^{\frac{1}{1-\sigma_h}}. \quad (3.2.1)$$

Since I compute this index at the individual level, each household h, has its own specific elasticity of substitution over products σ_h , which can be an important source of heterogeneity in inflation rates as some households are able to adjust more to price changes than others. I'm able to avoid estimating σ_h for each household since:

$$s_{k,h,t} = \left(\frac{p_{k,t}}{\varphi_{k,h} P_{h,t}} \right)^{1-\sigma_h}$$

which implies

$$\pi_{h,t} = \log \left(\frac{P_{h,t}}{P_{h,t-1}} \right) = \sum_{k=1}^N \omega_{k,h,t} \log \left(\frac{p_{k,t}}{p_{k,t-1}} \right) \quad (3.2.2)$$

where

$$\omega_{k,h,t} = \left(\frac{\frac{s_{k,h,t} - s_{k,h,t-1}}{\log(s_{k,h,t}) - \log(s_{k,h,t-1})}}{\sum_{\ell=1}^N \frac{s_{\ell,h,t} - s_{\ell,h,t-1}}{\log(s_{\ell,h,t}) - \log(s_{\ell,h,t-1})}} \right).$$

I can examine changes in the Sato-Vartia price index for each household, as long as I have expenditure share information for each product in period t and t-1. Since households stop buying some products from quarter to quarter, this means that my Sato-Vartia price index is not representative of all of the household's Nielsen spending (about 60% of a household's purchases are from products that a household buys in quarter t and t-1). This is a limitation of the Sato-Vartia index that the Laspyres and Paasche do not share since for these indexes I only need

price information in two quarters and share info for only one quarter (Laspeyres and Paasche indexes are representative of about 98% of Nielsen spending).

I do not consider the Feenstra (1994) index nor the more recent Redding and Weinstein (2018) index. Each of these indices contain a variety adjustment term that is meant to capture changes in the cost of living due to new products entering the market (they are also based on the CES utility and since CES has a preference for variety, new products equals lower cost of living). However, it is not clear how these new products should affect individual level cost of living³. When an individual purchases a product that they did not buy in the previous quarter it should not have the same effect on their cost of living as a new product entering the market. Practically, both Feenstra (1994) and Redding and Weinstein (2018) indexes require estimating elasticities of substitution, which would be both computationally heavy and introduce substantial measurement error if estimated at the household level.

The Nielsen Homescan data includes information on the actual price paid for the barcode level product. Argente and Lee (2017) and Kaplan and Schulhofer-Wohl (2017) use changes in the actual price paid at the barcode level for income groups and households respectively to compute changes in the cost of living. I do not follow their example for two reasons: (1) my definition of product is more aggregated (product module instead of barcode) so changes in the price paid by the household could possibly represent the household switching from a less expensive to a more expensive item within a category, (2) at the household level, one reason that a price paid for a product could change is that the household may buy the product at a cheaper location one period (say a bulk grocery store) and then buy that same product at more expensive location the next (a convenience store); although such household level decisions are interesting, they introduce substantial noise into my inflation measures and are outside of the scope of changes in the cost of living that I want to consider (the combination of a less-aggregated definition of product and using household level prices in Kaplan and Schulhofer-Wohl (2017)

³One can assume that the new products affect all household's cost of living equally, however, since I'm looking at differences in cost of living changes between households the Feenstra index would simply collapse to the Sato-Vartia.

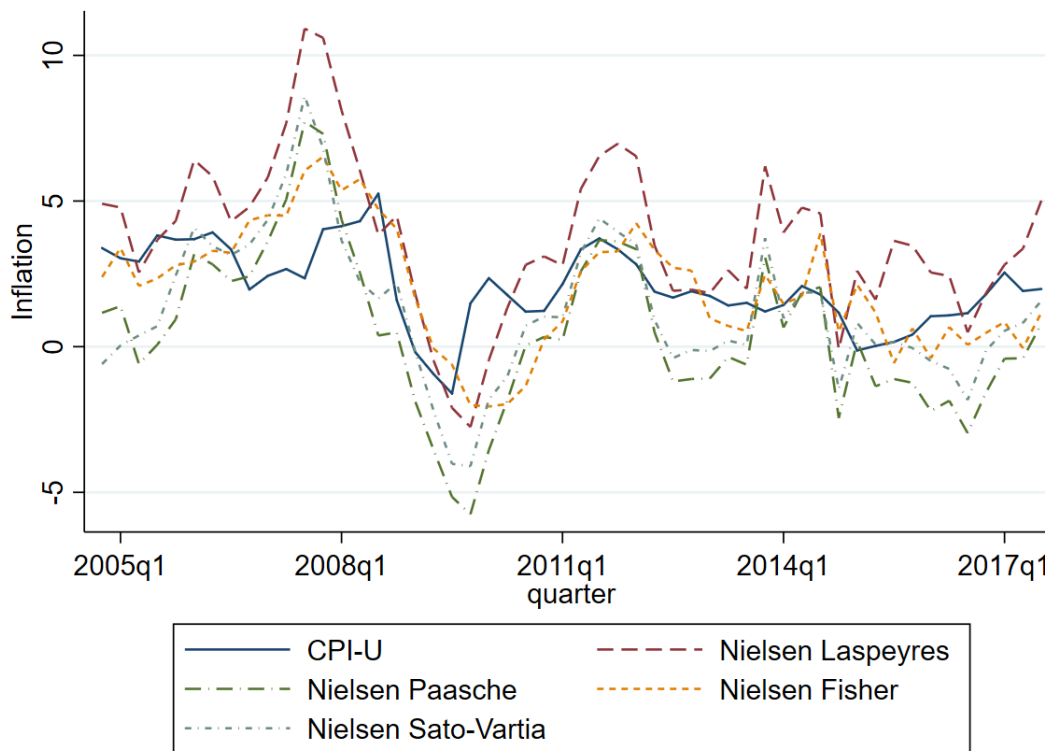


Figure 3.1. Nielsen Individual Inflation v. CPI-U

Note: Nielsen Inflation rates are the weighted average of the individual inflation rates for the quarter (democratic index). Prices for each product module are the average national prices.

leads their inflation rates to only be representative of 25% of Nielsen household spending). See appendix C6 for a more complete discussion on my use of common (national) versus effective (price paid) prices.

Figure 3.1 shows the weighted average of the Nielsen individual inflation rates computed using each of the four methods (with national prices) compared to the CPI-U. As expected, in each quarter the Laspeyres inflation measure has the highest inflation rate, while the Paasche measure is the lowest. The Sato-Vartia and the Fisher are in the middle. The CPI-U is less volatile than any of the other indexes likely because it includes a wider variety of products rather than just representing retail inflation as in the Nielsen data, however, it does roughly follow the same trends as the four Nielsen inflation measures.

3.2.2 Consumer Expenditure Survey

I also create individual inflation measures using the consumer expenditure survey from 1996-2017. The manner that I construct these measures is straightforward and is similar to what Hobijn and Lagakos (2005) does at the demographic-group level. Households are in the survey for up to four quarters and are asked about their expenditure on a variety of different product classes (gasoline, rent, electronics, etc.) during the quarter. I aggregate their responses into one of 26 categories for which the BLS provides category specific CPI's and then match household level expenditure shares with the BLS category level CPI data (see appendix for more details). For each household I construct sequential Laspeyres π^l and Paasche π^p indexes as:

$$\pi_{h,t}^l = 100\left(\sum_i^h s_{i,t-4} \frac{p_{i,t}}{p_{i,t-4}} - 1\right), \quad \pi_{h,t}^p = 100\left(\sum_i^h s_{i,t} \frac{p_{i,t}}{p_{i,t-4}} - 1\right), \quad (3.2.3)$$

where p_i is the CPI index for category i . Only my Paasche inflation rates line up with consumer expenditure information for the quarter since the Laspeyres individual inflation rates require prices a year after the share data, which is after the household has dropped from the sample.

Figure 3.2 shows the distribution of individual inflation rates over time. As seen in the figure, the distribution is quite wide, households experience inflation at very different rates. Also, the distribution narrows and widens over time. Periods of high inflation or deflation seem to have wider distributions than periods of tame inflation.

Finally, Table 3.2 shows the summary statistics of the Nielsen expenditures, Quarterly expenditures in the CEX measured in two ways, and my computed inflation measures.⁴ The CEX spending is considerably higher than the Nielsen spending, which is unsurprising since

⁴In the CEX survey they ask households their total spending in each of many categories during each month of the quarter. The total categorical CEX spending is the quarterly total of this measure and this is also the expenditure I use to create household level inflation rates. Households are also asked about total expenditures in the last three months. Since households are interviewed in either the first, second, or third month of the quarter, this latter question does not line-up with the calendar quarter.




Figure 3.2. Distribution of Individual Inflation Rates: CEX Laspeyres Index
 Note: Individual Inflation from the CEX using a Sequential Laspeyres Price Index for each household.

the Nielsen homescan only includes retail spending. Interestingly, the standard deviation of individual inflation rates (with national prices) is close to 5 percentage points regardless of the dataset or manner in which inflation is calculated. Inflation rates using regional prices are much more dispersed than those using national prices. I should note that the summary table pools all of the data together in the entire sample. When I first look at the standard deviation of inflation rates by quarter, the average standard deviation is smaller (ranging from 3.4-5.06 percentage points for Nielsen national prices).

3.3 How Consumers React to Higher Household Inflation

The previous section showed that there is a wide distribution of individual inflation rates around the aggregate/mean level of inflation. In this section, I present a simple household-level model that shows how a household reacts to inflation rates that are higher than the aggregate.

Table 3.2. Summary Statistics**Panel A: Individual Data**

	Mean	Median	SD	10 %	90 %
Quarterly Expenditure (Nielsen)	1082	925	716	378	1959
Quarterly Expenditure (CEX Last Three Months)	7284	4919	8258	1482	15360
Quarterly Expenditure (CEX total Categorical)	6755	4373	9387	1577	12888
Nielsen National Prices					
Laspeyres Inflation	3.57	2.85	5.80	-1.71	9.04
Paasche Inflation	0.33	0.48	5.11	-4.71	5.46
Fisher Inflation	1.87	1.71	3.95	-2.33	6.23
Sato-Vartia Inflation	1.10	0.97	4.53	-3.44	5.83
Nielsen Regional Prices					
Laspeyres Inflation	6.69	4.73	13.03	-3.99	9.04
Paasche Inflation	1.85	1.15	19.03	-9.06	12.15
Fisher Inflation	3.80	2.98	9.90	-4.58	12.59
Sato-Vartia Inflation	2.54	1.57	37.24	-7.15	11.94
CEX					
Laspeyres Inflation	1.69	2.36	5.27	-2.63	5.62
Paasche Inflation	1.57	2.24	5.65	-3.03	5.94

Panel B: Standard Deviation of Inflation Measures

	Quarterly		
	Average	Min	Max
Nielsen National Prices			
$\sigma(\text{Laspeyres } \pi)$	5.06	3.00	8.23
$\sigma(\text{Paasche } \pi)$	4.41	2.88	7.21
$\sigma(\text{Fisher } \pi)$	3.40	2.76	4.62
$\sigma(\text{Sato-Vartia } \pi)$	3.90	2.56	6.70
CEX			
$\sigma(\text{Laspeyres } \pi)$	3.75	1.43	8.06
$\sigma(\text{Paasche } \pi)$	4.01	1.71	8.47

Note: Statistics weighted by population projection factors from Nielsen and CEX. Inflation is in terms of percent change in the price index between quarter t-4 and t where the price index is computed as described in the text. Expenditure is in Nominal Dollars.

I then use my Nielsen and CEX expenditure data and household inflation rates to show how households react to higher individual inflation.

3.3.1 Simple Model

This simple model is meant to show how a single household behaves in an environment where their price index (and inflation) differ from the aggregate, but their wage, interest rate, and all other terms are indexed to aggregate inflation. I show that persistent increases to the household's own price index cause the household to reduce consumption, as if they had been hit by a negative wealth shock. For simplicity, I assume that the household buys a single consumption good, which is a stand-in for the household's unique basket.

I assume that households have concave utility $u(\cdot)$ strictly increasing over real consumption good c , for which they must pay household specific price p_h ; c is meant to represent the household's real consumption bundle and p_h its price index. I also assume that they provide labor in-elastically and receive wage w_t . They can also invest in a one-period bond b which the household can sell at real price Q_t . I assume that the households wage and savings are indexed to the aggregate price level (P_t) only.

The household's budget constraint is then:

$$p_{ht}c_t + P_t b_{t+1} = P_t w_t + P_t b_t Q_t \quad (3.3.1)$$

Dividing through by the aggregate price level, I denote the part of the household price index that is orthogonal to the aggregate price index as \tilde{p}_{ht} and real valued wage and bonds as \tilde{w} and \tilde{b} respectively. The real interest rate corresponding with Q is r_t^* .

The household's problem is:

$$\begin{aligned}
& \max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta u(c_t) \\
& \text{s.t.} \quad \tilde{p}_{ht} c_t + \tilde{b}_{t+1} = \tilde{w}_t + \tilde{b}_t (1 + r_t^*) \\
& \quad \quad \quad \lim_{s \rightarrow \infty} \left(\prod_{k=1}^s (1 + r_{t+k}^*) \right)^{-1} \tilde{b}_{t+s} = 0.
\end{aligned} \tag{3.3.2}$$

The Euler equation resulting from this simple model is:

$$\frac{u'(c_t)}{\tilde{p}_{ht}} = \mathbb{E}_t \frac{u'(c_{t+1})}{\tilde{p}_{ht+1}} \psi_t \tag{3.3.3}$$

where $\psi_t = \beta(1 + r_{t+1}^*)$. It follows that when a household expects that their price index will be lower tomorrow than today, they will decrease their consumption today and increase their consumption tomorrow, which reflect the fact that the effective real interest rate for the household is different than the real interest rate of the economy as a whole. Hence, in this model, shocks to a household's inflation rate will lead to shocks to their real consumption. If $u(\cdot)$ is such that the household is risk averse, the household would prefer that p is constant over-time (aka, that their inflation is always equal to aggregate inflation).

As an illustration, consider the simple case where the household has quadratic utility $u(c_t) = c_t - \gamma c_t^2$. Under perfect foresight of the path of their future price index and the assumption that $(1 + r^*)(\beta) = 1$ it can be shown that:

$$c_t = \frac{r}{1+r} \left(\sum_{j=0}^{\infty} \frac{\beta^j}{\tilde{p}_{ht+j}} \mathbb{E}_t (\tilde{w}_{t+j} + \tilde{b}_t) \right). \tag{3.3.4}$$

Consumption at time t is a fraction $\frac{r}{1+r}$ of expected future income streams, current assets, and the expected buying power of the income and assets. Increases in the future path of the price

index lead to decreases in real consumption, as higher price levels mean that the household can buy less with the same amount of income. See the appendix for more details.

3.3.2 Household consumption response in the data

I next turn to the data to show (a) how a household level inflation shock affects the household's price index over time, and (b) how households respond to this shock. I find that a shock to a household's price index is quite persistent for all price indexes and permanent for shocks to the Paasche and Sato-Vartia price indexes. Consistent with the simple model presented above, following an inflation shock households reduce their real consumption. I also find weaker evidence that a household reduces its nominal consumption expenditures following an inflation shock, which implies that real consumption falls more than one-for-one following a household level price shock. I should note that since I cannot reliably see changes in income in the Nielsen or CEX data, I cannot determine if the behavioral changes I see are because households are adjusting their savings or their labor inputs in response to an inflation shock.

I start by using the chained-price indexes and corresponding inflation rates in the Nielsen data to construct impulse responses of the price index, nominal consumption, and real consumption following a shock to the household's inflation rate. I construct the IRF's following the Jordà (2005) method:

$$y_{h,t+k} = \beta_0 + \beta_1 \pi_{h,t} + \beta_3 y_{h,t-1} + \gamma_h + \alpha_t + \varepsilon_{h,t+k}, \quad k \in \{0, 1, \dots, 12\}. \quad (3.3.5)$$

Above, $y_{h,t+k}$ is the price index, log of nominal consumption, or log of real consumption for household h at time $t+k$. The household specific Laspeyres, Paasche, Fisher, or Sato-Vartia inflation rate denoted by $\pi_{h,t+k}$ is normalized so that it results in a one unit increase in the price index at time t . Household fixed effects are denoted by γ_h and time fixed effects by α_t .

Figure 3.3 shows the results of the regressions of the price index on the inflation shock.

The time fixed effects are designed to absorb the movement of the aggregate price index, so that the IRF's of the individual price index can be interpreted as the difference between the individual price index and the aggregate index. I find that following an inflation shock, the individual's price index stays above the aggregate price index for at least 4 quarters and remains at a permanently higher level for the Paasche and Sato-Vartia style price indexes. Note that since I am using chained-price indices, all of these indices allow the household to adjust its basket each quarter. Even accounting for these behavioral changes the price index shock is quite persistent (lasting around 2-3 years). One way to interpret these household inflation shocks is as a persistent shock to their spending power or real wealth.

I next show how this same inflation shock affects the household's nominal consumption expenditures (Figure 3.4). If households are attempting to smooth their real consumption then we should expect that they would increase their nominal expenditures following an increase in their price index, however, I find some evidence that the opposite is happening. Household's appear to reduce their nominal consumption (for Paasche, Fisher, and Sato-Vartia style inflation shocks) in the period of the inflation shock. Later on, once their price index has reverted closer to the aggregate price index, household's increase their nominal consumption. Similar to the results from the simple model, household's seem to buy more when prices are low and buy less when prices are high.

Finally, I show how a household's real retail consumption expenditures react to an inflation shock. I calculate real consumption expenditures by dividing a household's nominal consumption by their chained price index. Figure 3.5 shows the results of this analysis. Following a one unit inflation shock, real consumption falls by more than one percent in each of the definitions of real consumption. The fall in real consumption is quite persistent and only recovers for Fisher style real consumption. It remains at a permanently lower level for the other measures of real consumption. This suggests that household level inflation shocks can have large welfare effects.

The results I have presented above using the Nielsen data only show a household's

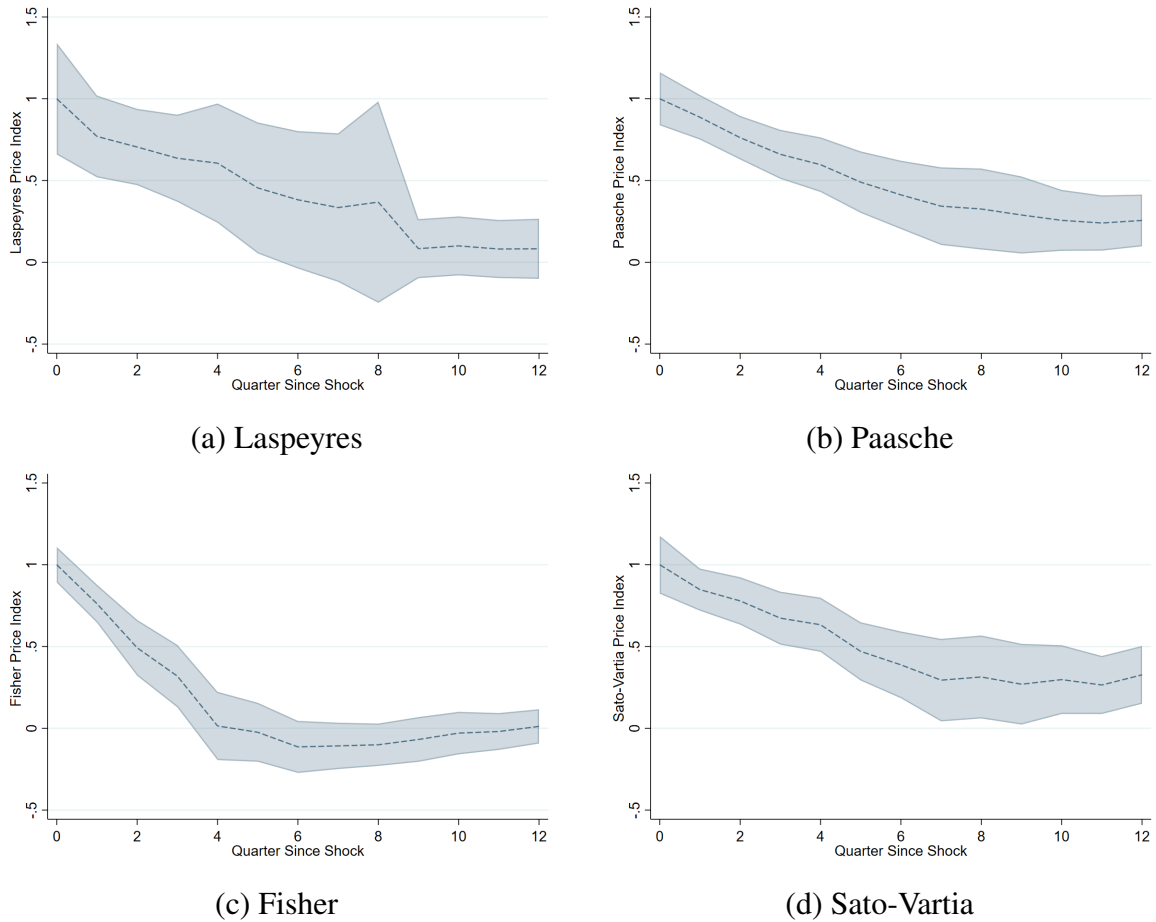


Figure 3.3. Path of Household Retail Price-Index Following Household Inflation Shock

Note: Standard errors robust to auto-correlation and are two-way clustered at the quarter and household level. Confidence interval (99 percent) is shown as the shaded area.

retail expenditure response to their retail price index. Since I only observe households for at most 4 quarters in the CEX, I cannot construct similar IRF's using all of the household's consumption expenditures. However, to test whether a household responds to an inflation shock by smoothing their real consumption or reducing their consumption as in my simple model I regress the household's CEX nominal consumption expenditures on their Paasche inflation rate, and lags of their expenditure and inflation rate along with household and time fixed effects. In my baseline model, I only include one lag⁵ I also check to see whether there is asymmetry in

⁵I can include at most two lags. I have up to four observations for each household, and the household fixed effect takes up one observation, the expenditure and two lags take up the other three.

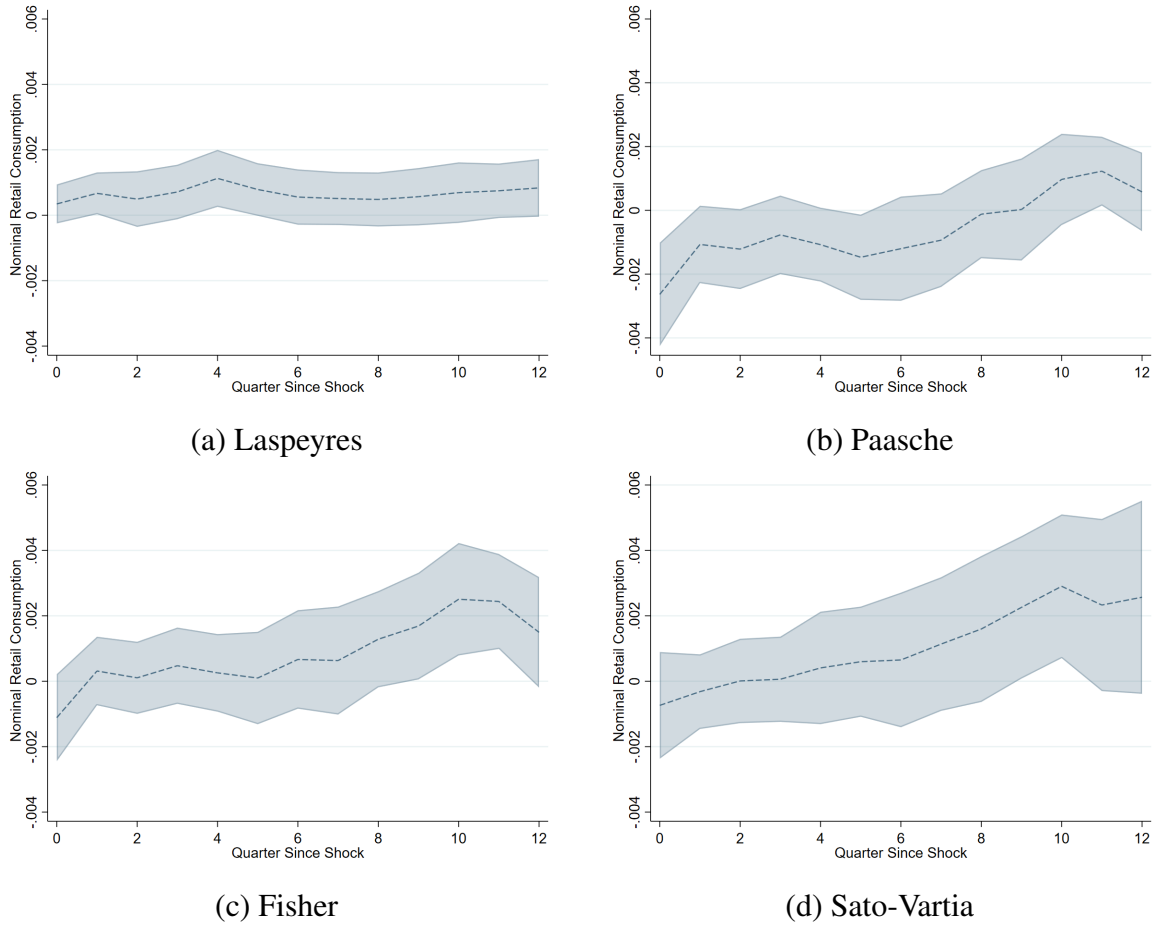


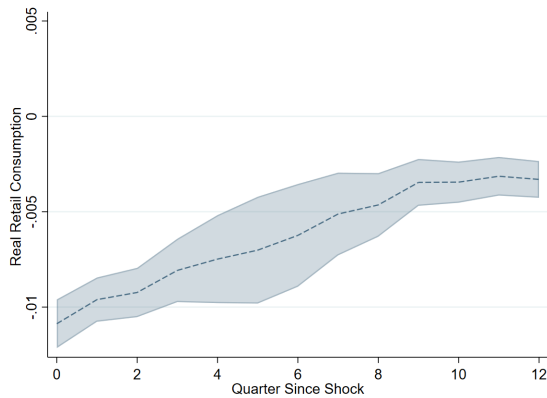
Figure 3.4. Response of Nominal Household Retail Consumption to one-unit Household Inflation Shock

Note: Standard errors robust to auto-correlation and are two-way clustered at the quarter and household level. Blue shaded area indicates 95 percent confidence interval.

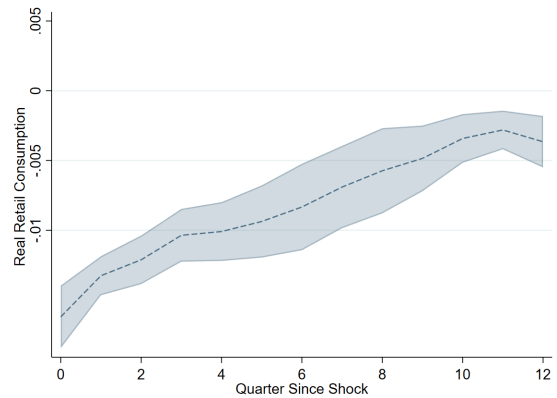
the household’s response to inflation shocks that are either above or below aggregate inflation.

I denote $\pi_{h,t}^+ = \max\{\pi_{h,t} - \bar{\pi}_t, 0\}$ and $\pi_{h,t}^- = -\min\{\pi_{h,t} - \bar{\pi}_t, 0\}$ where $\bar{\pi}_t$ is aggregate Paasche inflation for period t .


I show the results from this analysis in Table 3.3. Column 1 shows the response of a household’s nominal consumption to the household specific inflation rate. As with the Nielsen data, nominal consumption declines following a household inflation shock and I can reject that the household smooths its consumption (aka that nominal consumption increases by 1 percent following a 1 percent inflation shock) at the 99 percent level for lagged inflation and




(a) Laspeyres



(b) Paasche



(c) Fisher



(d) Sato-Vartia

Figure 3.5. Response of Real Household Retail Consumption to one-unit Household Inflation Shock

Note: Standard errors robust to auto-correlation and are two-way clustered at the quarter and household level. Blue shaded area indicates 95 percent confidence interval.

the 90 percent level for contemporaneous inflation. Column 3 shows the response of household consumption but this time allowing for asymmetric responses for positive and negative personal inflation shocks. I find that following a negative inflation shock, household's appear to take advantage of the lower price index and increase their consumption, however, I also find that following a positive inflation shock households decrease their consumption. However this second result is not significant and I can reject that households are able to smooth their real consumption following an increase in their price index at the 90 percent level for column 3.

It is possible that households would like to smooth their real consumption following an

Table 3.3. Response of Household Spending to Household Inflation Shock: CEX

	$\ln(P \cdot C)$	$\ln(P \cdot C)$
$\pi_{h,t}$	-0.00523* (0.00273)	
$\pi_{h,t-1}$	-0.00125** (0.000577)	
$\pi_{h,t}^+$		0.000895 (0.00487)
$\pi_{h,t-1}^+$		-0.000567 (0.00120)
$\pi_{h,t}^-$		0.00749*** (0.0164)
$\pi_{h,t-1}^-$		0.00161* (0.000920)
N	347,462	347,462

Note: Standard errors, in parentheses, clustered at time level and are robust to auto-correlation. Significance at the one percent, five percent and ten percent levels indicated by ***, **, and *. Individual and time fixed effects and a lag of the dependent variable are also included.

increase in their price index, but not all households are able to do so because of credit constraints. Table C2 in the appendix shows the difference in household's consumption responses to an inflation shock by income group. Richer households are far more responsive to inflation shocks than poorer households (to both positive and negative inflation shocks). This could imply that richer households are less credit constrained; less credit constrained households are more able to take advantage of low prices following a negative inflation shock and smooth their consumption following a positive income shock. However, other explanations, such as different preferences to smooth consumption over time can also explain this same pattern.

3.4 Inflation Dispersion and Inflation

In this section I discuss the relationship between inflation dispersion and aggregate inflation. I show that there is a robust relationship between the level of absolute aggregate

inflation and individual inflation dispersion.

I start by showing that the correlation between aggregate price changes and inflation dispersion is robust. Table 3.4 shows the results of regressing the standard deviation of individual inflation rates on the absolute value of aggregate inflation; the even columns include the expected inflation rate in that period (the difference between the TIPs 5 year bond yield and the 5 year treasury yield). Panel A shows the results using inflation rates calculated using the Nielsen data. I find that the relationship between aggregate inflation (defined as the weighted average of the individual inflation measures, a democratic index) is positive in all cases. It is statistically significant for the Paasche, and Sato-Vartia definitions of inflation and for the Laspeyres definition of inflation once expected inflation is included.

Panel B shows the results using the CEX data, which are positive and statistically significant in every specification; a one percent increase in Laspeyres inflation increases the standard deviation of individual inflation rates by 0.38 percentage points (off of a mean of 3.75). In the CEX data, I find that it is unexpected inflation that is driving this relationship; in fact increases in expected inflation seem to decrease inflation dispersion.

McLeay and Tenreyro (2019), use the fact that regions have their own labor market conditions and inflation rates to identify regional Phillips curves, since the national Phillips curve can be hard to identify due to central bank actions. In a similar way, I check the robustness of the relationship between inflation and inflation dispersion by running similar regressions at the regional level. This lets me confirm that it is not oil price shocks that are affecting inflation and some products thereby also causing inflation dispersion. It also significantly increases my statistical power. For these regressions I regress the standard deviation of the Nielsen Household inflation rates using regional prices on the regional inflation rate; I also include time and region fixed effects. Table 3.5 shows the results of this analysis and confirms that the relationship between inflation and inflation dispersion also exists at the regional level.

Vavra (2013) and Li (2019) discuss the relationship between price change dispersion and

Table 3.4. Household Inflation Dispersion and Aggregate Inflation**Panel A: Nielsen Homescan**

	Laspeyres		Paasche		Fisher		Sato-Vartia	
	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$
$ \bar{\pi} $	0.058 (0.054)	0.13** (0.053)	0.29** (0.12)	0.29** (0.11)	0.074 (0.058)	0.065 (0.064)	0.26*** (0.073)	0.30*** (0.071)
$ \mathbb{E}(\bar{\pi}) $		-0.75 (0.48)		0.12 (0.24)		0.069 (0.15)		-0.56*** (0.16)
N	52	52	52	52	52	52	52	52

Panel B: CEX

	Laspeyres		Paasche	
	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$
$ \bar{\pi} $	0.38** (0.17)	0.18*** (0.052)	0.40** (0.16)	0.17* (0.094)
$ \mathbb{E}(\bar{\pi}) $		-1.31*** (0.20)		-1.0*** (0.23)
N	83	60	87	60

Note: Newey-west HAC standard errors in Parentheses. Significance at the one, five and ten percent levels indicated by ***,**,and * respectively. Aggregate Inflation, $|\bar{p}_i|$ is a democratic index of individual inflation weighted using the population weights in the Nielsen Homescan or CEX. Expected inflation, $|\mathbb{E}(\bar{\pi})|$, is the difference between the yield on a 5-year treasury bond and the corresponding 5-year Inflation Protected bond (TIPs). Nielsen Homescan Data from 2004-2017, CEX from 1996-2017, TIPs Expected Inflation 2003-2017. National prices are used throughout.

the business cycle. In order to ensure that the pattern I find is related to inflation and inflation dispersion rather than simply a story of the cyclical pattern of inflation dispersion, I repeat the empirical exercise in tables 3.4, but I include the National unemployment rate as a control. Table 3.6 shows the results of this robustness check (with the Nielsen Data). Even when controlling for the business cycle, the relationship between inflation dispersion and the absolute value of inflation still holds.

Next, I investigate whether the relationship I find between inflation and inflation dispersion is driven primarily by differences in the sectoral composition of household purchases

Table 3.5. Regional Household Inflation Dispersion and Regional Inflation

	Laspeyres $\sigma^r(\pi_h)$	Paasche $\sigma^r(\pi_h)$	Fisher $\sigma^r(\pi_h)$	Sato-Vartia $\sigma^r(\pi_h)$
$ \bar{\pi}_r $	0.544*** (0.0260)	1.575*** (0.171)	0.516*** (0.0620)	4.020*** (0.466)
Region FE	X	X	X	X
Time FE	X	X	X	X
Observations	10,656	10,656	10,656	10,656
R-squared	0.568	0.459	0.453	0.791

Note: Standard errors, in parentheses, two-way clustered at region and time levels and are robust to auto-correlation. Significance at the one percent level indicated by ***. Aggregate Regional Inflation, $|\bar{p}_r|$ is a democratic index of individual inflation weighted using the population weights in the Nielsen Homescan. Regional prices are used throughout.

Table 3.6. Robustness Check with Unemployment Rate

	Laspeyres		Paasche		Fisher		Sato-Vartia	
	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$	$\sigma(\pi_h)$
$ \bar{\pi} $	0.45*** (0.15)	0.37** (0.15)	0.86*** (0.11)	0.85*** (0.11)	0.67*** (0.21)	0.66*** (0.22)	5.94*** (1.16)	4.99*** (1.09)
$ \mathbb{E}(\bar{\pi}) $		1.40 (1.32)		0.83 (0.95)		0.73 (0.50)		0.34 (0.40)
UR	0.011 (0.23)	0.097 (0.20)	0.082 (0.20)	0.14 (0.15)	-0.060 (0.12)	-0.0068 (0.11)	0.098* (0.053)	0.11** (0.042)
N	52	52	52	52	52	52	52	52

Note: Newey-west HAC standard errors in Parentheses. Significance at the one, five and ten percent levels indicated by ***, **, and * respectively. Aggregate Inflation, $|\bar{p}_i|$ is a democratic index of individual inflation weighted using the population weights in the Nielsen Homescan or CEX. Expected inflation, $|\mathbb{E}(\bar{\pi})|$, is the difference between the yield on a 5-year treasury bond and the corresponding 5-year Inflation Protected bond (TIPs). Nielsen Homescan Data from 2004-2017, CEX from 1996-2017, TIPs Expected Inflation 2003-2017. National prices are used throughout.


Table 3.7. Module level Household Inflation Dispersion and Module Inflation

	Laspeyres $\sigma^r(\pi_h)$	Paasche $\sigma^r(\pi_h)$	Fisher $\sigma^r(\pi_h)$	Sato-Vartia $\sigma^r(\pi_h)$
$ \bar{\pi}_t $	4.048*** (1.199)	0.799*** (0.0716)	0.526*** (0.0393)	1.099*** (0.0393)
Time FE	X	X	X	X
Product Module FE	X	X	X	X
R-squared	0.732	0.592	0.505	0.838


Note: Standard errors, in parentheses, two-way clustered at module and time levels and are robust to auto-correlation. Significance at the one percent level indicated by ***. Aggregate Regional Inflation, $|\bar{p}_t^r|$ is a democratic index of individual inflation weighted using the population weights in the Nielsen Homescan. Regional prices are used throughout.

(i.e. some households spend more money on education a sticky sector than others) or a more fundamental element of price setting behavior that also exists within sectors (i.e. difference in price setting patterns of the individual education goods that the individual purchases). To do this, I exploit the fact that the Nielsen Homescan data has information on the actual bar-code of the product the household purchases along with the group (product module in the Nielsen data) to which that product belongs (the Nielsen inflation measure that I use in the rest of the paper is based on purchases in the broader product module category). I construct quarterly individual specific inflation rates for each of the 1,235 product modules in the Nielsen data. I then compare the standard deviation of the inflation rate in each of these product modules compared to the average inflation rate in the category. The results of this analysis is shown in table 3.7; I find a strong correlation between inflation and inflation dispersion even at the very narrow category level!

In order to ensure that this relationship is not only the product of trends in the late 1990s and 2000s, I repeat a similar analysis using city level CPI data going back until 1915. I cannot calculate individual inflation rates that far back, but to the extent that cities have differences in consumption baskets then one may expect to see a similar relationship between city level



(a) CPI-W 1915-2018



(b) CPI-U 1978-2018

Figure 3.6. Inflation and SD of City Inflation Rates

Note: SD is the standard deviation of city-level inflation rates. From city-level CPI data from the BLS. The 95 percent confidence interval is the shaded area.

inflation dispersion and aggregate inflation as that seen for individuals. The results from this analysis are shown in Figure 3.6. I note that there is a relationship between the variance of city inflation rates and aggregate absolute inflation going back until 1915.

3.5 Menu-cost model

I need a model that generates a positive relationship between absolute inflation and inflation dispersion, but does not introduce a relationship between inflation and price dispersion (like in the New Keynesian model). In this section I present a partial-equilibrium menu cost model that reproduces the link between aggregate inflation and household inflation dispersion that I find in the data. In addition to heterogeneous household preferences, a key assumption in this model is that the elasticity of substitution for a product differs across sectors⁶. This leads to the firm level profit function being more convex with respect to prices in some sectors than in others, which results in sector level differences in price-sensitivity to the aggregate price level

⁶If the elasticity of substitution also differs between higher and lower quality products within a sector then this could explain the results found in appendix F Shampoo prices where prices change differently between higher and lower quality goods, as well as the module level results found in table 3.7

(Barro 1972). The model takes the path of the aggregate price-level P_t as given and then models the firm-level price responses and the corresponding distribution of household level inflation rates.

3.5.1 Households

Households are modeled simply. There are N households with heterogeneous preferences across sectors and common CES preferences within each sector. Households supply one unit of labor inelastically and consume their entire income each period.

Formally, households have Cobb-Douglas utility with heterogeneous preferences across sectors m so that the utility for household h is given by:

$$u_t(h) = \prod_{m=1}^M C_t(m)^{\alpha_{h,m}}, \quad (3.5.1)$$

where $C_t(m)$ is an aggregated product of purchases from sector m defined implicitly by:

$$C_t(m) = \left(\int_0^1 c_t(z_m)^{\frac{\theta_m-1}{\theta_m}} \right)^{\frac{\theta_m}{\theta_m-1}}. \quad (3.5.2)$$

this implies a sector specific price index that is common for all households and a household specific price index across sectors:

$$P_{h,t} = \prod_{m=1}^M \left(\frac{p_{m,t}}{\alpha_{h,m}} \right)^{\alpha_{h,m}}. \quad (3.5.3)$$

A household's inflation rate at time t is then:

$$\pi_{h,t} = \sum_{m=1}^M \alpha_{h,m} \log \left(\frac{p_{m,t}}{p_{m,t-1}} \right). \quad (3.5.4)$$

Cobb-douglas sectorial utility was chosen so that the final inflation rate equation (3.5.4) would correspond exactly to the fixed weight Laspeyres and Paasche measures that I examine empirically.

I take the aggregate price level as given and make the simplifying assumption that total household demand for firm z 's product aggregates to:

$$c_t(z_m) = C_t \left(\frac{p_t(z_m)}{P_t} \right)^{-\theta_m} \quad (3.5.5)$$

where P_t is the exogenous aggregate price level and C_t is aggregate consumption, which is fixed in real terms and normalized to one.

3.5.2 Firms

There are M sectors and in each sector m , there is a continuum of monopolistic firms each producing a differentiated good. Firm z sets its price $p_t(z)$ and then produces consumption good $y(z)$ according to the following production function:

$$y_t(z) = A_t(z)L_t(z) \quad (3.5.6)$$

where $A_t(z)$ is firm z 's total factor productivity at time t and $L_t(z)$ is the quantity of labor employed.

The per-period firm profit function is:

$$\Pi_t(z) = p_t(z)y_t(z) - W_t L_t(z) - \chi_j W_t I_t(z). \quad (3.5.7)$$

here W_t is the wage rate, $I_t(z)$ is an indicator equal to one if the firm changes their price in period t and χ_j is the proportion/multiple of the prevailing wage rate that the firm must spend to change their price.

Demand for firm z 's product comes from equation (3.5.5). Combining firm demand with the firm's profit function yields a maximization problem with one choice variable: the product price.

Recursively, the firm's problem becomes:

$$V\left(A_t(z), \frac{p_{t-1}(z)}{P_t}\right) = \max_{p_t(z)} \left(\Pi_t(z)^R + E_t \left[D_{t,t+1}^R V\left(A_{t+1}(z), \frac{p_t(z)}{P_{t+1}}\right) \right] \right), \quad (3.5.8)$$

where $D_{t,t+1}$ is the firm's discount factor and the R superscript denotes real valued.

The firm problem differs for each firm because of different TFP values $A_t(z)$ and starting prices p_{t-1} , as well as a different curvature of the profit function because of varying elasticities of substitution for their products.

3.5.3 Simulation

I test whether this model is able to explain the relationship between the absolute value of aggregate inflation and inflation dispersion that I find in the data by simulating the model with 10 sectors. I use value function iteration to solve for the value function (and corresponding policy function) for firms in each sector. I then simulate the firm's responses to changes in TFP and the aggregate price level for 250,000 periods. In order to focus on changes in aggregate inflation, I assume that the TFP shocks are common across all firms. Finally, I construct price indexes for

Table 3.8. Model Simulation Results

	$\sigma(\pi_h)$	$\sigma(p(z,t))$
$ \bar{\pi} $	0.158*** (0.0003)	0.0001 (0.000015)

1000 households by assigning each household random uniformly distributed preference weights for each sector. I assume that each sector has elasticities of substitution ranging from 2 to 26 spaced evenly. I calibrate the rest of the model following Nakamura and Steinsson (2010).

I use the model results to regress the standard deviation of household level inflation rates on the absolute value of democratic aggregate inflation, which exactly mirrors my empirical exercise. The results of this experiment are shown in table 3.8. I find that a one percentage point increase in the absolute value of aggregate inflation is related to a 0.158 increase in the standard deviation of household inflation rates, which is around the mid-range of my empirical estimates. I also check whether there is a relationship between inflation and price dispersion in my model, column 2 shows the results; while the relationship between inflation and price dispersion is statistically significant, it is extremely small.

3.6 Conclusion

This project enlisted two large household-level datasets (the Consumer Expenditure Survey and the Nielsen Consumer Panel) to calculate household level inflation rates. I have found that household level inflation dispersion is costly: (1) shocks to a household's price index are persistent and are permanent in the case of Paasche and Sato-Vartia inflation; (2) households appear to reduce both their nominal and real consumption following shocks to their inflation rate; (3) poorer households are less able to smooth their consumption following shocks to their inflation rate. I also found that there is a robust relationship between aggregate inflation

and inflation dispersion distinct from the business cycle, which can be explained by a partial equilibrium menu-cost model where the elasticity of substitution (and therefore the curvature of the firm's profit function) varies across sectors.

Inflation is costly because not everyone experiences inflation in the same way and increased inflation expands the cross-sectional volatility of household-level inflation. This is a cost that cannot be fixed by indexing paychecks, loans, etc. to the aggregate inflation rate. The increase in volatility can be especially costly for low-asset households that cannot smooth their consumption, so the increase in inflation volatility is matched by an increase in real consumption volatility.

Also, given the wide distribution of household inflation rates it should come at no surprise that expectations of future inflation rates in the Michigan survey and others are so widely dispersed (Mankiw et al. 2003). While some have speculated this may be due to the inattention or sticky information on the part of respondents (Carroll 2003, Mankiw et al. 2003), households may be simply responding based rational beliefs that their inflation rates may not conform to projected aggregate rates.

This research is particularly pertinent to the current debate over whether central banks should increase their inflation targets from around 2 percent to a higher level (the ECB is currently considering this proposal in their strategy review). Several papers have considered the welfare costs of increasing an inflation target in the US from 2 to 4 percent (Ascari et al. 2018, Ascari 2004, Amano et al. 2007) and find that this would lead to a consumption equivalent welfare loss between 0.25% and 4 %. Given the relationship between inflation dispersion and inflation that I find, an unexpected increase in the inflation target should lead to an additional welfare loss due to increased cross-sectional volatility of individual inflation rates.

3.7 Chapter Acknowledgements

Chapter 3 is calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

Bibliography

- Aguiar, M. and Bilal, M. (2015). Has consumption inequality mirrored income inequality? *American Economic Review*, 105(9):2725–56.
- Alvarez, F., Beraja, M., Gonzalez-Rozada, M., and Neumeyer, P. A. (2018). From hyperinflation to stable prices: Argentina’s evidence on menu cost models. *The Quarterly Journal of Economics*, 134(1):451–505.
- Amano, R., Ambler, S., and Rebei, N. (2007). The macroeconomic effects of nonzero trend inflation. *Journal of Money, Credit and Banking*, 39(7):1821–1838.
- Amble, N. and Stewart, K. (1994). Experimental price index for elderly consumers. *Monthly Lab. Rev.*, 117:11.
- Argente, D. and Lee, M. (2017). Cost of living inequality during the great recession.
- Argente, D. and Lee, M. (2021). Cost of living inequality during the Great Recession. *Journal of the European Economic Association*, 19(2):913–952.
- Ascari, G. (2004). Staggered prices and trend inflation: some nuisances. *Review of Economic Dynamics*, 3(7):642–667.
- Ascari, G., Phaneuf, L., and Sims, E. R. (2018). On the welfare and cyclical implications of moderate trend inflation. *Journal of Monetary Economics*, 99:56–71.
- Atkin, D., Faber, B., Fally, T., and Gonzalez-Navarro, M. (2020). Measuring welfare and inequality with incomplete price information. *National Bureau of Economic Research*.
- Attanasio, O. P. and Weber, G. (1995). Is consumption growth consistent with intertemporal optimization? evidence from the consumer expenditure survey. *Journal of political Economy*, 103(6):1121–1157.
- Auclert, A., Bardóczy, B., and Rognlie, M. (2020). Mpcs, mpes and multipliers: A trilemma for new keynesian models. Technical report, National Bureau of Economic Research.

- Auclert, A., Bardóczy, B., and Rognlie, M. (forthcoming). Mpc, mpes and multipliers: A trilemma for new keynesian models. *Review of Economics and Statistics*.
- Auclert, A., Rognlie, M., and Straub, L. (2018). The intertemporal keynesian cross. Technical report, National Bureau of Economic Research.
- Barro, R. J. (1972). A theory of monopolistic price adjustment. *The Review of Economic Studies*, 39(1):17–26.
- Barsky, R. B., House, C. L., and Kimball, M. S. (2007). Sticky-price models and durable goods. *American Economic Review*, 97(3):984–998.
- Barth III, M. J. and Ramey, V. A. (2001). The cost channel of monetary transmission. *NBER macroeconomics annual*, 16:199–240.
- Bauer, M. D. and Swanson, E. T. (2020). The fed’s response to economic news explains the” fed information effect”. Technical report, National Bureau of Economic Research.
- Berger, D. and Vavra, J. (2015). Consumption dynamics during recessions. *Econometrica*, 83(1):101–154.
- Bils, M. and Klenow, P. J. (1998). Using consumer theory to test competing business cycles models. *Journal of Political Economy*, 106(2):233–261.
- BLS (2020). Handbook of methods: Consumer price index.
- Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196.
- Borusyak, K. and Jaravel, X. (2017). Revisiting event study designs. *Available at SSRN 2826228*.
- Borusyak, K., Jaravel, X., and Spiess, J. (2022). Revisiting event study designs: Robust and efficient estimation. *arXiv preprint arXiv:2108.12419*.
- Broda, C. and Parker, J. A. (2014). The economic stimulus payments of 2008 and the aggregate demand for consumption. *Journal of Monetary Economics*, 68:S20–S36.
- Broer, T., Harbo Hansen, N.-J., Krusell, P., and Öberg, E. (2020). The new keynesian transmission mechanism: A heterogeneous-agent perspective. *The Review of Economic Studies*, 87(1):77–101.
- Browning, M. (2005). A working paper from april 1985: Which demand elasticities do we know and which do we need to know for policy analysis? *Research in Economics*, 59(4):293–320.

- Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *the Quarterly Journal of economics*, 118(1):269–298.
- Cavallo, A. (2020). Inflation with covid consumption baskets. *National Bureau of Economic Research*.
- Cloyne, J. S., Jordà, Ò., and Taylor, A. M. (2020). Decomposing the fiscal multiplier. *National Bureau of Economic Research*.
- Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2017). Innocent bystanders? monetary policy and inequality. *Journal of Monetary Economics*, 88:70–89.
- Coibion, O., Gorodnichenko, Y., and Wieland, J. (2012). The optimal inflation rate in new keynesian models: should central banks raise their inflation targets in light of the zero lower bound? *Review of Economic Studies*, 79(4):1371–1406.
- Colciago, A. (2011). Rule-of-thumb consumers meet sticky wages. *Journal of money, credit and banking*, 43(2-3):325–353.
- Comin, D., Lashkari, D., and Mestieri, M. (2021). Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374.
- Comin, D. A., Danieli, A., and Mestieri, M. (2020). Income-driven labor market polarization. *National Bureau of Economic Research*.
- Cravino, J., Lan, T., and Levchenko, A. A. (2018). Price stickiness along the income distribution and the effects of monetary policy. *NBER Working Papers*.
- Cravino, J., Lan, T., and Levchenko, A. A. (2020). Price stickiness along the income distribution and the effects of monetary policy. *Journal of Monetary Economics*, 110:19–32.
- Cravino, J. and Levchenko, A. A. (2017). The distributional consequences of large devaluations. *American Economic Review*, 107(11):3477–3509.
- Crawford, I. and Oldfield, Z. (2002). *Distributional aspects of inflation*. Citeseer.
- Crossley, T. F. and Low, H. W. (2011). Is the elasticity of intertemporal substitution constant? *Journal of the European Economic Association*, 9(1):87–105.
- De Chaisemartin, C. and d’Haultfoeuille, X. (2015). Supplement to fuzzy differences-in-differences. Technical report.
- De Chaisemartin, C. and d’Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–96.

- Deaton, A. and Muellbauer, J. (1980). An almost ideal demand system. *The American economic review*, 70(3):312–326.
- Diercks, A. M. (2017). The reader’s guide to optimal monetary policy.
- Doepke, M. and Schneider, M. (2006). Inflation and the redistribution of nominal wealth. *Journal of Political Economy*, 114(6):1069–1097.
- Edelstein, P. and Kilian, L. (2009). How sensitive are consumer expenditures to retail energy prices? *Journal of Monetary Economics*, 56(6):766–779.
- Einav, L., Leibtag, E., and Nevo, A. (2010). Recording discrepancies in nielsen homescan data: Are they present and do they matter? *QME*, 8(2):207–239.
- Erceg, C. J., Henderson, D. W., and Levin, A. T. (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of monetary Economics*, 46(2):281–313.
- Fajgelbaum, P. D. and Khandelwal, A. K. (2016). Measuring the unequal gains from trade. *The Quarterly Journal of Economics*, 131(3):1113–1180.
- Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, pages 157–177.
- Feenstra, R. C. and Weinstein, D. E. (2017). Globalization, markups, and us welfare. *Journal of Political Economy*, 125(4):1040–1074.
- Feiveson, L., Gornemann, N., Hotchkiss, J. L., Mertens, K., and Sim, J. (2020). Distributional considerations for monetary policy strategy.
- Feldstein, M. (2008). The tax rebate was a flop; obama’s stimulus plan won’t work either. *The Wall Street Journal*, August 8.
- Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco.
- Galí, J. (2015). *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*. Princeton University Press.
- Galí, J. and Gambetti, L. (2019). Has the us wage phillips curve flattened? a semi-structural exploration. *National Bureau of Economic Research*.
- Galí, J., López-Salido, D., and Vallés, J. (2007). Understanding the effects of government spending on consumption. *Journal of the European Economic Association*, 5(1):227–270.

- Garner, T. I., Johnson, D. S., and Kokoski, M. F. (1996). An experimental consumer price index for the poor. *Monthly Lab. Rev.*, 119:32.
- Gelman, M., Gorodnichenko, Y., Kariv, S., Koustas, D., Shapiro, M. D., Silverman, D., and Tadelis, S. (2016). The response of consumer spending to changes in gasoline prices. Technical report, National Bureau of Economic Research.
- Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American economic review*, 102(4):1692–1720.
- Golosov, M. and Lucas Jr, R. E. (2007). Menu costs and phillips curves. *Journal of Political Economy*, 115(2):171–199.
- Goolsbee, A. (1998). Investment tax incentives, prices, and the supply of capital goods. *The Quarterly Journal of Economics*, 113(1):121–148.
- Gürer, E. and Weichenrieder, A. (2020). Pro-rich inflation in europe: Implications for the measurement of inequality. *German Economic Review*, 21(1):107–138.
- Gürkaynak, R. S., Sack, B. P., and Swanson, E. T. (2004). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *The Response of Asset Prices to Monetary Policy Actions and Statements (November 2004)*.
- Hamilton, J. D. (2018). Why you should never use the hodrick-prescott filter. *Review of Economics and Statistics*, 100(5):831–843.
- Heathcote, J., Perri, F., and Violante, G. L. (2020). The rise of us earnings inequality: Does the cycle drive the trend? *Review of economic dynamics*, 37:S181–S204.
- Herbst, E. and Johannsen, B. K. (2021). Bias in local projections.
- Hobijn, B. and Lagakos, D. (2005). Inflation inequality in the united states. *Review of income and Wealth*, 51(4):581–606.
- Hottman, C. J. and Monarch, R. (2020). A matter of taste: Estimating import price inflation across us income groups. *Journal of International Economics*, 127:103382.
- House, C. L. and Shapiro, M. D. (2008). Temporary investment tax incentives: Theory with evidence from bonus depreciation. *American Economic Review*, 98(3):737–68.
- Hoynes, H., Miller, D. L., and Schaller, J. (2012). Who suffers during recessions? *Journal of Economic perspectives*, 26(3):27–48.
- Jaimovich, N., Rebelo, S., and Wong, A. (2019). Trading down and the business cycle. *Journal*

- of Monetary Economics*, 102:96–121.
- Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the us retail sector. *The Quarterly Journal of Economics*, 134(2):715–783.
- Jaravel, X. (2021). Inflation inequality: Measurement, causes, and policy implications. *Annual Review of Economics*, 13.
- Jaravel, X. and O’Connell, M. (2020). Real-time price indices: Inflation spike and falling product variety during the great lockdown. *Journal of Public Economics*, 191:104270.
- Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5):1589–1610.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, 95(1):161–182.
- Kaplan, G. and Menzio, G. (2015). The morphology of price dispersion. *International Economic Review*, 56(4):1165–1206.
- Kaplan, G., Moll, B., and Violante, G. L. (2018a). Monetary policy according to HANK. *American Economic Review*, Forthcoming.
- Kaplan, G., Moll, B., and Violante, G. L. (2018b). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.
- Kaplan, G. and Schulhofer-Wohl, S. (2017). Inflation at the household level. *Journal of Monetary Economics*, 91:19–38.
- Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.
- Kim, D.-H. and Lin, S.-C. (2012). Inflation and inflation volatility revisited. *International finance*, 15(3):327–345.
- Kim, S. (2019). Quality, price stickiness, and monetary policy. *Journal of Macroeconomics*, page 103129.
- Krueger, D., Mitman, K., and Perri, F. (2016). On the distribution of the welfare losses of large recessions. *National Bureau of Economic Research*.
- Lauper, C. and Mangiante, G. (2021). Monetary policy shocks and inflation inequality.
- Leahy, J. V. and Thapar, A. (2019). Demographic effects on the impact of monetary policy.

National Bureau of Economic Research.

- Li, M. (2019). Countercyclical dispersion of price changes: Causes and implications.
- Mankiw, N. G., Reis, R., and Wolfers, J. (2003). Disagreement about inflation expectations. *NBER macroeconomics annual*, 18:209–248.
- McCarthy, P. S. (1996). Market price and income elasticities of new vehicle demands. *The Review of Economics and Statistics*, pages 543–547.
- McGranahan, L. and Paulson, A. L. (2005). Constructing the Chicago Fed income based economic index-consumer price index: Inflation experiences by demographic group: 1983-2005.
- McKay, A. and Wieland, J. F. (2021). Lumpy durable consumption demand and the limited ammunition of monetary policy. *Econometrica*.
- McKay, A. and Wieland, J. F. (2022). Forward guidance and durable goods demand. *American Economic Review: Insights*, 4(1):106–22.
- McLeay, M. and Tenreyro, S. (2019). Optimal inflation and the identification of the Phillips curve. Technical report, National Bureau of Economic Research.
- Meyer, B. D. and Sullivan, J. X. (2013). Consumption and income inequality and the great recession. *American Economic Review*, 103(3):178–83.
- Misra, K. and Surico, P. (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics*, 6(4):84–106.
- Montiel Olea, J. L. and Plagborg-Møller, M. (2021). Local projection inference is simpler and more robust than you think. *Econometrica*, 89(4):1789–1823.
- Muellbauer, J. (1975). Aggregation, income distribution and consumer demand. *The Review of Economic Studies*, 42(4):525–543.
- Nakamura, E. and Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4):1415–1464.
- Nakamura, E. and Steinsson, J. (2010). Monetary non-neutrality in a multisector menu cost model. *The Quarterly Journal of Economics*, 125(3):961–1013.
- Nakamura, E., Steinsson, J., Sun, P., and Villar, D. (2018). The elusive costs of inflation: Price dispersion during the US great inflation. *The Quarterly Journal of Economics*, 133(4):1933–1980.

- Orchard, J. (2021). Household inflation and aggregate inflation. *Working Paper*.
- Pakko, M. R. (1998). Shoe-leather costs of inflation and policy credibility. *Federal Reserve Bank of St. Louis Review*, 80(6):37.
- Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2011). Consumer spending and the economic stimulus payments of 2008. Technical report, National Bureau of Economic Research.
- Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, 103(6):2530–53.
- Ramey, V. (2020). The macroeconomic consequences of infrastructure investment. Technical report, National Bureau of Economic Research.
- Ramey, V. A. (2016). Macroeconomic shocks and their propagation. *Handbook of macroeconomics*, 2:71–162.
- Ramey, V. A. (2021a). Comment. *NBER Macroeconomics Annual*, 35:232–241.
- Ramey, V. A. (2021b). The macroeconomic consequences of infrastructure investment. In Glaeser, E. L. and Poterba., J. M., editors, *Economic Analysis and Infrastructure Investment*, page Chapter 4. Chicago: University of Chicago Press.
- Redding, S. J. and Weinstein, D. E. (2018). Measuring aggregate price indexes with demand shocks: Theory and evidence for CES preferences.
- Sahm, C. R., Shapiro, M. D., and Slemrod, J. (2010). Household response to the 2008 tax rebate: Survey evidence and aggregate implications. *Tax Policy and the Economy*, 24(1):69–110.
- Sahm, C. R., Shapiro, M. D., and Slemrod, J. (2012). Check in the mail or more in the paycheck: does the effectiveness of fiscal stimulus depend on how it is delivered? *American Economic Journal: Economic Policy*, 4(3):216–50.
- Schettkat, R. and Yocarini, L. (2006). The shift to services employment: A review of the literature. *Structural change and economic dynamics*, 17(2):127–147.
- Schmitt-Grohé, S. and Uribe, M. (2003). Closing small open economy models. *Journal of international Economics*, 61(1):163–185.
- Shapiro, M. D. and Slemrod, J. (2003). Consumer response to tax rebates. *American Economic Review*, 93(1):381–396.
- Shapiro, M. D. and Slemrod, J. (2009). Did the 2008 tax rebates stimulate spending? *American*

Economic Review, 99(2):374–79.

Sun, L. and Abraham, S. (2020). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*.

Swanson, E. T. (2021). Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics*, 118:32–53.

Taylor, J. B. (2009). The lack of an empirical rationale for a revival of discretionary fiscal policy. *American Economic Review*, 99(2):550–55.

Vavra, J. (2013). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *The Quarterly Journal of Economics*, 129(1):215–258.

Wilcox, D. W. (1992). The construction of us consumption data: some facts and their implications for empirical work. *The American economic review*, pages 922–941.

Wolf, C. K. (2021). The missing intercept: A demand equivalence approach. Technical report, National Bureau of Economic Research.

Appendix: Chapter 1

A1 Mathematical Appendix

A1.1 Proof of Proposition 1

Lemma 1 *If $F(H) : [0, \infty) \rightarrow [0, \infty)$ is homogeneous of degree $k \in (0, 1)$ then $\frac{\partial \frac{F'(H_j)}{F'(H_i)}}{\partial \frac{F(H_i)}{F(H_j)}} > 0$.*

First I show that a function that is homogeneous of degree $k \in (0, 1)$ is strictly increasing.

Suppose $H_i > H_j$ then:

$$F(H_i) = H_i^k F(1) > H_j^k F(1) = F(H_j)$$

For notational convenience, let $Y_i := F(H_i)$. By Euler's Homogeneous Function Theorem, $F(H_i) = F'(H_i)H_i$, which implies that:

$$\begin{aligned} \frac{F'(H_j)}{F'(H_i)} &= \frac{Y_j}{Y_i} \left(\frac{H_i}{H_j} \right) \\ &= \frac{Y_j}{Y_i} \left(\frac{F^{-1}(Y_i)}{F^{-1}(Y_j)} \right), \end{aligned}$$

where the inverse function must exist since F is strictly increasing. Next, I take the derivative with respect to the output ratio:

$$\frac{\partial}{\partial \frac{F(H_i)}{F(H_j)}} \frac{F'(H_j)}{F'(H_i)} = \frac{Y_j}{Y_i} \frac{\partial}{\partial \frac{F(H_i)}{F(H_j)}} \left(\frac{F^{-1}(Y_i)}{F^{-1}(Y_j)} \right) - \frac{F^{-1}(Y_i)}{F^{-1}(Y_j)} \quad (\text{A1.1})$$

Since the inverse of a homogeneous function of degree k , is a homogeneous function of degree $1/k$ it follows that:

$$\frac{\partial}{\partial \frac{Y_i}{Y_j}} \left(\frac{F^{-1}(Y_i)}{F^{-1}(Y_j)} \right) = \frac{\partial}{\partial \frac{Y_i}{Y_j}} \left(\left(\frac{Y_i}{Y_j} \right)^{1/k} \frac{F^{-1}(1)}{F^{-1}(1)} \right) \quad (\text{A1.2})$$

$$= \frac{1}{k} \left(\frac{Y_i}{Y_j} \right)^{(1-k)/k} . \quad (\text{A1.3})$$

By substituting equation (A1.3) into equation (A1.1) I find that:

$$\begin{aligned} \frac{\partial}{\partial \frac{F(H_i)}{F(H_j)}} \frac{F'(H_j)}{F'(H_i)} &= \frac{Y_j}{Y_i} \frac{1}{k} \left(\frac{Y_i}{Y_j} \right)^{(1-k)/k} - \left(\frac{Y_i}{Y_j} \right)^{1/k} \\ &= \left(\frac{Y_i}{Y_j} \right)^{1/k} \left(\frac{1}{k} - 1 \right), \end{aligned}$$

which is > 0 if and only if $k < 1$. ■

Corollary 1 *If $F(H) : [0, \infty) \rightarrow [0, \infty)$ and $G(H) : [0, \infty) \rightarrow [0, \infty)$ are both homogeneous of degree $k \in (0, 1)$ then $\frac{\partial \frac{G'(H_j)}{F'(H_i)}}{\partial \frac{F(H_i)}{G(H_j)}} > 0$.*

This proof follows from the proof above, except replace $\frac{F^{-1}(1)}{F^{-1}(1)}$ in equation (A1.2) with $\frac{F^{-1}(1)}{G^{-1}(1)}$, which implies that:

$$= \frac{F^{-1}(1)}{G^{-1}(1)} \left(\frac{Y_i}{Y_j} \right)^{1/k} \left(\frac{1}{k} - 1 \right),$$

■

Proposition 1 *In a two-sector competitive economy with a representative household that has preferences satisfying equation (1.4.4), production function in each sector $F_i(H_i) : [0, \infty) \rightarrow [0, \infty)$ both homogeneous of degree $k \in (0, 1)$ and standard market clearing conditions, then an decrease/increase in household expenditure will lead to an increase/decrease in the relative price of necessities.*

Due to market clearing, it follows that

$$C^i(X, p_N, p_L) = F_i(H_i) \quad \forall i$$

From equation (1.4.4) we know that

$$\frac{\partial C^L(X, p_N, p_L)}{\partial X C^N(X, p_N, p_L)} > 0.$$

This implies that:

$$\frac{\partial F_L(H_L)}{\partial X F_N(H_N)} = \frac{\partial Y_L}{\partial X Y_N} > 0. \quad (\text{A1.4})$$

Relative prices can be expressed as:

$$\frac{p_L}{p_N} = \frac{F_{N,H}(H_N)}{F_{L,H}(H_L)}$$

From lemma and corollary 1, we get that:

$$\frac{\partial F_{N,H}(H_N)}{\partial \frac{Y_L}{Y_N} F_{L,H}(H_L)} > 0. \quad (\text{A1.5})$$

Combining equation (A1.4) with equation (A1.5) and the chain-rule implies that:

$$\frac{\partial p_L}{\partial X p_N} > 0$$

So the price of the expanding sector (luxuries in this case) must increase. ■

A1.2 Derivation of Wage-Phillips Curve

I add sticky wages by following the convention in the literature and creating market power in the labor market via a labor union (see Erceg et. al. 2000, Auclert et. al. 2018, Auclert et. al. 2020, Broer et. al. 2020, Ramey 2020).

Specifically, each worker (i) in the economy provides h_{ikt} hours of labor to each of a continuum of unions indexed by $k \in (0, 1)$. Total labor for person (i) is then:

$$h_{it} = \int_k h_{ikt} dk. \quad (\text{A1.6})$$

Each union k aggregates units of work into a union specific task $H_{kt} = \int_i h_{ikt} di$.

There is a competitive labor packer that takes labor from unions and packages it into one unit of “usable” labor following a CES function. Aggregate labor is then:

$$H_t = \left(\int_k H_{kt}^{\frac{\epsilon_w - 1}{\epsilon_w}} \right)^{\epsilon_w / (\epsilon_w - 1)}, \quad (\text{A1.7})$$

where ε_w is the elasticity of substitution between different types of labor.

Unions set a common wage w_{kt} for all members and require each member household to supply uniform hours: $h_{ikt} = H_{kt}$.

Following (Auclert et al. 2018,2020) I add an extra disutility term for households, so that households dislike adjusting wages:

$$\frac{\Psi_w}{2} \int_k \left(\frac{w_{kt}}{w_{kt-1}} - 1 \right)^2 dk, \quad (\text{A1.8})$$

where Ψ_w scales the degree of wage stickiness.

At time t, union k sets wage w_{kt} to maximize (on behalf of all union workers):

$$\begin{aligned} \max_{w_{kt}} \mathbb{E}_t \sum_{\tau > 0} \beta^{t+\tau} \left(\int [V(X_{it+\tau}, \mathbf{p}_{t+\tau}) - g(h_{i,t+\tau})] d\psi_{it+\tau} - \frac{\Psi_w}{2} \int_k \left(\frac{w_{kt}}{w_{kt-1}} - 1 \right)^2 dk \right) \\ \text{s.t. } H_{kt} = \left(\frac{w_{kt}}{W_t} \right)^{-\varepsilon_w} H_t \end{aligned} \quad (\text{A1.9})$$

The union takes as given the distribution ψ_{it} of workers (in this version of the model, all workers are identical) and all prices excluding w_{kt} (note that $W_t = \left(\int_k w_{kt}^{1-\varepsilon_w} dk \right)^{1/(1-\varepsilon_w)}$.)

The envelope theorem allows me to ignore both the intertemporal reoptimization of saving or spending in response to a marginal change in wages, along with the intratemporal reoptimization of spending across sectors. I treat any change in income as a change in consumption expenditure:

$$\begin{aligned}
\frac{\partial X_{it}}{\partial w_{kt}} &= \frac{\partial}{\partial w_{kt}} \int_0^1 w_{kt} h_{ikt} dk \\
&= \int_0^1 \frac{\partial}{\partial w_{kt}} w_{kt} \left(\frac{w_{kt}}{W_t} \right)^{-\varepsilon_w} H_t dk \\
&= (1 - \varepsilon_w) \left(\frac{w_{kt}}{W_t} \right)^{-\varepsilon_w}.
\end{aligned}$$

I next derive the change in hours worked to a change in wages for household (i) using the labor rule that $H_{kt} = h_{ikt} \forall i$ and the demand constraint:

$$\begin{aligned}
\frac{\partial h_{it}}{\partial w_{kt}} &= -\varepsilon_w \left(\frac{w_{kt}^{-\varepsilon_w - 1}}{W_t^{-\varepsilon_w}} \right) \\
&= -\varepsilon_w \frac{H_{kt}}{w_{kt}}.
\end{aligned}$$

It follows that the first order condition of the union's maximization problem equation (A1.9) becomes:

$$\begin{aligned}
\int H_{kt} \left[V_X(X_{it}, \mathbf{p}_t) (1 - \varepsilon_w) \left(\frac{w_{kt}}{W_t} \right)^{-\varepsilon_w} + \frac{\varepsilon_w}{w_{kt}} g'(h_{it}) \right] d\psi_{it} - \psi_w \left(\frac{w_{kt}}{w_{kt-1}} - 1 \right) \frac{1}{w_{kt-1}} \\
+ \beta \psi_w \mathbb{E}_t \left[\left(\frac{w_{k,t+1}}{w_{k,t}} - 1 \right) \left(\frac{w_{k,t+1}}{w_{kt}^2} \right) \right] = 0.
\end{aligned}$$

This simplifies when we note that the maximization problem for all unions is identical, so in equilibrium $w_{kt} = w_t$. Denoting $\pi_t^w \equiv \left(\frac{w_t}{w_{t-1}} - 1 \right)$ and using the functional forms for $V[\cdot]$ and $g(\cdot)$ provided in section 6 yields:

$$\psi_w \pi_t^w (1 + \pi_t^w) = \beta \mathbb{E}_t \left(\psi_w \pi_{t+1}^w (1 + \pi_{t+1}^w) \right) + H_t w_t \int \left[\frac{1}{a(\mathbf{p}_t) b(\mathbf{p}_t)} \left(\frac{X_t}{a(\mathbf{p}_t)} \right)^{((1-\eta)/b(\mathbf{p}_t))-1} (1 - \varepsilon_w) + \frac{\varepsilon_w}{W_t} \varphi H_{it}^\phi \right] d\psi_{it}.$$

In the representative agent model that I am considering here, this further simplifies to:

$$(1 + \pi_t^w) \pi_t^w = \beta \mathbb{E}_t \left[(1 + \pi_{t+1}^w) \pi_{t+1}^w \right] + \left(\frac{\varepsilon_w}{\psi_w} \right) \left(\varphi H_t^\phi - \left(\frac{\varepsilon_w - 1}{\varepsilon_w} \right) \frac{W_t}{a(\mathbf{p}_t) b(\mathbf{p}_t)} \left(\frac{X_t}{a(\mathbf{p}_t)} \right)^{((1-\eta)/b(\mathbf{p}_t))-1} \right) \quad (\text{A1.10})$$

It follows that the union will adjust wages in expectations of future wage inflation or when the marginal disutility of labor is higher than the product of marginal utility of expenditure and the optimal wage.

A1.3 A Note on Aggregation

In general, it is not true that if micro-households have non-homothetic preferences then the aggregate household will also have non-homothetic preferences of the same form. Very few types of non-homothetic preferences are Gorman-Polar (Stone-Geary is a notable exception), so these type of preferences cannot simply be added up across households to create an aggregate household with the same preference structure and parameters as the micro households (Muellbauer 1975).

Muellbauer (1975) shows that a necessary and sufficient condition for there to exist an income/expenditure level such that a representative household with that income/expenditure level to have preferences identical to the average of all households is that households must have Generalized Linear (GL) preferences. The expenditure/income of a slightly less general version of these preferences, Price Independent Generalized Linear (PIGL) is shown to depend positively

on both aggregate income/expenditure and the inequality of the income/expenditure distribution. Intuitively, this is because in a more unequal economy, all else equal, will have a higher portion of aggregate income/expenditure concentrated in a few hands, which means that more luxuries will be consumed. Hence, the representative household should have higher income/expenditure than that implied by the aggregate expenditure in the economy.

If the representative household proceeds to purchase relatively more necessity goods, then this will cause necessity prices to increase. Since poorer households have lower expenditure than rich households, these households will have a larger percentage of their basket devoted to the necessity good. This increase in necessity prices will increase their price index relative to rich households.

It has been documented that both recessions (Heathcote et al. 2020) and contractionary monetary policy (Coibion, Gorodnichenko, Kueng, Silvia 2018) increase inequality. Since demand for the necessity good depends both on aggregate expenditure (decreasing) and inequality (decreasing), a shock that simultaneously lowers aggregate expenditure and raises inequality would have ambiguous effects on relative necessity demand. To fix ideas, if representative expenditure x^r is a function $F(\cdot)$ of aggregate expenditure \bar{x} and expenditure inequality Σ_x then the elasticity of representative expenditure to a macroeconomic shock, $\mathcal{E}_{x^r, shock}$, would be:

$$\mathcal{E}_{x^r, shock} = \mathcal{E}_{x^r, \bar{x}} \mathcal{E}_{\bar{x}, shock} + \mathcal{E}_{x^r, \Sigma_x} \mathcal{E}_{\Sigma_x, shock}. \quad (\text{A1.11})$$

In equation (A1.11), the elasticity of representative expenditure to a shock depends both on the elasticity of aggregate expenditure to the shock and the elasticity of inequality to the shock, where each term is scaled by the elasticity of representative expenditure to either aggregate expenditure or inequality.⁷ In the empirical section, I show that following a monetary policy

⁷In the PIG-Log (AIDS) specification I adopt in the main text, the elasticity of x^r with respect to both aggregate expenditure and inequality (as measured by the Theil Index) is one, so equation (A1.11) reduces to just $\mathcal{E}_{\bar{x}, shock} + \mathcal{E}_{\Sigma_x, shock}$. Coibion et al. (2017) finds that the elasticity of the standard deviation of expenditure increases by .03

shock the effect coming through aggregate expenditure dominates.

A2 Data Appendix

A2.1 Cross-walk between CPI and CEX

The US Bureau of Labor Statistics (BLS) uses weights computed from the Consumer Expenditure Survey in calculating the official Consumer Price Index. In principle, this means that I could match each of the 243 item strata used to compute the CPI with corresponding consumption expenditures in the CEX. However, the BLS neither publishes the cross-walk between the CPI and the CEX, nor do they publish the price indices for each item strata. So, for this project I create my own cross-walk between the CEX and the publicly available price index series from the BLS. Given this crosswalk, I pull the CPI price series and the CEX data directly from the BLS website using their API.

I match expenditure in the CEX with prices in the CPI using the CEX UCC product hierarchy (available from the BLS) alongside the BLS CPI data finder. The goal is to create the most disaggregated product categories for which I have data in both the CPI and CEX. In general, the CEX has reported purchases at a more disaggregated level than the CPI. For example, the CPI price series “Women’s suits and separates” matches with 5 different UCC codes (for 2019) in the CEX. I aggregate UCC codes from the CEX so they match the more aggregated CPI series. In the cases where the CPI data was more disaggregated, e.g., types of gasoline, I choose a more aggregate CPI series series—e.g.,gasoline. Where CPI series only exist for a subset of years, I choose the most disaggregated series available for which prices are available over the latter part of the sample (since 2007).

There are many UCC codes in the CEX that are only available for certain years. For example, Women’s pants are available from 1990Q2-2007Q1. From, 2007-2019 Women’s pants

four months after a one-s.d. monetary policy shock, while consumption falls by approximately 0.5 percent. Given that the Theil Coefficient for a log-normal distribution is $\sigma^2/2$ it follows that the aggregate expenditure elasticity dominates the inequality elasticity.

are included in the more aggregated category Women’s pants and shorts. My final product categorization insures that the products represent the same breadth of spending in each year.

The complete cross-walk between the CEX UCC codes and the CPI price series is available from the author.

In addition to classifying products as necessities or luxuries, the crosswalk can also be used to construct income-level cost-of-living indexes. For example, figure A6 displays income level Laspeyres indices from 2007-2013:

$$P_t^I = \sum w_{Ij} p_{jt}$$

where the weights on each category for each income-group, w_{Ij} come from income level expenditure shares from 2005-06 in each of the 119 non-housing products in the crosswalk along with income level expenditure shares of rent and owners equivalent rent. The inflation gap in core-cpi from 2007q2-2009q3 between Low- and High-income households is 0.86 percentage points, almost exactly the 0.85 pp. gap over that same period in the model.

A3 Alternate Calibrations

As mentioned in the main text, I consider several alternative calibrations. I consider three different values for α ; (1) $\alpha = 0.366$ from Fernald (2014), (2) $\alpha = 0.3$, which is implied by letting the marginal elasticity of marginal cost to quantity supplied in the model equal the median estimated value in Hottman and Monarch (2020), and (3) $\alpha = 0.16$, which is implied by the median results for ω in Feenstra and Weinstein (2017). I also directly estimate β_L and γ_{LN} from the micro-data, and use these values. The method of estimation is described in the next subsection.

A3.1 Demand Parameter Estimation

I follow Deaton and Muellbauer (1980) and Fajgelbaum and Khandelwal (2016) when estimating the parameters in the AIDs. Specifically, I estimate equation (1.6.6) directly from the micro data by replacing $a(\mathbf{p})$ with a known price index (I use the CPI) so that the coefficient β_j represents changes in the share of expenditure on product j with changes in real expenditure, so that equation (1.6.6) becomes:

$$s_j = a_j^* + \sum_k \gamma_{jk} \log(p_k) + \beta_j (x_h^*). \quad (\text{A3.1})$$

Where x_h^* is real household expenditure, and a_j^* is a transformation of a^j .⁸ Since there are only two sectors, I can estimate equation (A3.1) directly via OLS by treating the price of one sector (necessities) as the numeraire and following the parameter restrictions defined earlier: $\sum_{j=1}^N a_j = 1$, $\sum_{j=1}^N \beta_j = \sum_{j=1}^N \gamma_{jk} = 0$ and $\gamma_{ij} = \gamma_{ji} \quad \forall i, j$. Similar to the rest of the analysis, I control for household size, age of the household head, and the number of wage earners. I use the full household sample (1991-2019) and define the necessity good as the composite good of products with relative expenditure ratio greater than one.

Results from this estimation are shown in table A.1. Column one reports the OLS results. I estimate that $\beta^N = -0.18$, which implies a luxury sector expenditure elasticity for the representative household of 1.4. I also estimate a positive cross-price elasticity, implying that necessities and luxuries are gross-complements. Column 2 shows an alternate estimation using household log-income and income quintiles as instruments for expenditure (Aguiar and Bils (2015) estimate expenditure elasticities using income as an instrument to correct for large under-reporting in the CEX).

⁸In this framework, the a^j cannot be separately identified.

Table A.1. Almost Ideal Demand System Parameter Estimation

	OLS	IV – Income
	$s_{n,t}^h$ (1)	$s_{n,t}^h$ (2)
Parameter Estimates:		
γ_{NL}	9.5×10^{-6} (.17) $\times 10^{-6}$	1.1×10^{-5} (018) $\times 10^{-5}$
β^N	-.18 (000075)	-.24 (00013)
Luxury Expenditure Elasticity	1.39	1.52
Necessity Expenditure Elasticity	.66	.55
Luxury Own Price Elasticity	-1.86	-2.45
Necessity Own Price Elasticity	-1.41	-1.8
Luxury Cross-Price Elasticity	-.86	-1.45
Necessity Cross-Price Elasticity	-.41	-.8
Observations	273,545	273,537

Notes: The unit of observation is the household-quarter. Robust standard errors in parentheses.

A3.2 Results

Here I show similar figures as those in the main text, but the the alternate 6 calibrations alongside the baseline calibration.

A4 Additional Figures and Tables




Figure A1. Model and Data: Necessity Shares by Income Group

Note: Data from 2005-06. Model income-group shares at steady state. The baseline calibration is described in the main text.




Figure A2. Model v. Data: Necessity Share




Figure A3. Model v. Data: Relative Necessity Prices

Table A.2. Relationship Unemployment and Relative Real Expenditure

	Log-Real Expenditure					
	(1)	(2)	(3)	(4)	(5)	(6)
Right hand side variables:						
UR × Exp. Ratio	0.006 (0.0063)	0.013*** (0.0042)	0.020*** (0.0066)	0.013*** (0.0041)	0.006* (0.0038)	0.017*** (0.0045)
UR × Energy				-0.001 (0.0039)		
UR × Durable					-0.033*** (0.010)	
UR × Service						0.016*** (0.0039)
Sector FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
Weighted	No	Yes	Yes	Yes	Yes	Yes
Balanced Sample	No	No	Yes	No	No	No
Observations	36,788	36,788	22,166	36,788	36,788	36,788

Notes: The unit of observation is the sector-month. Exp. ratio is the ratio of expenditure shares of poor over rich households for the sector. Standard errors, in parentheses, are clustered at the time level and are robust to auto-correlation. Significance at the 1, 5, and 10 percent levels indicated by ***, **, and *. Real Expenditure is aggregate expenditure on sector j normalized by the sector specific price index.




Figure A4. Necessity Share: Non-durables

Source: Consumer Expenditure Survey, Personal Consumption Expenditures (BEA) and Author's own calculations. Excludes housing and durable consumer goods.




Figure A5. Necessity and Luxury Expenditures Normalized by Sector Level Prices

Source: Consumer Expenditure Survey, Personal Consumption Expenditures (BEA) and Author's own calculations. Real expenditure in 2007 Q1 dollars. Necessity and luxury expenditure normalized by sector specific prices. Excludes housing.






Figure A6. Income Level CPI's during the Great Recession and Slow Recovery

Source: Bureau of Labor Statistics and Author's own calculations.

Notes: NBER recession indicated by shaded area. Weights are democratic and based on 2005-06 consumption patterns. The base period is 2007Q1.

Panel A: Unfiltered Series



Panel B: Filtered Series





Figure A7. Relative Necessity Prices including energy

Source: Bureau of Labor Statistics and Author's own calculations. Excludes housing. Data filtered following Hamilton (2018). Shaded area indicates NBER Recessions.

Panel A: Unfiltered Series



Panel B: Filtered Series





Figure A8. Relative Necessity Prices 1990-2000 Consumption Patterns

Source: Bureau of Labor Statistics and Author's own calculations. Excludes housing and energy. Data filtered following Hamilton (2018). Shaded area indicates NBER Recessions. Uses consumption shares data from 1990-2000 only when defining products as luxuries or necessities.

Panel A: Unfiltered Series



Panel B: Filtered Series





Figure A9. Relative Necessity Prices 2000-2010 Consumption Patterns

Source: Bureau of Labor Statistics and Author's own calculations. Excludes housing and energy. Data filtered following Hamilton (2018). Shaded area indicates NBER Recessions. Uses consumption shares data from 2000-2010 only when defining products as luxuries or necessities.

Panel A: Unfiltered Series



Panel B: Filtered Series




Figure A10. Relative Necessity Prices 2010-2020 Consumption Patterns
Source: Bureau of Labor Statistics and Author's own calculations. Excludes housing and energy. Data filtered following Hamilton (2018). Shaded area indicates NBER Recessions. Uses consumption shares data from 2010-2019 only when defining products as luxuries or necessities.




Figure A11. Great Recession: Change in Necessity Share by Income Quintile
 Source: Consumer Expenditure Survey and Author's own calculations. Excludes housing.

Pre-2008 Shock Data

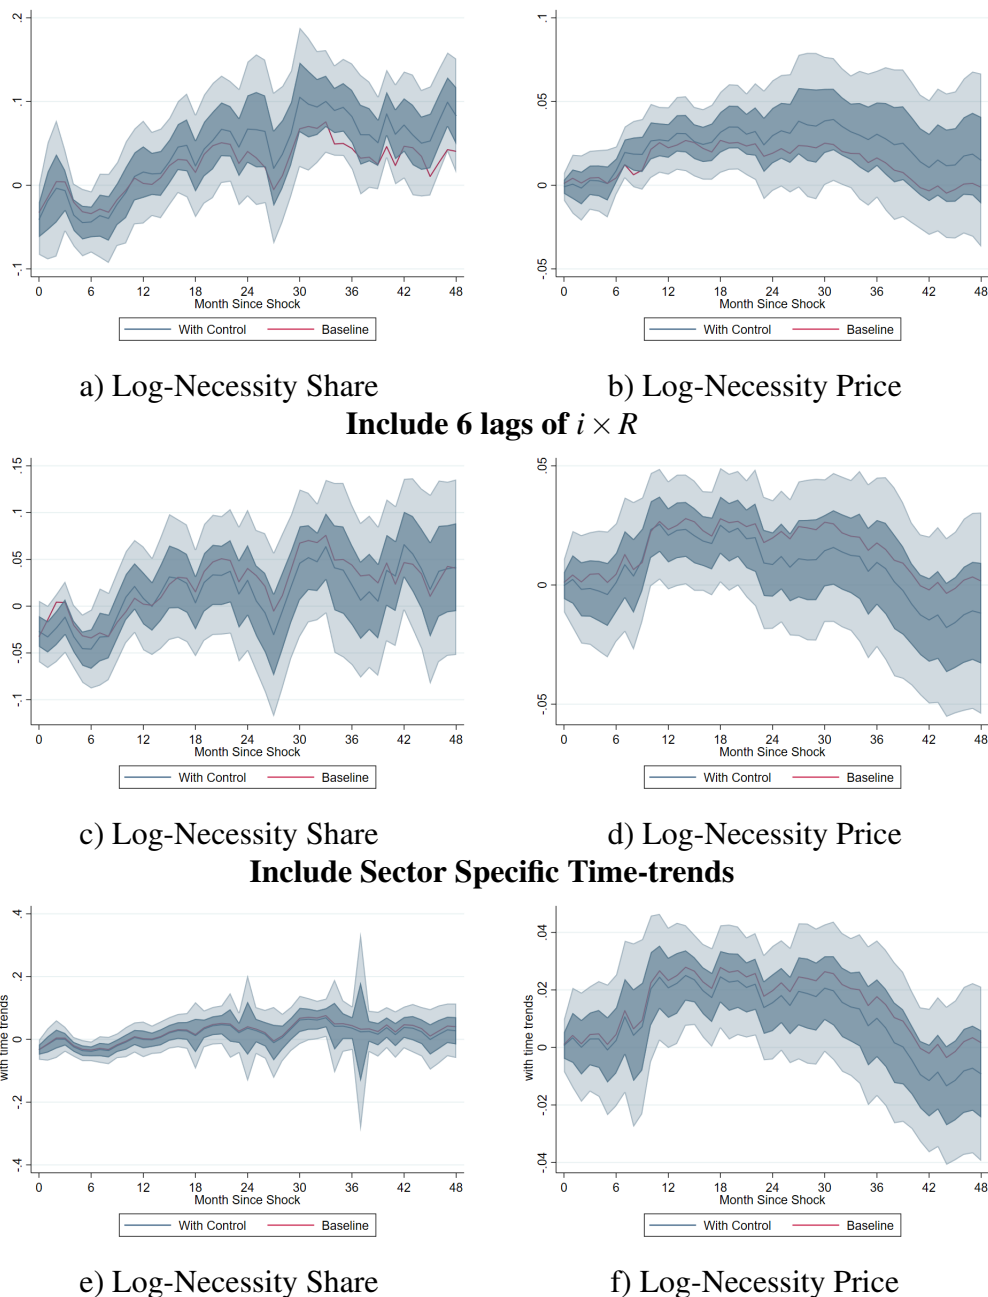


Figure A12. Additional IRF Robustness Checks

Note: Data from 1991-2019. Estimated coefficients from Local Projections explained in section 5. The unit of observation is the sector-month. Robust standard errors clustered at the monthly level are shown by one- and two-standard error confidence bands indicated by the dark and light shaded areas respectively. Sectors weighted by their share in pooled aggregate expenditure. Monetary Policy shock normalized to 25-basis point increase in 1-year treasury in month $t = 0$. When the dependent variable is log-price a balanced sample is used of 60 sectors with price data available for the entire period.

Appendix: Chapter 2

B1 Model

B1.1 Households

A continuum of identical households maximizes utility subject to their budget constraints.

The utility function for the representative household is:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma}}{1-\sigma} + \psi \frac{D_t^{1-\sigma^d}}{1-\sigma^d} - \nu \frac{H_t^{1+\phi}}{1+\phi} - \frac{\xi}{2} A_t^2 \right]$$

where C_t is nondurable consumption, D_t is the durable stock, H_t is hours worked, and A_t are household holdings of nominal bonds. We allow the curvature of nondurable consumption utility σ and the durable service flow σ^d to differ. We include bonds in the utility function to generate empirically estimated patterns of MPCs and intertemporal MPCs (?). The magnitude of MPCs in the model are determined by the parameter ξ .

The household budget constraint is

$$A_t = \frac{R_{t-1}}{\Pi_t} A_{t-1} - C_t + W_t H_t - X_t - T_t + \text{Profits}_t^k + \text{Profits}_t^s$$

where R_t is the gross nominal interest rate, Π_t is the gross inflation rate, W_t is the real wage, X_t is durable expenditure denominated in nondurable goods, T_t are transfers, Profits_t^k are profits of the capital good producing firms, and Profits_t^s are profits of the sticky-price firms.

Durables follow a standard accumulation equation

$$D_t = (1 - \delta^d)D_{t-1} + \frac{X_t}{p_t^d}$$

where δ^d is the depreciation rate of household durables and p_t^d is the relative price of durable goods.

Households pick an optimal plan $\{C_t, D_t, A_t, X_t\}_{t=0}^{\infty}$ to maximize utility. Labor supply is not chosen by the household, but instead by a union as shown below. The first order conditions for the household problem are:

$$\begin{aligned}\lambda_t &= (C_t)^{-\sigma} \\ \lambda_t &= \beta \frac{R_t}{\Pi_{t+1}} \lambda_{t+1} - \xi A_t \\ p_t^d \lambda_t &= \mu_t \\ \mu_t &= \beta(1 - \delta^d) \mu_{t+1} + \psi(D_t)^{-\sigma^d}\end{aligned}$$

B1.2 Wages

A continuum of unions indexed by j provide differentiated labor services to the final good firm that are substitutable with elasticity ε^w . Each period there is a iid probability θ^w that the union cannot adjust the contract wage. In this case, wages will adjust by a fraction χ^w of last periods inflation.

The demand for hours from union j at time $t + s$ conditional on having last reset wages at time t is

$$H_{t+s}^d(j) = H_{t+s}^d \left(\frac{W_t(j) \left(\frac{P_{t+s-1}}{P_{t-1}} \right) \chi^w \left(\frac{P_t}{P_{t+s}} \right)}{W_{t+s}} \right)^{-\varepsilon^w} = H_{t+s}^d W_{t+s}^{\varepsilon^w} \left(\frac{P_{t+s}}{P_t} \right)^{\varepsilon^w} \left(\frac{P_{t+s-1}}{P_{t-1}} \right)^{-\varepsilon^w \chi^w} W_t(j)^{-\varepsilon^w}$$

where P_t is the price level at time t .

If the union can adjust its wage at time t it picks the optimal wage to maximize the

expected discounted utility of the representative household while this wage prevails:

$$\begin{aligned} & \max_{w_t^*} \sum_{s=0}^{\infty} (\beta \theta^w)^s H_{t+s}^d W_{t+s}^{\varepsilon^w} \left(\frac{P_{t+s}}{P_t} \right)^{\varepsilon^w} \left(\frac{P_{t+s-1}}{P_{t-1}} \right)^{-\varepsilon^w \chi^w} \\ & \times \left[\tilde{\lambda}_{t+s} \left(\frac{P_{t+s-1}}{P_{t-1}} \right)^{\chi^w} \left(\frac{P_{t+s}}{P_t} \right)^{-1} (W_t^*)^{1-\varepsilon^w} - v H_{t+s}^\phi (W_t^*)^{-\varepsilon^w} \right] \end{aligned}$$

where $\tilde{\lambda} = (1 - \gamma)\lambda_t^o + \gamma\lambda_t^r$

The first order condition for the union is:

$$\begin{aligned} & (\varepsilon^w - 1) \sum_{s=0}^{\infty} (\beta \theta^w)^s H_{t+s}^d W_{t+s}^{\varepsilon^w} \left(\frac{P_{t+s}}{P_t} \right)^{\varepsilon^w - 1} \left(\frac{P_{t+s-1}}{P_{t-1}} \right)^{-\chi^w (\varepsilon^w - 1)} \tilde{\lambda}_{t+s} (W_t^*)^{1-\varepsilon^w} \\ & = \varepsilon^w v \sum_{s=0}^{\infty} (\beta \theta^w)^s H_{t+s}^d H_{t+s}^\phi W_{t+s}^{\varepsilon^w} \left(\frac{P_{t+s}}{P_t} \right)^{\varepsilon^w} \left(\frac{P_{t+s-1}}{P_{t-1}} \right)^{-\varepsilon^w \chi^w} (W_t^*)^{-\varepsilon^w} \end{aligned}$$

We write it recursively using

$$\begin{aligned} F_{1t} &= v H_t^d H_t^\phi W_t^{\varepsilon^w} (W_t^*)^{-\varepsilon^w} + \beta \theta^w \Pi_{t+1}^{\varepsilon^w} \Pi_t^{-\chi^w \varepsilon^w} \left(\frac{W_t^*}{W_{t+1}^*} \right)^{-\varepsilon^w} F_{1,t+1} \\ F_{2t} &= H_t^d W_t^{\varepsilon^w} \tilde{\lambda}_t (W_t^*)^{1-\varepsilon^w} + \beta \theta^w \Pi_{t+1}^{\varepsilon^w - 1} \Pi_t^{-\chi^w (\varepsilon^w - 1)} \left(\frac{W_t^*}{W_{t+1}^*} \right)^{1-\varepsilon^w} F_{2,t+1} \\ \varepsilon^w F_{1t} &= (\varepsilon^w - 1) F_{2t} \end{aligned}$$

Wage dispersion across unions lead to inefficiency in the labor types used by firms. This creates a wedge between hours worked H_t and effective hours worked H_t^d , which we denote by s_t^w ,

$$H_t = s_t^w H_t^d,$$

and which evolves according to,

$$s_t^w = (1 - \theta^w) \left(\frac{W_t^*}{W_t} \right)^{-\varepsilon^w} + \theta \left(\frac{W_{t-1}}{W_t} \right)^{-\varepsilon^w} \Pi_t^{\varepsilon^w} s_{t-1}^w$$

B1.3 Production of capital goods

The representative capital goods firm chooses investment I_t , the capital stock K_t , and the utilization rate u_t to maximize profits,

$$\begin{aligned} \max_{\{K_{t+s}, I_{t+s}, u_{t+s}\}} & \sum_{s=0}^{\infty} \beta^s \lambda_{t+s}^o \text{Profits}_t^k \\ \text{s.t. Profits}_t^k &= R_{t+s}^k u_{t+s} K_{t+s-1} - I_t \\ K_t &= (1 - \delta(u_t)) K_{t-1} + I_t \left[1 - S \left(\frac{I_t}{I_{t-1}} \right) \right] \end{aligned}$$

where R_{t+s}^k is the rental rate of capital paid by the final goods firm, $S \left(\frac{I_t}{I_{t-1}} \right)$ is an investment adjustment cost, and $\delta(u)$ is the depreciation rate of capital which is increasing in utilization.

Let ζ_t be the Lagrange multiplier on the capital accumulation equation and define Tobin's q as the relative value of capital to nondurable consumption,

$$q_t = \frac{\zeta_t}{\lambda_t^o}.$$

Then the first order conditions for the representative capital producing firms are,

$$\begin{aligned} 1 &= q_t \left[1 - S \left(\frac{I_t}{I_{t-1}} \right) - \left(\frac{I_t}{I_{t-1}} \right) S' \left(\frac{I_t}{I_{t-1}} \right) \right] + \beta \frac{\lambda_{t+1}^o}{\lambda_t^o} q_{t+1} \left(\frac{I_{t+1}}{I_t} \right)^2 S' \left(\frac{I_{t+1}}{I_t} \right) \\ q_t &= \beta \frac{\lambda_{t+1}^o}{\lambda_t^o} R_{t+1}^k u_{t+1} + \beta (1 - \delta(u_{t+1})) \frac{\lambda_{t+1}^o}{\lambda_t^o} q_{t+1} \\ R_t^k &= \delta'(u_t) q_t \end{aligned}$$

B1.4 Production of final goods

Final output Y_t is produced using a Cobb-Douglas production function with capital share α ,

$$s_t Y_t = Z_t (u_t K_{t-1})^\alpha (H_t^d)^{1-\alpha}$$

where Z_t is aggregate TFP. The wedge s_t captures a distortion from price dispersion, which is described below.

The cost minimization for the representative final goods firm is

$$\begin{aligned} \min R_t^k u_t K_{t-1} + W_t H_t^d \\ \text{s.t. } Z_t (u_t K_{t-1})^\alpha (H_t^d)^{1-\alpha} = s_t Y_t \end{aligned}$$

which yields the following first order conditions for capital and labor,

$$\begin{aligned} R_t^k &= \xi_t \alpha \frac{s_t Y_t}{u_t K_{t-1}} \\ W_t &= \xi_t (1 - \alpha) \frac{s_t Y_t}{H_t^d} \end{aligned}$$

where ξ_t is the Lagrange multiplier on the production function. Dividing the two first order conditions yields the optimal capital-labor ratio,

$$\frac{u_t K_{t-1}}{H_t^d} = \frac{\alpha}{1 - \alpha} \frac{W_t}{R_t^k},$$

which in turn yields the marginal cost of output is,

$$MC_t = \alpha^{-\alpha} (1 - \alpha)^{-(1-\alpha)} (R_t^k)^\alpha W_t^{1-\alpha} \frac{1}{Z_t}$$

With perfect competition among final goods firms, the real final goods price is equal to marginal cost,

$$p_t^f = MC_t,$$

and final good firms make zero profits.

B1.5 Prices

A continuum of retailers purchases final goods at price p_t^f and differentiates these goods with elasticity of substitution ε . Retailers can only reset their price with probability θ . The profit maximization problem for setting the reset price is

$$\max_{p_t^*} \sum_{s=0}^{\infty} \beta^s \left(\frac{\lambda_{t+s}^o}{\lambda_t^o} \right) \theta^s Y_{t+s} \left[(p_t^*)^{1-\varepsilon} \left(\frac{P_{t+s}}{P_t} \right)^{\varepsilon-1} - (p_t^*)^{-\varepsilon} \left(\frac{P_{t+s}}{P_t} \right)^{\varepsilon} p_{t+s}^f \right]$$

The first order condition for the optimal reset price is

$$\begin{aligned} \varepsilon \sum_{s=0}^{\infty} \beta^s \left(\frac{\lambda_{t+s}^o}{\lambda_t^o} \right) \theta^s Y_{t+s} \left(\frac{P_{t+s}}{P_t} \right)^{\varepsilon} (p_t^*)^{-\varepsilon-1} p_{t+s}^f \\ = (\varepsilon - 1) \sum_{s=0}^{\infty} \beta^s \left(\frac{\lambda_{t+s}^o}{\lambda_t^o} \right) \theta^s Y_{t+s} \left(\frac{P_{t+s}}{P_t} \right)^{\varepsilon-1} (p_t^*)^{-\varepsilon} \end{aligned}$$

which we write recursively as

$$\begin{aligned} X_{1t} &= Y_t p_t^f (p_t^*)^{-\varepsilon-1} + \beta \theta \left(\frac{\lambda_{t+1}^o}{\lambda_t^o} \right) \left(\frac{P_{t+1}}{P_t} \right)^{\varepsilon} \left(\frac{p_t^*}{p_{t+1}^*} \right)^{-\varepsilon-1} X_{1,t+1} \\ X_{2t} &= Y_t (p_t^*)^{-\varepsilon} + \beta \theta \left(\frac{\lambda_{t+1}^o}{\lambda_t^o} \right) \left(\frac{P_{t+1}}{P_t} \right)^{\varepsilon-1} \left(\frac{p_t^*}{p_{t+1}^*} \right)^{-\varepsilon} X_{2,t+1} \\ \varepsilon X_{1t} &= (\varepsilon - 1) X_{2t} \end{aligned}$$

The optimal reset price determines aggregate inflation

$$1 = (1 - \theta)(p_t^*)^{1-\varepsilon} + \theta\Pi_t^{-(1-\varepsilon)}$$

as well as the relative price distortion

$$\begin{aligned} s_t &= \int_0^1 \left(\frac{P_t(i)}{P_t} \right)^{-\varepsilon} di \\ &= (1 - \theta)(p_t^*)^{-\varepsilon} + \theta \int_0^1 \left(\frac{P_{t-1}(i)}{P_t} \right)^{-\varepsilon} di \\ &= (1 - \theta)(p_t^*)^{-\varepsilon} + \theta\Pi_t^\varepsilon s_{t-1} \end{aligned}$$

Due to monopoly power, the sticky-price firms make non-zero profits in equilibrium equal to

$$\text{Profits}_t^s = Y_t(1 - p_t^f)$$

B1.6 Government

The central bank sets the gross nominal interest rate according to the following interest rate rule,

$$R_t = (1 - \rho_r)R_{t-1} + \rho_r \left[R + \phi_\pi(\Pi_t - \bar{\Pi}) + \phi_y \left(\frac{Y_t}{\bar{Y}} - 1 \right) \right]$$

where ρ_r determines the degree of interest rate smoothing, ϕ_π the response to deviations of inflation from target, and ϕ_y the response to deviations of output from target.

The government issues one-period nominal bonds at gross interest R_t to cover debt

repayment and any fiscal deficit.

$$B_t = \frac{R_{t-1}}{\Pi_t} B_{t-1} - T_t$$

To balance the budget over time, taxes are an increasing function of the debt level,

$$T_t = \bar{T} + \phi_b (B_{t-k} - \bar{B}) - \varepsilon_t.$$

We allow for a lag of k periods in the response of taxes to debt. The shock ε_t represents a one-time deficit financed transfer from the government to households.

B1.7 Durable Goods Production

Durable goods are produced competitively using nondurables N_t as inputs,

$$\frac{X_{it}}{p_t^d} = N_{it} \left(\frac{X_t}{\bar{X}} \frac{1}{p_t^d} \right)^{-\zeta}$$

where $\frac{X_{it}}{p_t^d}$ is the real production of durable goods by firm i and ζ is a negative production externality.

Real profits from the sale of durable goods are

$$\max_{N_{it}} X_{it} - N_{it} = \max_{N_{it}} p_t^d N_{it} \left(\frac{X_t}{\bar{X}} \frac{1}{p_t^d} \right)^{-\zeta} - N_{it}$$

Profit maximization yields an upward sloping supply curve,

$$p_t^d = \left(\frac{X_t}{\bar{X}} \right)^{\frac{\zeta}{1+\zeta}}$$

where \bar{X} is steady state durable expenditure, so the steady relative durable price is normalized to 1. Since durable expenditure is denominated in units of nondurable consumption, the supply

elasticity of real durable goods is given by $\frac{1}{\zeta}$.

B1.8 Market Clearing

The goods market clears if total expenditure equals output.

$$Y_t = C_t + I_t + X_t$$

The bond market clears if bonds supplied by the government equal bonds held by households,

$$B_t = A_t \tag{B1.1}$$

B1.9 Functional Forms

We assume the following functional forms:

$$\begin{aligned} \delta(u_t) &= \delta_0 + \delta_1(u_t - 1) + \delta_2(u_t - 1)^2 \\ S\left(\frac{I_t}{I_{t-1}}\right) &= \frac{\kappa}{2} \left(\frac{I_t}{I_{t-1}} - 1\right)^2 \end{aligned}$$

B2 Nominal Counterfactuals

Figures B1 and B2 display the counterfactuals for nominal PCE and nominal motor vehicle expenditure.

B3 Data Appendix

B3.1 Details for Figure 2.1

The following are details of the Sahm et al. (2012) calculation and our update. Sahm et al. (2012) use Parker et al. (2013) estimate of a marginal propensity to spend on new motor vehicles

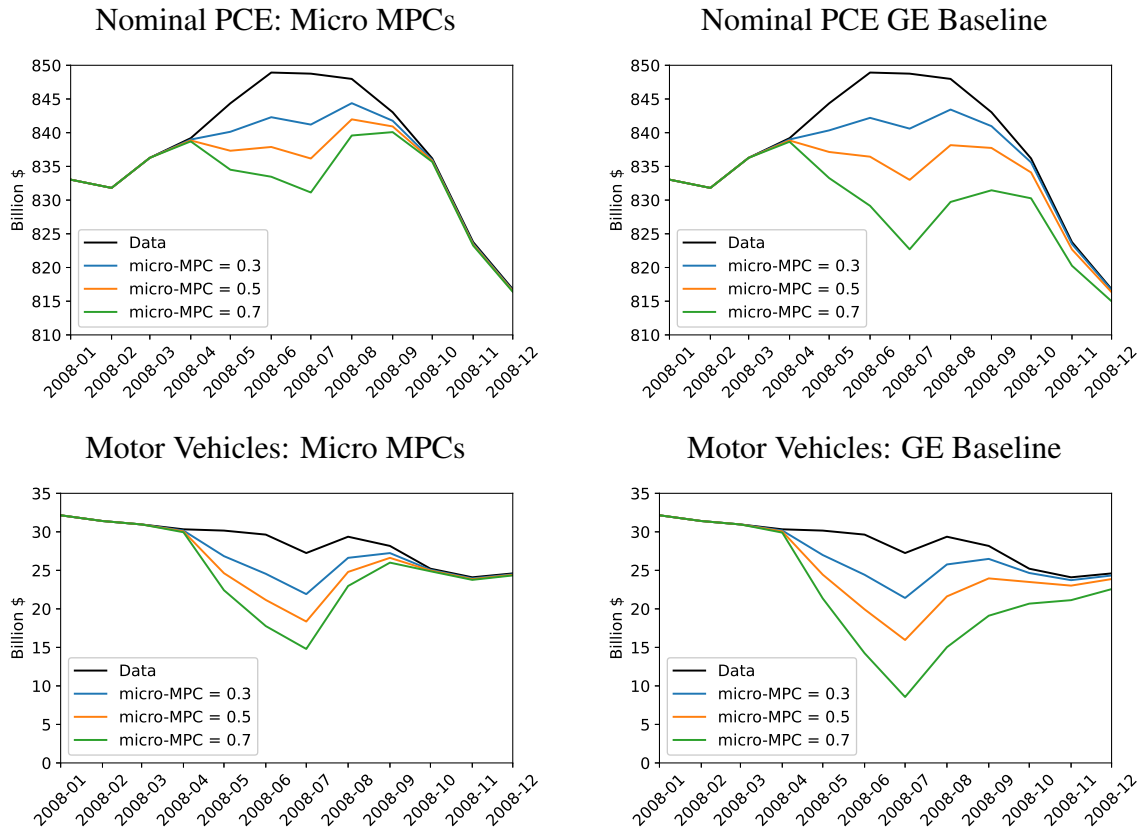


Figure B1. Counterfactual Nominal Consumption Expenditures: Baseline Model

Notes. Based on NK model simulations and actual data on rebates and consumption. The micro MPC value refers to the MPC for total consumption.

of 0.357 (from Table 7 of Parker et al. (2013)) to calculate induced spending. Following Parker et al. (2013), they assume that the spending is evenly distributed between the current and the next month. They use seasonal factors to seasonally adjust the induced spending. We follow the same procedure to calculate induced spending and then subtract it from actual spending to create the implied counterfactual, which does not account for partial or general equilibrium effects.

The following graph shows counterfactuals from the motor vehicle accounting exercise for different assumptions of how much the spending is smoothed.

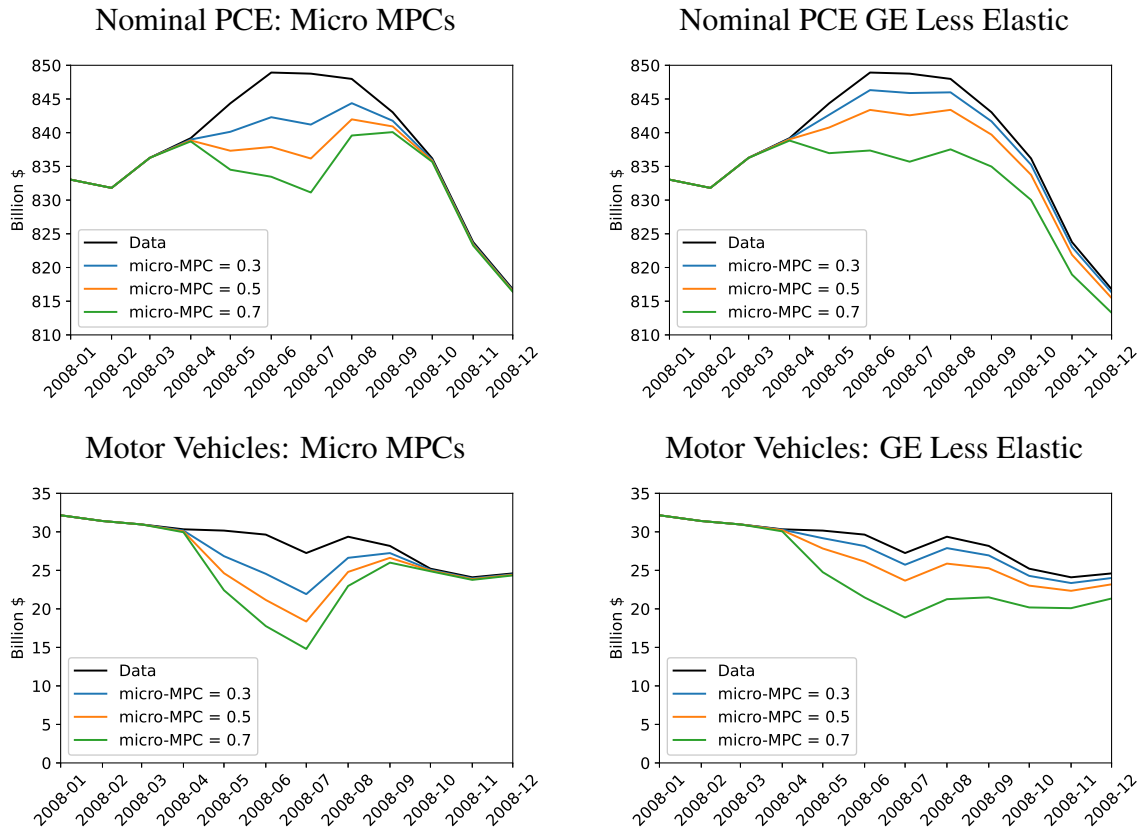


Figure B2. Counterfactual Nominal Consumption Expenditures: Less Elastic Durable Supply Model

Notes. Based on NK model simulations and actual data on rebates and consumption. The micro MPC value refers to the MPC for total consumption.

B4 Decomposing the Difference Between OLS and DID Imputation

In section 2.4, our implementation of Borusyak et al. (2022) DID imputation method yields a much smaller MPC for total expenditure (.3) compared to our OLS replication of Parker et al. (2013) (0.52). We use Sun and Abraham (2020) method to decompose OLS event studies and show that the difference between the OLS and DID imputation coefficients can be explained by OLS applying negative weights to past-treated units.

We first apply Sun and Abraham (2020) to decompose the differences and differences




Figure B3. Expenditures on New Motor Vehicles: Alternative Counterfactuals

Note. The baseline counterfactual assumes that rebate-induced spending is spread over two months. The two alternatives show the counterfactual with the induced spending spread over three or four months.

coefficient (β_2 from 2.4.1) as a linear combination of cohort average treatment effects on the treated (CATT) from the period households receive the rebate and from other periods. Where the CATT from each period ($\gamma_{e,h}$) are estimated in the following saturated regression:

$$\begin{aligned}
 C_{i,t+1} - C_{i,t} = & \sum_s \beta_{0s} \text{month}_{s,i} + \beta_1' \mathbf{X}_{i,t} \\
 & + \sum_e \gamma_{e,0} (I(ESP_{i,t+1}) \times I(t+1=e)) \\
 & + \sum_{h \neq 0} \sum_e \gamma_{e,h} (I(ESP_{i,t+1+h}) \times I(t+1+h=e)) \\
 & + \varepsilon_{i,t+1}.
 \end{aligned} \tag{B4.1}$$

In the above expression, $\gamma_{e,0}$, represents the contemporaneous treatment effect for house-

holds that report receiving their rebate in interview e .⁹ Each $\gamma_{e,h}$ represent separate CATT for different horizons around the treatment date. For example, if $h = 1$ then $\gamma_{e,h}$ would be the estimated impact of treatment on the period after receiving the rebate. We could also estimate separate effects for the never-treated units in each interview, however, in practice these would be collinear with the interview-month fixed effects, so we treat the never-treated households as the excluded category. Hence, $\gamma_{e,\infty} = 0 \forall e$.

We can express the original OLS coefficient, β_2 , as a linear combination of these $\gamma_{e,h}$ (Sun and Abraham 2020):

$$\beta_2 = \sum_h \sum_e \omega_{e,h} \gamma_{e,h}$$

Where the weights, $\omega_{e,h}$, are the coefficients in the following series of regressions:

$$\begin{aligned} (I(ESP_{i,t+1+h}) \times I(t+1+h=e)) &= \sum_s \tilde{\beta}_{0s} month_{s,i} + \tilde{\beta}'_1 \mathbf{X}_{i,t} \\ &+ \omega_{e,h} (I(ESP_{i,t+1}) \times I(t+1=e)) + \varepsilon_{i,t+1}. \end{aligned}$$

The weights on the period the rebate is received sum to 1, $\sum_e \omega_{e,0} = 1$, while the weights on the other sum to -1, $\sum_e \sum_{h \neq 0} \omega_{e,h} = -1$. In each period, the treatment weights and the other period weights are symmetric i.e. $\omega_{e,0} = -\sum_{h \neq 0} \omega_{e,h}$.¹⁰

In the right panel of figure B4 we plot the estimated weights ($\omega_{e,h}$), separately for each period. Where:

⁹In keeping with the notation in Sun and Abraham (2020), e could also represent the household's rebate cohort. This results in a similar decomposition, but figure 2.13 would then represent treatment cohorts rather than interview dates. We find that the decomposition via interview date is more intuitive for our application.

¹⁰The never treated units are included in the weights for the other periods.

$$\text{Weight Treated} = \omega_{e,0}$$

$$\text{Weight Not-yet Treated} := \omega_{e,h<0} = \omega_{e,\infty} + \sum_{h<0} \omega_{e,h}$$

$$\text{Weight Past Treated} := \omega_{e,h>0} = \sum_{h>0} \omega_{e,h}$$

The treated weight each period is symmetric with the non-treated and past-treated weights: $\omega_{e,0} = -(\omega_{e,h<0} + \omega_{e,h>0})$. Since these weights are symmetric, in figure 2.13 in the main text, we show only the per-period treatment weights in the upper-left panel.

With our estimated weights ($\omega_{e,h}$) and CATT ($\gamma_{e,h}$) we can decompose the relative contribution of each period and horizon of treatment to the final OLS DID coefficient (β_2). We can also estimate average coefficients for past-treated, not-yet treated, and treated units in each period:

$$\text{Coefficient Treated} = \gamma_{e,0}$$

$$\text{Coefficient Not-yet Treated} := \gamma_{e,h<0} = \frac{\sum_{h<0} \omega_{e,h} \gamma_{e,h}}{\sum_{h<0} \omega_{e,h}}$$

$$\text{Coefficient Treated} - \text{Not Yet Treated} = \gamma_{e,0} - \frac{\omega_{e,h<0} \gamma_{e,h<0}}{\omega_{e,h<0} + \omega_{e,h>0}}$$

$$\text{Coefficient Past Treated} := \gamma_{e,h>0} = \frac{\sum_{h>0} \omega_{e,h} \gamma_{e,h}}{\sum_{h>0} \omega_{e,h}}$$

$$\text{Average Coefficient} := \gamma_e = \gamma_{e,0} - \frac{\omega_{e,h<0} \gamma_{e,h<0} + \omega_{e,h>0} \gamma_{e,h>0}}{\omega_{e,h<0} + \omega_{e,h>0}}$$

The right panel of figure B4 shows the estimated coefficients at each horizon as described above, while the upper-right panel in the main text (figure 2.13 shows the average period-coefficients (γ_e)). The main text also shows the relative contribution to the average coefficient coming from the difference between the treated and the not-yet treated and the past treated in

the bottom right panel of figure 2.13. The relative contributions from each period and horizon (bottom panel in figure 2.13 in the main text) are simply the period weights multiplied by the period coefficients.

The reason why the past-treated units in September had such a large contribution to the overall OLS coefficient (see figure 2.13) is because the past treated units receive a sizable negative weight and because these past treated units have such a large negative average coefficient ($\gamma_{e,h} > 0$). Part of this negative coefficient could be explained by households that report receiving their rebate in the June interview reverting back to regular spending over the next interview. What is really interesting is that even the treated units have a negative estimated coefficient in September, while its the non-treated units that have a positive coefficient (although this is very small). For these two reasons, we do not believe that the past-treated units serve as valid controls for the treated units.

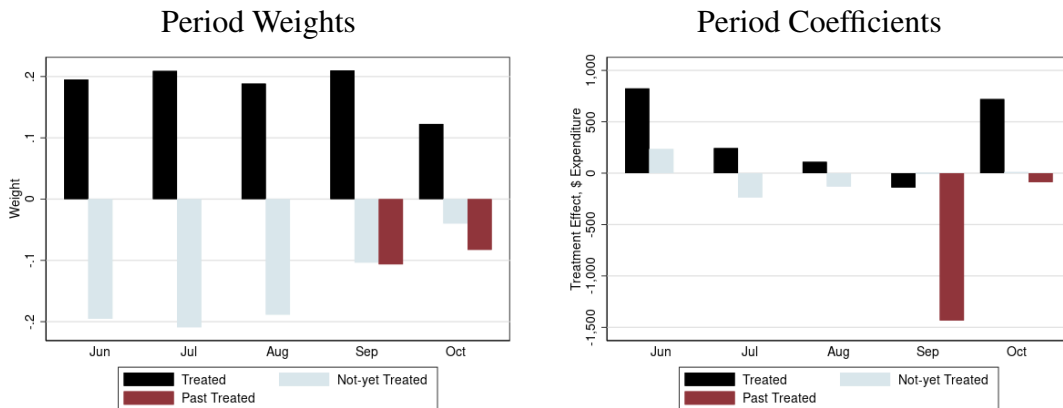


Figure B4. OLS: Weights and CATT

Notes. The dependent variable is the change in total expenditure. Based on estimations of equation 2.4.1 via OLS.

B5 Additional Regression Table

B6 DID Decomposition for Non-durable and New Vehicle Expenditure

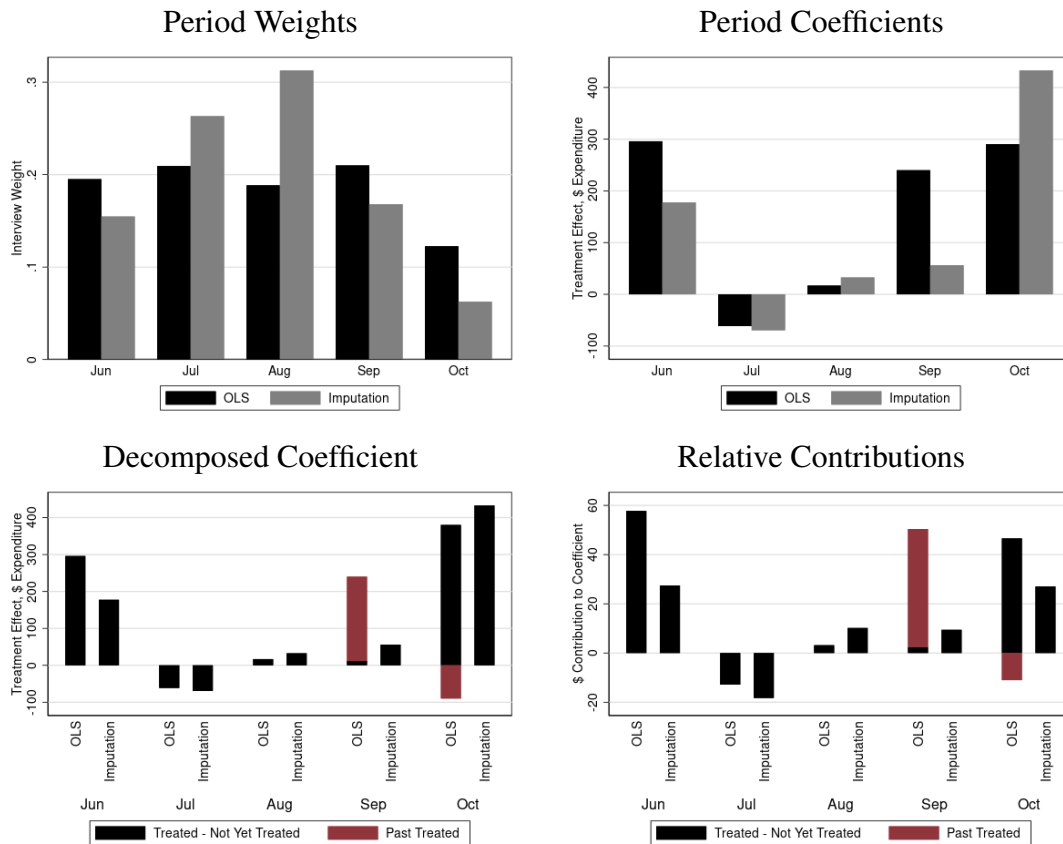


Figure B5. Non-Durable: Decomposing the OLS and DID Imputation Coefficients

Notes. The dependent variable is the change in non-durable expenditure. Based on estimations of equation 2.4.1 via OLS and the DID imputation method described in section 2.4. Periods after October, 2008, also receive positive weight, however, these weights are quite small and are not shown here.

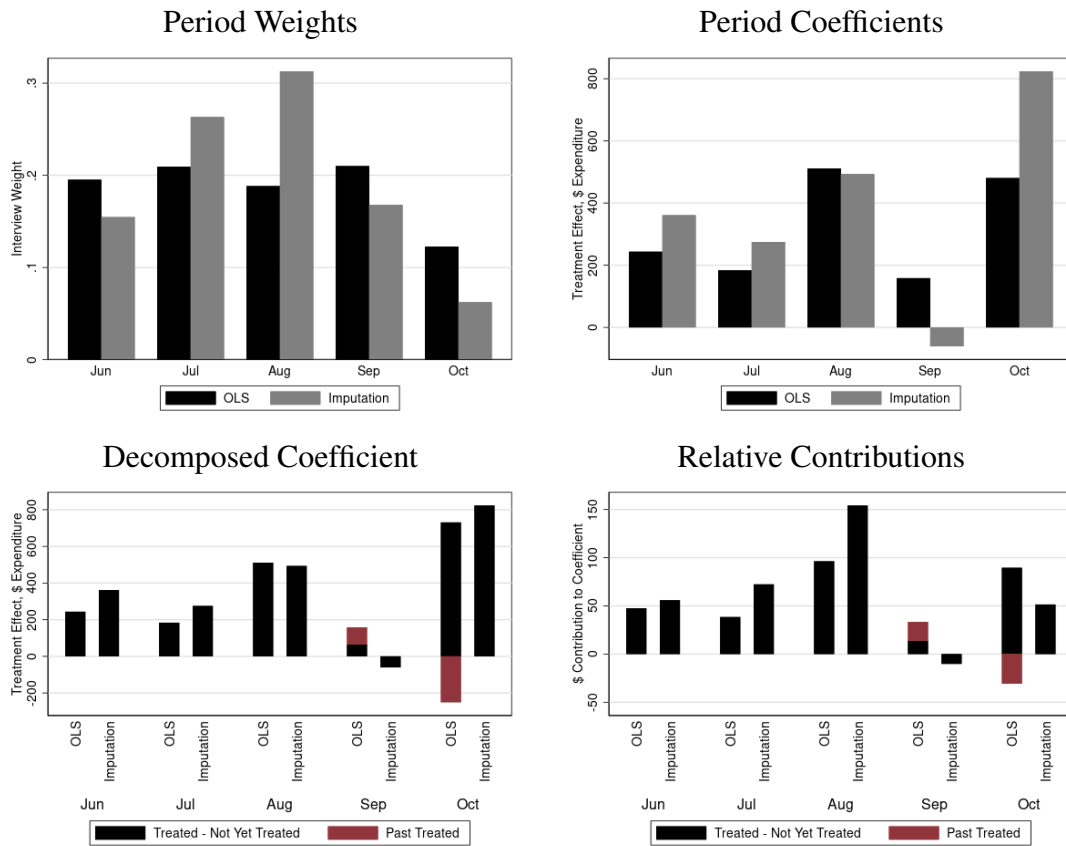


Figure B6. New Vehicles: Decomposing the OLS and DID Imputation Coefficients

Notes. The dependent variable is the change in new vehicle expenditure. Based on estimations of equation 2.4.1 via OLS and the DID imputation method described in section 2.4. Periods after October, 2008, also receive positive weight, however, these weights are quite small and are not shown here.

Table B1. Expenditure Response to Rebate: Full Sample (monthly frequency)

	FOOD	SND	ND	EXP	CAR	ln(FOOD)	ln(SND)	ln(ND)	ln(EXP)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Right hand side variables:									
Lead 1 Rebate Amount	0.00040 (00052)	0.013 (0011)	0.032** (0014)	0.14** (0062)	0.11** (0045)				
Rebate Amount	0.013* (00067)	0.035** (0014)	0.054*** (0018)	0.13* (0067)	0.081* (0048)				
Lag 1 Rebate Amount	0.0095 (00074)	0.035** (0015)	0.038** (0019)	0.079 (0065)	0.039 (0039)				
Lag 2 Rebate Amount	-0.00065 (00066)	0.020 (0014)	0.025 (0018)	0.16* (0084)	0.055 (0058)				
Lead 1 Rebate Indicator						0.0036 (00068)	0.0067 (00056)	0.015** (00060)	0.023*** (00079)
Rebate Indicator						0.013* (00080)	0.017** (00066)	0.025*** (00068)	0.039*** (00084)
Lag 1 Rebate Indicator						0.0078 (00084)	0.018*** (00067)	0.019*** (00072)	0.034*** (00092)
Lag 2 Rebate Indicator						-0.0023 (00075)	0.0077 (00063)	0.0074 (00066)	0.033*** (00096)
Time Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Household Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R ²	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Observations	66,394	66,394	66,394	66,394	66,394	66,391	66,394	66,386	66,393

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

Appendix: Chapter 3

C1 Household Model

Consider a household with quadratic utility over consumption good c_t , stochastic income stream $\{w_t\}_{t=0}^{\infty}$, household specific price level p_t and the ability to invest in a risk free bond b_t at interest rate r . My goal is to see how the household responds to unanticipated and anticipated increases to their own price index.

The household's problem becomes:

$$\max \mathbb{E}_0 \beta u(c_t) \tag{C1.1}$$

$$s.t. p_t c_t + b_{t+1} \left(\frac{1}{1+r} \right) = w_t + b_t \tag{C1.2}$$

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta_t b_t^2 \right] < \infty \tag{C1.3}$$

where $\beta < 1$ is the household's discount factor. I assume that the household has per-period quadratic utility that takes the form:

$$u(c_t) = c_t - \gamma c_t^2. \tag{C1.4}$$

I further assume that the household consumes in the region in which marginal utility is positive.

The Household Euler equation is:

$$\frac{1 - 2\gamma c_t}{p_t} = \beta \mathbb{E}_t \frac{1 - 2\gamma c_{t+1}}{p_{t+1}} (1 + r). \quad (\text{C1.5})$$

I make a further simplifying assumption that the interest rate $r = 1/\beta - 1$. I consider three types of changes to the household's price index when the household has perfect foresight: (1) an unanticipated permanent increase in p at time T from p^A to p^B , (2) an anticipated permanent shock in p from p^A to p^B at time T , and (3) an unanticipated temporary shock to p at time T . Finally, using a slightly modified model, I show the household's response to price index shocks when they do not have perfect foresight.

C1.1 Permanent Unanticipated Shock

In case (1), the household now assumes that $p_t = p^B \forall t > T$, so the Euler equation equation (C1.5) becomes:

$$\begin{aligned} \frac{1 - 2\gamma c_t}{p^A} &= \mathbb{E}_t \frac{1 - 2\gamma c_{t+1}}{p^A} \text{ if } t \leq T - 1 \\ \frac{1 - 2\gamma c_t}{p^B} &= \mathbb{E}_t \frac{1 - 2\gamma c_{t+1}}{p^B} \text{ if } t > T - 1. \end{aligned} \quad (\text{C1.6})$$

In either case, equation (C1.5) becomes $c_t = \mathbb{E}[c_{t+1}]$.

From the transversality condition equation (C1.3), I get that

$$\lim_{t \rightarrow \infty} \beta^{t/2} b_{t+1} = 0 \quad (\text{C1.7})$$

I combine equation (C1.7) with the budget constraint equation (C1.2) to get that:

$$-\frac{b_t}{p^t} = \sum_{j=0}^{\infty} \frac{\beta^j}{p^B} \mathbb{E}_t \approx (w_{t+j} - c_{t+j}) \quad \forall t > T - 1 \quad (\text{C1.8})$$

Next I use the martingale property of the euler equation, and the fact that $\beta < 1$ to solve for c_t as a function of future income streams and current bond holdings:

$$\frac{c_t}{1 - \beta} = \sum_{j=0}^{\infty} \frac{\beta^j}{p^B} \mathbb{E}_T (w_{t+j} + b_t) \quad \forall t > T - 1 \quad (\text{C1.9})$$

$$c_t = \frac{r}{1 + r} \sum_{j=0}^{\infty} \frac{\beta^j}{p^B} \mathbb{E}_T (w_{t+j} + b_t) \quad \forall t > T - 1. \quad (\text{C1.10})$$

It is straightforward to show that as the household's price index increases, current consumption decreases. This is because an increase in the household's price index is equivalent to a decrease in the household's lifetime income. A positive price index shock is a negative wealth shock and consumption is a normal good.

I can also show what happens to current real consumption regardless of the path of future price index shocks assuming that the household has perfect foresight. If the price index changes again at time $T + 1$ to p^C then consumption at time $T+1$ should be:

$$c_T = \frac{r}{1 + r} \left(\frac{\beta^j}{p^B} \mathbb{E}_T (w_t + b_t) + \sum_{j=1}^{\infty} \frac{\beta^j}{p^C} \mathbb{E}_T (w_{t+j} + b_t) \right). \quad (\text{C1.11})$$

Continuing this process, but letting each time $t + j$ have it's own price index p_{t+j} , then under perfect foresight the household's consumption at time t becomes:

$$c_t = \frac{r}{1 + r} \left(\sum_{j=0}^{\infty} \frac{\beta^j}{p_{t+j}} \mathbb{E}_t (w_{t+j} + b_t) \right). \quad (\text{C1.12})$$

Nominal consumption does not change with respect to changes in p as

$$p^B c_t = \frac{r}{1+r} \sum_{j=0}^{\infty} \beta^j (w_{t+j} + b_t) \quad (\text{C1.13})$$

does not depend on p^B .

C1.2 Permanent Anticipated Shock

Under the second case (2), I show what happens to consumption and nominal consumption in period $T-1$ the period prior to the price index shock, if the household expects the increase in the price index.

Then the Euler equation becomes:

$$\frac{1 - 2\gamma c_{T-1}}{p^A} = \mathbb{E}_{T-1} \left[\frac{1 - 2\gamma c_T}{p^B} \right] \quad (\text{C1.14})$$

$$1 - 2\gamma c_{T-1} = \mathbb{E}_{T-1} \left[\frac{p^A}{p^B} \left(1 - 2\gamma \frac{r}{1+r} \sum_{j=0}^{\infty} \frac{\beta^j}{p^B} (w_{T+j} + b_T) \right) \right] \quad (\text{C1.15})$$

$$c_{T-1} = \frac{1}{2\gamma} - \frac{p^A}{p^B} \frac{1}{2\gamma} + \frac{p^A}{p^B} \left(\frac{r}{1+r} \sum_{j=0}^{\infty} \mathbb{E}_{T-1} \frac{\beta^j}{p^B} (w_{T+j} + b_T) \right). \quad (\text{C1.16})$$

Both nominal and real consumption decrease at time $T - 1$ in anticipation of the increased price index if and only if:

$$\frac{1}{2\gamma} < \frac{2}{p^B} \left(\frac{r}{1+r} \sum_{j=0}^{\infty} \mathbb{E}_{T-1} \beta^j (w_{T+j} + b_T) \right). \quad (\text{C1.17})$$

C1.3 Transitory Shock

This subsection shows how a household would respond to a transitory shock to their price level. I start by showing analytically the household's response to a one period transitory shock (at T from p^A to p^B) that they expect to subside the next period. Next, I use numerical methods (dynare) to show the household's response if shocks to their price index follow an AR(1) process.

First, if there is a one-period shock to the household's price level from p^A to p^B at time T, then the euler equation at time T becomes:

$$\frac{1 - 2\gamma c_T}{p^B} = \mathbb{E}_T \left[\frac{1 - 2\gamma c_{T+1}}{p^A} \right] \quad (\text{C1.18})$$

Similar to the previous subsection, consumption at time T is then:

$$c_T = \frac{1}{2\gamma} - \frac{p^B}{p^A} \frac{1}{2\gamma} + \frac{p^B}{p^A} \left(\frac{r}{1+r} \sum_{j=1}^{\infty} \mathbb{E}_T \frac{\beta^j}{p^A} (w_{T+j} + b_T) \right). \quad (\text{C1.19})$$

As before, we have different responses of nominal consumption depending on the values of γ and w ; however, if

$$\frac{1}{\gamma} < \frac{p^A}{p^B 2\gamma} + \frac{2}{p^A} \left(\frac{r}{1+r} \sum_{j=1}^{\infty} \mathbb{E}_T \beta^j (w_{T+j} + b_T) \right) \quad (\text{C1.20})$$

then nominal consumption increases in period T. If the household is sufficiently risk averse then they respond to the increase in their price level by increasing their nominal consumption to then smooth their real consumption. Finally, real consumption also increases after an increase in price index if and only if:

$$\frac{1}{2\gamma p^A} < \frac{1}{(p^A)^2} \left(\frac{r}{1+r} \sum_{j=1}^{\infty} \mathbb{E}_T \beta^j (w_{T+j} + b_{T+j}) \right) \quad (\text{C1.21})$$

C1.4 Stochastic Shock Process

In the previous subsections, the household “knew” about the future path of their price index. Consider instead the case where the price index follows a stochastic path:

$$p_t = e^{\pi_t} \quad (\text{C1.22})$$

$$\pi_t = \rho \pi_{t-1} + \varepsilon_t. \quad (\text{C1.23})$$

The household knows the distribution of future price shocks, but not the path itself. In this case, consumption is no longer a martingale, so I use numerical methods to solve the model as I can no longer solve for current consumption analytically.

This household model is similar to some open economy models, since the interest rate in the model is not dependent on the household’s bond holdings. This means that the steady state bond-holdings in the model are dependent on initial bond holdings and the history of price shocks. To solve the model, I follow Schmitt-Grohé and Uribe (2003) by introducing a debt-elastic interest rate in the model. The higher the household’s level of debt, the higher premium the household will have to pay to service their debt. This is mostly a practical addition to solve the model, but does have economic meaning in the sense that extremely indebted households are riskier for lenders and so must pay higher interest.

The household’s interest rate is then:

$$r_{h,t} = r^* e^{-b_t}. \quad (\text{C1.24})$$

Figure C1 shows the response of the household to a one percent shock to their price index ($\beta = .99, \gamma = .5, \rho = .95, w_t = .1 \forall t$). A persistent shock to the household's price index shock leads to a fall in real and nominal consumption, and an increase in savings.

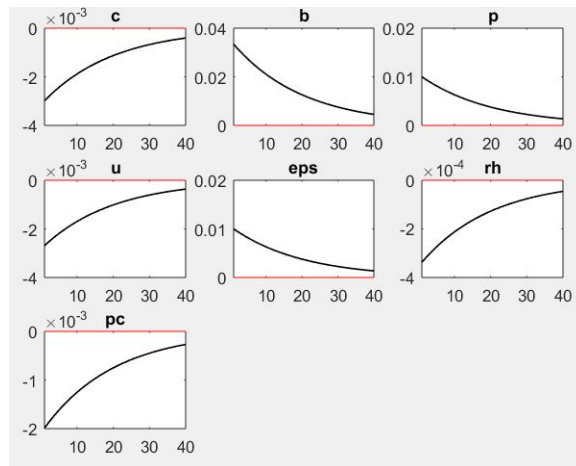



Figure C1. Response to Price Index Shock


C2 Asymmetric Responses to Household Inflation Shocks

In this section I perform a similar analysis as the one used to construct figures 3.3, 3.4, and 3.5, but I allow for asymmetric responses to inflation and deflation shocks. I define a positive inflation shock as $\max\{0, \pi_{ht}\}$ and a negative inflation shock (deflation) as $\min\{0, \pi_{ht}\}$.


C2.1 Price Index




(a) Laspeyres



(b) Paasche




(c) Fisher




(d) Sato-Vartia


Figure C2. Path of Household Retail Price-Index Following Positive Household Inflation Shock
Note: Standard errors robust to auto-correlation and are two-way clustered at the quarter and household level. Confidence interval (99 percent) is shown as the shaded area.




(a) Laspeyres



(b) Paasche




(c) Fisher




(d) Sato-Vartia

Figure C3. Path of Household Retail Price-Index Following Negative Household Inflation Shock
 Note: Standard errors robust to auto-correlation and are two-way clustered at the quarter and household level. Confidence interval (99 percent) is shown as the shaded area.


C2.2 Nominal Consumption




(a) Laspeyres



(b) Paasche




(c) Fisher




(d) Sato-Vartia

Figure C4. Response of Nominal Household Retail Consumption to a Positive one-unit Household Inflation Shock


Note: Standard errors robust to auto-correlation and are two-way clustered at the quarter and household level. Blue shaded area indicates 95 percent confidence interval.




(a) Laspeyres



(b) Paasche



(c) Fisher




(d) Sato-Vartia


Figure C5. Response of Nominal Household Retail Consumption to a Negative one-unit Household Inflation Shock

Note: Standard errors robust to auto-correlation and are two-way clustered at the quarter and household level. Blue shaded area indicates 95 percent confidence interval.


C2.3 Real Consumption




(a) Laspeyres



(b) Paasche




(c) Fisher




(d) Sato-Vartia

Figure C6. Response of Real Household Retail Consumption to a Positive one-unit Household Inflation Shock


Note: Standard errors robust to auto-correlation and are two-way clustered at the quarter and household level. Blue shaded area indicates 95 percent confidence interval.




(a) Laspeyres



(b) Paasche



(c) Fisher



(d) Sato-Vartia

Figure C7. Response of Real Household Retail Consumption to a Negative one-unit Household Inflation Shock

Note: Standard errors robust to auto-correlation and are two-way clustered at the quarter and household level. Blue shaded area indicates 95 percent confidence interval.

C3 Response of Consumption to Inflation Shock by Income Group

C4 Simple model of household inflation rates

From an earlier version of the paper. Shows how the relationship between inflation and inflation dispersion is not simply mechanical.

Suppose firms (k) produce one consumption good and set prices so that:

Table C2. Response of Spending to Household Inflation Shock by Income Group: CEX

	$\ln(P \cdot C)$	$\ln(P \cdot C)$
$\pi_{h,t}$	-0.00149 (0.00220)	
$\pi_{h,t}^+$		-0.00991* (0.00520)
$\pi_{h,t}^-$		-0.00235 (0.00155)
$\pi_{h,t} \times$ 2nd Quartile	-0.00439*** (0.00108)	
$\pi_{h,t} \times$ 3rd Quartile	-0.00660*** (0.00123)	
$\pi_{h,t} \times$ 4th Quartile	-0.00645*** (0.00172)	
$\pi_{h,t}^+ \times$ 2nd Quartile		0.00900*** (0.00312)
$\pi_{h,t}^+ \times$ 3rd Quartile		0.00774** (0.00316)
$\pi_{h,t}^+ \times$ 4th Quartile		0.0169*** (0.00576)
$\pi_{h,t}^- \times$ 2nd Quartile		0.0111*** (0.00252)
$\pi_{h,t}^- \times$ 3rd Quartile		0.0205*** (0.00463)
$\pi_{h,t}^- \times$ 4th Quartile		0.0257*** (0.00633)
N	382,058	382,058

Note: The omitted group is the first income quartile. Standard errors, in parentheses, clustered at the individual and time level and are robust to auto-correlation. Significance at the one percent, five percent and ten percent levels indicated by ***, **, and *. I also include income group and time fixed effects, as well as a lag of the dependent variable.

$$p_{k,t} = \varphi_k \bar{P}_t e^{z_{k,t}}. \quad (\text{C4.1})$$

Where φ_k is a constant markup, \bar{P}_t is the aggregate price level at time t and $z_{k,t}$ is a firm and time specific cost-push shock. ($\sum_{k=1}^N \alpha_k z_{k,t} = 0$). Think of the time specific cost-push shock as coming from the firms decision to change prices in a Golosov and Lucas Jr (2007) style model. Suppose further that this idiosyncratic shock evolves in the following way:

$$z_{k,t} = \rho_k z_{k,t-1} + \varepsilon_{k,t} \quad (\text{C4.2})$$

$$\varepsilon_{k,t} \sim (0, \sigma_{k,t}) \quad (\text{C4.3})$$

Each period the firm gradually adjusts prices back to its normal markup over the aggregate price level, but is also subject to a new shock ($\varepsilon_{k,t}$). Change in prices for firm k can then be given by:

$$\log(p_{k,t}/p_{k,t-1}) = \log\left(\frac{\bar{P}_t}{\bar{P}_{t-1}}\right) + (\rho_k - 1)z_{k,t-1} + \varepsilon_{k,t}.$$

For simplicity, assume that households have Cobb-Douglas style utility over the goods provided by N firms (this corresponds exactly to Laspeyres or Paasche inflation).

Then their price index is:

$$P_{h,t} = \prod_{k=1}^N \left(\frac{p_{k,t}}{\alpha_{h,k}}\right)^{\alpha_{h,k}} \quad (\text{C4.4})$$

Combining the household's price index and the law of motion of firm prices allows me to

solve for the expected value and variance of the households inflation rate:

$$\mathbb{E}[\pi_{h,t}] = \log\left(\frac{\bar{P}_t}{P_{t-1}}\right) + \sum_{k=1}^N \alpha_{h,k}(\rho_k - 1)z_{k,t-1} \quad (\text{C4.5})$$

$$\text{Var}[\pi_{h,t}] = \sum_{k=1}^N \alpha_{h,k}^2 \sigma_{k,t}^2 \quad (\text{C4.6})$$

On average, households should expect their inflation rate to equal aggregate inflation plus the weighted average of the firms prices adjusting back toward their normal level. If I make the simplifying assumption, that the distribution of the cost-push shock is the same for all firms then:

$$\text{Var}[\pi_{h,t}] = \sigma_t^2 \sum_{k=1}^N \alpha_{h,k}^2 \quad (\text{C4.7})$$

The variance of the household's inflation shock (which corresponds to inflation dispersion) depends on two things: (1) the variance of the firms price shocks (price change dispersion, which is a rate and different then the level price dispersion in New Keynesian models). (2) The household's Herfindahl index (preference intensity) over goods.

I test whether this model can explain the relationship I find between inflation dispersion and inflation by testing whether σ_t^2 varies with aggregate inflation. I extract the firm's cost push shocks ($\varepsilon_{t,k}$ from the Nielsen data by running the following series of regressions:

$$\log(p_{k,t}) - \log(p_{k,t-4}) = \beta_0 \bar{\pi}_t + v_{k,t} \quad (\text{C4.8})$$

where $p_{k,t}$ is average price of a unit in Nielsen product module k in quarter t and $\bar{\pi}_t$ is average annual household inflation rate in quarter t . Note that:

$$v_{k,t} = (\rho_k - 1)z_{k,t-4} + \varepsilon_{k,t} = (\rho_k - 1)(\rho_k z_{k,t-8} + \varepsilon_{k,t-4}) + \varepsilon_{k,t} \quad (\text{C4.9})$$

So $\varepsilon_{k,t}$ should be the residuals from regression of $v_{k,t}$ on $v_{k,t-4}$. Note that the coefficient on this regression should be negative (it is). Finally, I regress the standard deviation of the firm's cost push shocks on aggregate inflation. My resulting coefficient is 0.006 with a standard error of 0.002. While this result is positive and statistically significant indicating that this model can help explain the relationship between inflation dispersion and aggregate inflation, it is orders of magnitude too small.

One key assumption I made was that $\sigma_{k,t} = \sigma_t \forall k$. I do not expect that firms should all have the same relationship between the “cost-push” shock and aggregate inflation. For example, firms in some sectors may be more able to adjust their prices to inflation than firms in other sectors. Households buy different goods of varying stickiness (Cravino et al. 2018), which could help explain the large relationship between inflation and inflation dispersion that I see; however, this does not explain the relationship between inflation and inflation dispersion within narrow product categories. Kim (2019) shows that as product price (and by extension quality) within categories increases, price changes become more infrequent. Differences in the average product quality of a household's bundle may then also contribute to inflation dispersion.

Despite its failure to explain the inflation dispersion and inflation relationship, the model's prediction that households with higher preference intensity should have more volatile inflation rates is quite accurate.

I measure each household's inflation volatility as deviations from the aggregate inflation rate as $(\pi_{h,t} - \bar{\pi}_t)^2$ and regress this measure on their Herfindahl index. Table C3 shows the results of this regression using sequential Laspeyres inflation rates in the Nielsen data (note that this is an older result and will be updated).

Figure C8 shows that poorer households on average have a higher preference intensity

over products, which suggests that poor households may have more volatile inflation rates. However, the actual average difference in preference intensity between income groups is not that large.

Table C3. High Preference Intensity \implies Volatile Inflation

	$(\pi_{h,t} - \bar{\pi}_t)^2$	$(\pi_{h,t} - \bar{\pi}_t)^2$
$\sum_{k=1}^N \alpha_{h,k}^2$	20.71	21.05
	(4.121)	(4.207)
$ \bar{\pi}_t $		1.572
		(0.412)
Observations	2,404,480	2,404,480

Standard Errors (in parentheses) robust to auto-correlation and are two-way clustered at the household and quarter levels.




Figure C8. Preference Intensity by Income Group

C5 Inflation Rates and Volatility by Income Group

Table C4. Household Inflation Volatility and Income

Panel A: Nielsen Homescan

Household Income	Laspeyres $(\pi_{h,t} - \bar{\pi}_t)^2$	Paasche $(\pi_{h,t} - \bar{\pi}_t)^2$	Fisher $(\pi_{h,t} - \bar{\pi}_t)^2$	Sato-Vartia $(\pi_{h,t} - \bar{\pi}_t)^2$
\$25k-50k	1.973*** (0.664)	1.452** (0.567)	0.441* (0.254)	0.433 (0.517)
\$50k-100k	1.909*** (0.693)	1.758** (0.684)	0.555* (0.285)	0.741 (0.611)
> \$100k	1.441* (0.858)	1.297* (0.664)	0.377 (0.295)	0.991 (0.720)
Time Fixed Effects	X	X	X	X
N	1,894,135	1,894,135	1,894,135	1,894,135

Panel B: CEX

Household Income	Laspeyres $(\pi_{h,t} - \bar{\pi}_t)^2$	Paasche $(\pi_{h,t} - \bar{\pi}_t)^2$
2nd Quartile	-15.72*** (2.316)	-16.79*** (2.335)
3rd Quartile	-17.13*** (2.299)	-17.81*** (2.300)
4th Quartile	-18.56*** (2.422)	-19.39*** (2.364)
Time Fixed Effects	X	X
N	243,468	243,468

Note: Omitted category is less than \$25,000 household income for the Nielsen Homescan and 1st quartile for the CEX. Standard errors, in parentheses, clustered at the household and quarter levels (Newey-west HAC standard errors were too much for my computer; I will run this on the server later). Significance at the one, five and ten percent levels indicated by ***, **, and * respectively. Aggregate Inflation, $|\bar{p}_t|$ is a democratic index of individual inflation weighted using the population weights in the Nielsen Homescan or CEX. National prices are used throughout.

Table C5. Average Inflation Rates of Income Group Compared to National Rate**Panel A: Nielsen Homescan**

Household Income	Laspeyres $\pi_{h,t} - \bar{\pi}_t$	Paasche $\pi_{h,t} - \bar{\pi}_t$	Fisher $\pi_{h,t} - \bar{\pi}_t$	Sato-Vartia $\pi_{h,t} - \bar{\pi}_t$
\$25k-50k	0.0375* (0.0194)	-0.0917*** (0.0112)	-0.0292*** (0.00981)	-0.0883*** (0.0116)
\$50k-100k	-0.00714 (0.0227)	-0.164*** (0.0145)	-0.0845*** (0.0126)	-0.157*** (0.0169)
> \$100k	-0.0526* (0.0275)	-0.159*** (0.0154)	-0.103*** (0.0141)	-0.150*** (0.0198)
Time Fixed Effects	X	X	X	X
N	2,325,692	2,325,692	2,325,692	2,325,692

Panel B: CEX

Household Income	Laspeyres $\pi_{h,t} - \bar{\pi}_t$	Paasche $\pi_{h,t} - \bar{\pi}_t$
2nd Quartile	-0.120*** (0.0277)	-0.0780*** (0.0248)
3rd Quartile	-0.205*** (0.0357)	-0.146*** (0.0322)
4th Quartile	-0.314*** (0.0364)	-0.266*** (0.0342)
Time Fixed Effects	X	X
N	243,468	243,468

Note: Omitted category is less than \$25,000 household income for the Nielsen Homescan and 1st quartile for the CEX. Standard errors, in parentheses, clustered at the household and quarter levels (Newey-west HAC standard errors were too much for my computer; I will run this on the server later). Significance at the one, five and ten percent levels indicated by ***, **, and * respectively. Aggregate Inflation, $|\bar{p}_t|$ is a democratic index of individual inflation weighted using the population weights in the Nielsen Homescan or CEX. National prices are used throughout.

C6 Effective versus Common Prices

In this project I calculated household inflation rates using the household shares for the particular product as weights and the change in the national average price paid for the product as the price change p_t/p_{t-1} ; however, it is not obvious that using the common/average prices for the product is the most accurate way to calculate changes in a household's inflation rate. Kaplan

and Schulhofer-Wohl (2017) use the price the household actually paid for the product instead (effective price). As I discussed in the text, there are several problems with using the effective versus the common price:

- In order to calculate the change in price between two periods the household must buy that same product in both periods, which is only a small fraction of their total basket (around 25 percent in Kaplan and Schulhofer-Wohl (2017)). Since household behavior is likely shaped by changes in prices this would lead to households endogenously sorting to products that have either not changed their price or lowered their price rather than products that have increased their price; which would bias the household's inflation rate downward.
- Changes in the household's effective price paid could be a result of the household switching stores or using coupons in one period and not using them in the next. In terms of the standard cost of living index (COLI), those using effective prices to calculate household inflation rates should take a stance on the consequences of store-switching and the change in effort (coupons or no coupons, shopping during sales during the quarter, etc) on household utility.
- In my main analysis, I use either the Nielsen product module or the CEX expenditure category as my definition of product. Use of effective prices virtually necessitates using the upc code or the brand as the definition of product, else changes in effective price from one period to another could simply be the result of product switching. I will discuss my rationale for using a slightly larger category of product in another section of the appendix.

I favor using the common (national) rather than the effective price in my analysis for the reasons above and because the common price is exogenous to the household's behavior. However, In the remainder of this section I justify my decision: (1) I show that my results are robust to repeating my main analysis in table 3.4, but using effective rather than common prices; (2) I follow Kaplan and Menzio (2015) in creating a relative price index for each household which

Table C6. Household Inflation Dispersion and Aggregate Inflation with Effective Prices

	Laspeyres $\sigma(\pi_h)$	Paasche $\sigma(\pi_h)$	Fisher $\sigma(\pi_h)$	Sato-Vartia $\sigma(\pi_h)$
$ \bar{\pi} $ (0.372)	0.587 (0.0355)	0.321*** (0.231)	0.462* (0.0511)	0.232***
N	52	52	52	52

Note: Newey-west HAC standard errors in Parentheses. Significance at the one, five and ten percent levels indicated by ***, **, and * respectively. Aggregate Inflation, $|\bar{\pi}|$ is a democratic index of individual inflation weighted using the population weights in the Nielsen Homescan. Effective prices are used throughout.

is based exclusively on the prices that the household pays for their bundle relative to national prices. I show that the distribution of changes in these “relative” price indexes is stable over time; (3) I attempt to create another version of household inflation with effective prices by combining the Nielsen Retail Scanner data with the Consumer panel data (I do this only for one product module due to the massive amount of computing power this requires). I construct inflation rates where the product is defined at the upc-store level, which fixes the household’s choice of store over time. I show that this new measure of inflation is highly correlated with my household level inflation rates using common prices, but not to the Kaplan and Schulhofer-Wohl (2017) style inflation rates using effective prices.

C6.1 Robustness check using effective prices

Table C6 shows the results of repeating my main analysis from table 3.4, but using effective prices rather than common prices. My results are similar in magnitude to the results using common prices, although they are not quite as statistically significant. Here, I define a product as the product module as I do in my main text.

C6.2 Distribution of Relative Price Indexes over time

When measuring changes in the cost of living, how important is using exactly the price that a household pays for a given product? To help answer this question, I follow Kaplan and Menzio (2015) and create relative price indexes in the following manner:

$$RP_{ht} = \frac{\sum_n s_{n,h,t} p_{n,ht}}{\sum_n s_{n,h,t} p_{n,t}}, \quad (\text{C6.1})$$

where $s_{n,h,t}$ is the household (h) specific consumption share of good n at time t, $p_{n,ht}$ is the price that the household actually pays for product n, and $p_{n,t}$ is the national average price for the product. RP_{ht} is equal to one when the household on average pays the same prices for products in their bundle as national prices, while $RP_{ht} < 1$ or $RP_{ht} > 1$ implies that the household is buying their bundle at a discount or a premium respectively. The relative price index is a convenient way to separate changes in the cost of living into changes in prices the household pays compared to national prices and differences in consumption shares compared to the national average.

Figure C9 shows the distribution of relative price indexes for all households in 2017Q1 in blue. For comparison, I also included the standard normal pdf in red. The relative price index has less variance and heavier tails than the normal and is slightly skewed to the right.

Does the household's position in the relative price index distribution change over time? That is, do households frequently switch from paying a premium for their bundle to getting their bundle at a discount? Table C7 shows the results of regressing a household's RP_{ht} on their relative price index from the previous quarter. Column 1 presents the baseline results while column 2 adds additional controls for changing marital status, household size and income level. I find that there is a strong correlation between a household's relative price index this period and last period, however this correlation is not one; so households do shift around in the distribution




Figure C9. Distribution of Relative Price Indexes: 2017Q1

Note: Nielsen Consumer Expenditure Survey 2017Q1.

Table C7. Household relative price distribution AR(1)

	(1) RP_{ht}	(2) RP_{ht}
RP_{ht-1}	0.746*** (0.00920)	0.743*** (0.00963)
Δ Household Size		0.000235 (0.000156)
Δ Marital Status		-0.000493*** (0.000180)
Time FE	X	X
Household Income FE		X
N	2,972,187	2,972,187

Note: Newey-west HAC standard errors in Parentheses clustered at the quarter household level. Significance at the one, five and ten percent levels indicated by ***, **, and * respectively.




Figure C10. Distribution of changes in RP_{ht} over time

Note: Nielsen Consumer Expenditure Survey 2004-2017.

of relative prices (possibly because of store switching, using coupons, etc.).

How important are changes in household RP_{ht} for shifts in the distribution of household inflation rates? Figure C10 shows the distribution of changes in the relative price index (defined as $\frac{RP_{ht} - RP_{ht-1}}{RP_{ht-1}}$) over time. The distribution is relatively constant except for a spike in 2012 (heretofore unidentified data issues with prices in the Nielsen Homescan). Since the distribution is constant, changes in the distribution of relative price indexes cannot explain the relationship we find between inflation and inflation dispersion.

C6.3 Household inflation rates using store-specific prices

When a household buys a product the price the household pays for the product n can be decomposed into three different components:

$$p_{n,h,t} = \underbrace{p_{n,t}}_{\text{Common Price}} \times \underbrace{S_{n,h,t}}_{\text{Store Premium}} \times \underbrace{\varphi_{n,h,t}}_{\text{Shopping Effort Premium}} . \quad (\text{C6.2})$$

As discussed before, when trying to calculate household inflation rates using changes in $p_{n,h,t}$ the prices $p_{n,h,t}$ are only observed for products the household buys in both periods (common products). Furthermore, the change in $p_{n,h,t}$ could be a result of changes in shopping effort $\varphi_{n,h,t}$ or a change in $S_{n,h,t}$ because of store switching. Rather than taking a stand on household utility from shopping effort or store switching or trying to extrapolate changes in the cost of living for common products on non-common products, I use the Nielsen Retail Scan data to create a more complete measure of household inflation.

The Nielsen Retail Scan (RMS) data is point of sale data from over 30,000 retail stores that includes the revenue and quantity sold for every product (defined at the upc level) during each week (it is from these data that the Nielsen Consumer Panel extrapolates household level prices for goods purchased at participating stores). The Nielsen Consumer panel has a store code for purchases from stores that are also listed in the RMS data. For select product modules (currently cereal) I combine the Consumer Panel and the RMS data and calculate household level prices as:

$$p_{n,h,t-k}^t = p_{n,t-k} \times S_{n,t-k}^{n,h,t} . \quad (\text{C6.3})$$

Here, $S_{n,t-k}^{n,h,t}$ is the store premium for the store that the household bought good n at at time t . This allows me to fix shopping effort and store rather than take a stance on the relationship of shopping effort and store switching with household utility. Another benefit is that since stores stop selling products much less often than households stop buying products I no longer need to rely on goods that are common to the household's basket in both periods.

So household inflation is defined as:

$$\pi_{h,t}^S = \sum_n w_{n,h,t} \frac{P_{n,h,t}^t}{P_{n,h,t-4}^t}, \quad (\text{C6.4})$$

where the period is a quarter and I look at the change in prices four quarters back to control for seasonality. Finally, $w_{n,h,t}$ is the household weight on product n (in this example I use Laspeyres weights). Not all household purchases are from a store that is in the RMS data. Non RMS-store purchases simply use the common price instead of $P_{n,h,t-k}^t$.

Constructing this household inflation measure is computationally intense. Rather than construct this measure for all product modules I test how similar this measure is to my previously constructed inflation measures: π^c common prices, π_r regional prices, and π_f effective prices (which corresponds to the inflation rates in Kaplan and Schulhofer-Wohl (2017)). Table C8 shows the results of this comparison. There is a high correlation between the common and region price inflation rate measures that I use in the text and the store price measure. The effective price measure is actually negatively correlated to the store price inflation measure (perhaps because of the selection bias issue that I raised earlier).

Table C8. Household Inflation Rates with Store Specific Prices v. Other Inflation Measures

	π^s	π^s	π^s
π^c	0.713*** (0.0270)		
π^r		0.730*** (0.0134)	
π^f			-0.0131*** (0.00317)
Household FE	X	X	X
Time FE	X	X	X
N	428,879	428,879	428,878
R-squared	0.421	0.753	0.389

Note: Nielsen Consumer Panel and Nielsen Retail Scan data. Only includes cereal products. Products defined at the upc level. Standard errors in parentheses. Significance at the one, five and ten percent levels indicated by ***, **, and * respectively.