

THE UNIVERSITY OF CHICAGO

ESSAYS IN MACROECONOMICS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF ECONOMICS

BY

MARTIN BERAJA

CHICAGO, ILLINOIS

JUNE 2016

To Eze, my friend, who lives on in my memory.

I miss you every day...

CONTENTS

List of Figures	vi
List of Tables	viii
Acknowledgements	ix
Abstract	x
1 A Semi-Structural Methodology for Policy Counterfactuals with an Application to Fiscal Unions	1
1.1 Introduction	1
1.2 Literature review	6
1.3 A Simple Fiscal Union Model	9
1.3.1 Firms and Households	10
1.3.2 Federal government	11
1.3.3 Exogenous shocks and processes	12
1.3.4 Equilibrium	12
1.3.5 Aggregation	13
1.4 Transfers Policy in Alternative Models of Fiscal Unions	14
1.4.1 Different mechanisms, same counterfactual...	16
1.5 Semi-Structural Transfer Counterfactual in Fiscal Union Models	18
1.5.1 A set of linear fiscal union models	18
1.5.2 Equilibrium representations	20
1.5.3 Semi-structural transfer policy counterfactual	20
1.6 A Counterfactual United States Economy without Fiscal Integration	22
1.6.1 The sufficient semi-structure ξ^s	22
1.6.2 Identification of $\Lambda(\xi, \Theta)$ and shocks	23
1.6.3 Data description	25
1.6.4 Estimation of transfer policy rule Θ	28
1.6.5 Main Findings	31
1.6.6 Channels of regional stabilization	34
1.6.7 Robustness to alternative policy specifications	36
1.7 Semi-Structural Policy Counterfactuals: A General Methodology	38
1.7.1 Implementation: How to describe the semi-structure ξ^s	41
1.8 Falsifying a Set of Models via the Semi-Structural Methodology	44
1.9 Conclusions	45

2	The Aggregate Implications of Regional Business Cycles (with Erik Hurst and Juan Ospina)	47
2.1	Introduction	47
2.2	Creating State-Level Price And Wage Indices	52
2.2.1	State-Level Wage Index	52
2.2.2	State-Level Price Index	54
2.3	Comparing Cross-State Patterns to Aggregate Time-series Patterns	59
2.4	A Monetary Union Model	63
2.4.1	Firms and Households	64
2.4.2	Sticky Wages	65
2.4.3	Equilibrium	66
2.4.4	Shocks	67
2.4.5	Aggregation	67
2.4.6	Aggregate vs. Local Shock Elasticities	69
2.5	A Procedure for Identifying Aggregate Shocks	72
2.6	Estimating the Wage Setting Equation Using Regional Data	75
2.6.1	Estimating Equation and Identification Assumptions	76
2.6.2	Estimates of λ and ϕ	80
2.7	The US Great Recession: From Regional to Aggregate	82
2.7.1	Findings in the Aggregate	82
2.7.2	Sensitivity to Alternative Parameters λ and ϕ	85
2.7.3	Sensitivity to Alternative Identifying Assumptions	86
2.7.4	Discussion: Aggregate v. Regional Shock Decomposition	89
2.8	Conclusion	89
3	From Hyperinflation to Stable Prices: Argentina’s Evidence on Menu Cost Models (with Fernando Alvarez, Martin Gonzalez-Rozada and Pablo Andres Neumeyer)	92
3.1	Introduction	92
3.2	Comparative Static Properties of Menu Cost Models	96
3.2.1	The Basic Menu Cost Model	96
3.2.2	Illustrating the theory with an example	108
3.3	Description of the Dataset	116
3.4	Estimating the Frequency of Price Changes	118
3.5	Argentina’s Evidence on Menu Cost Models of Price Dynamics	119
3.5.1	The Frequency of Price Changes and Inflation	120
3.5.2	Inflation and the Dispersion of the Frequency of Price Changes	125
3.5.3	Inflation and the Intensive and Extensive Margins of Price Adjustments	126
3.5.4	Inflation and the Dispersion of Relative Prices	130
3.6	Conclusions	132
	Bibliography	134

A	A Semi-Structural Methodology for Policy Counterfactuals with an Application to Fiscal Unions	145
A.1	Proof of Lemma 1	145
A.2	Proof of Claim 1	147
A.3	Proof of Proposition 1	147
A.4	Proof of Proposition 2	149
A.5	Proof of Lemma 2	150
A.6	Proof of Lemma 3	150
A.7	Figures and Tables	151
B	The Aggregate Implications of Regional Business Cycles	152
B.1	Proof of Lemma 4 and 5	152
B.2	Alternative wage setting specifications	154
B.2.1	Preferences with wealth effects in labor supply	154
B.2.2	Forward looking wages	155
B.3	Figures and Tables	156
C	From Hyperinflation to Stable Prices: Argentina’s Evidence on Menu Cost Models	164
C.1	Proofs of comparative static results	164
C.1.1	Proof of Proposition 3	164
C.1.2	Proof of Lemma 6	166
C.1.3	Proof of Lemma 7	167
C.1.4	Proof of Proposition 5	167
C.2	Analytical characterization of the model in Section 3.2.2	174
C.2.1	Some graphical illustrations of the model in Section 3.2.2	189
C.3	Classification of Gand Services and Data Collection	189
C.3.1	Instructions to CPI’s Pollsters	193
C.4	Robustness	194
C.4.1	Missing data, substitution, sales and aggregation	194
C.4.2	Estimation of the elasticities at a more disaggregated level	208
C.4.3	Sampling Periodicity	209
C.4.4	Expected Inflation	210
C.5	Background on Inflation and Economic Policy	212
C.6	Comparison to other studies	217

LIST OF FIGURES

1.1	Net Transfers growth v. Wage Income growth (2006-2010)	29
1.2	Impulse Response Function to a Demand shock	31
1.3	Employment Gains from Fiscal Integration by State in 2010	33
2.1	State Employment Growth vs. State Nominal and Real Wage Growth, 2007- 2010	60
2.2	Counterfactual Employment Response	84
3.1	Optimal thresholds for different inflation rates	112
3.2	Frequency of Price Adjustments λ_a and of Price Increases λ_a^+	113
3.3	Standard deviation of log relative prices, conditional on $z = 0$	115
3.4	Number Outlets per Good, and Frequencies of Substitution and Sales	117
3.5	Estimated Frequency of Price Changes λ and Expected Inflation	119
3.6	The Frequency of Price Changes (λ) and Expected Inflation.	122
3.7	The Frequency of Price Changes (λ) and Expected Inflation: International Evidence	124
3.8	Decomposition of inflation for low inflation rates	127
3.9	Inflation and the Extensive Margin of Price Adjustments	129
3.10	Inflation and the Intensive Margin of Price Adjustments	130
3.11	Average Dispersion of Relative Prices and Inflation	131
A.1	Employment in the Great Recession: Channels Decomposition	151
B.1	Nielsen Retail Price Index vs. CPI Food Price Index	156
B.2	The Evolution of Aggregate Real and Nominal Composition Adjusted Wages	157
B.3	Impulse Response to a Discount rate Shock	157
B.4	Impulse Response to a Productivity / Markup Shock	158
B.5	Impulse Response: Leisure Shock	158
B.6	Shock time-series	159
B.7	Counterfactual Price Response	159
B.8	Counterfactual Wage Response	160
B.9	State Net Migration Rate 2009-2010 vs. State Employment Growth 2007-2010	160
C.1	Average size of price increases Δ_p^+ and decreases Δ_p^-	190
C.2	Unconditional standard deviation of log relative prices $\bar{\sigma}$	191
C.3	Estimators $\hat{\lambda}$ and f , sampled every two weeks	210
C.4	Estimators λ and f , homogeneous goods sampled twice a month	211

C.5	Inflation, Money Growth and Deficits	213
C.6	Inflation and Depreciation Rate during HyperInflations	216

LIST OF TABLES

1.1	Policy rule baseline estimates	30
1.2	Employment statistics: Fiscal Integration v. Fiscal Autonomy	32
1.3	Employment in the Great Recession: Channels Decomposition	36
2.1	Comparison of Cross-State and Time-Series Estimates of Wage Elasticities During the Great Recession	62
2.2	Estimates of λ and $\frac{\lambda}{\phi}$ using Cross- Region Data	80
3.1	Cross Industry Dispersion of Duration $\frac{1}{\lambda}$	126
A.1	Counterfactual Employment Dispersion without Fiscal Integration	151
B.1	Calibration	161
B.2	Aggregate v. Regional Employment Impact Elasticities	161
B.3	Aggregate Elasticities ($\lambda=0.7, \phi=2$)	162
B.4	Regional Elasticities ($\lambda=0.7, \phi=2$)	162
B.5	Discount/Interest rate (γ) and Productivity/Markup (z) shocks' contribution to aggregate employment change	163
C.1	Example of the Harmonized Index of Consumer Price Classification	190
C.2	Goods ordered by weight whose prices are gathered once and twice per month	192
C.3	Weights for Divisions of the Harmonized Index of Consumer Price	192
C.4	The Frequency of Price Adjustment and Inflation: Robustness Checks.	196
C.5	Example of a spell of constant prices without missing prices	198
C.6	Example of a spell of constant prices with observed and missing prices	199
C.7	Example of a spell of constant prices, w/ missing prices and counters	200
C.8	Example of a spell of constant prices removing sales	202
C.9	Example: spell of constant prices w/ counters for substitutions	203
C.10	Distribution of fitted coefficients at 5 digit level	209
C.11	Elasticities and implied duration for different estimates of expected inflation.	212
C.12	Comparison with other Studies in Countries with High Inflation	217

ACKNOWLEDGEMENTS

I owe a debt of gratitude to my advisors—Fernando Alvarez, Erik Hurst, and Veronica Guerrieri—that I cannot possibly repay. For their encouragement, insightful comments, and infinite patience in making sense out of my many incoherent ideas over the years. I can only hope I am a mentor to my students as they were to me.

I thank Juliette, *mon coeur*, for giving me a caring home when I had left mine, for our many challenging and fun conversations, and for carefully reading this dissertation to improve its clarity. I came in pursue of a PhD and found someone to share my life with. I could not have made this journey without her.

I thank Monica and Victor, my parents, for every hug and loving word, and for always giving me the freedom to choose my way. They taught me of family and generosity. I am, because of them.

I thank Leila, my sister, for her support and kindness. I wish we keep growing separately without growing apart.

I thank Diego, Fede and Javi, *mis hermanos de la vida*, for being part of almost every story I can remember.

I thank Jonathan Adams, Chanont Banterghansa, Philip Barrett, Artur Bezerra de Carvalho, Jörn Boehnke, Mariano Lanfranconi, Alex McKay, Roberto Robatto and Francisco Roch, my friends and officemates in “The Basement” and “The Tower”, for our many hours discussing economics and sharing laughs.

Finally, I would like to thank the University of Chicago for giving me the chance and financial support to thrive in such an outstanding scholarly environment. My years here changed me forever.

ABSTRACT

In the first chapter of this dissertation, I propose a methodology to conduct counterfactual analysis in a way that is robust to specific assumptions about primitives of linear models of dynamic stochastic economies. Then, I apply the methodology to quantify how fiscal unions contribute to regional stabilization. I start by showing how to identify a set of models that yield the same counterfactual equilibrium after a policy change by imposing restrictions directly on equilibrium equations. Next, I describe how to construct this counterfactual equilibrium using data under a benchmark policy. The methodology allows obtaining quantitative predictions with respect to policy changes with minimal *a-priori* structural assumptions while being immune to Lucas critique, enhancing credibility of the analysis. In the application to fiscal unions, I focus on models where the federal government redistributes resources via a *transfer policy rule*, which is a function of local variables, in order to smooth local shocks. Using US state-level data, I construct a counterfactual US economy without the rule in place. This counterfactual is identical in many fiscal union models with rich features, such as nominal rigidities and asset market incompleteness. My primary finding is that during the Great Recession fiscal integration significantly reduced cross-state employment differences by redistributing resources from states that were doing relatively well to states that were doing relatively poorly. Finally, I discuss how the methodology can further be used to falsify a set of models, provided data before and after a policy change are available.

In the second chapter of this dissertation (joint with Erik Hurst and Juan Ospina), we study the aggregate implications of regional business cycles. We argue that it is difficult to make inferences about the drivers of aggregate business cycles using regional variation alone because (i) the local and aggregate elasticities to the same type of shock are quantitatively different and (ii) purely aggregate shocks are differenced out when using cross-region variation. Then, we highlight the importance of these issues in a monetary union model, and by contrasting the behavior of US aggregate time-series and cross-state patterns during the Great Recession. In particular, using household and retail scanner data for the US, we

document a strong relationship across states between local employment growth and local nominal and real wage growth. These relationships are much weaker in US aggregates. Finally, in order to identify the shocks driving aggregate (and regional) business cycles, we develop a methodology that combines regional and aggregate data. The methodology uses theoretical restrictions implied by a wage setting equation that holds in many monetary union models with nominal wage stickiness. We show how to estimate this equation using cross-state variation—thus linking particular regional patterns to particular aggregate shock decompositions. Applying the methodology to the US, we find that a combination of both “demand” and “supply” shocks are necessary to account for the joint dynamics of aggregate prices, wages and employment during the 2007-2012 period while only “demand” shocks are necessary to explain most of the observed cross-state variation. We conclude that the wage stickiness necessary for demand shocks to be the primary cause of aggregate employment decline during the Great Recession is inconsistent with the flexibility of wages estimated from cross-state variation.

In the third and last chapter of this dissertation (joint with Fernando Alvarez, Martin Gonzalez-Rozada, and Pablo Andres Neumeyer), we analyze how inflation affects firms’ price setting behavior. In a class of menu cost models, we derive several predictions about how price setting changes with inflation at very high and near-zero inflation rates. Then, we present evidence supporting these predictions using firm-level-data underlying Argentina’s consumer price index from 1988 to 1997—a unique experience where monthly inflation ranged from almost 200 percent to less than zero. First, we show that the frequency of price changes, the dispersion of relative prices, and the absolute size of price changes do not change with inflation when inflation is low. Second, we find that the frequency and size of price increases and decreases are symmetric around zero inflation. Third, we find that inflation changes near zero-inflation are mostly accounted for by changes in the difference between the frequency of price increases and decreases. Finally, we show that, when inflation is high, the frequency of price changes across different products becomes similar and the relation between the frequency of price changes, the dispersion of relative prices, and the absolute size of price changes are consistent with those predicted by menu cost models with no idiosyncratic shocks. Our findings reflect how inflation swamps idiosyncratic firm shocks as a motive for changing prices in high inflation economies whereas the opposite is true in low inflation economies. Moreover, they confirm and extend available evidence for countries that experienced either very high or near-zero inflation.

Chapter 1

A SEMI-STRUCTURAL METHODOLOGY FOR POLICY COUNTERFACTUALS WITH AN APPLICATION TO FISCAL UNIONS

1.1 INTRODUCTION

During the Great Recession many US states were hit with large negative shocks that depressed their economies. How would they have fared had they not been members of a fiscal union? The most common approaches to answer this question, as well as many others in economics, are the *structural approach* and the *reduced-form approach*. The former relies on specifying primitives of a particular model and identifying its parameters whereas the later is model-free but typically subject to the critique in Lucas (1976).

In this paper, I propose a methodology to conduct counterfactual analysis in a way that is robust to particular assumptions about model primitives and parameterizations while being immune to Lucas critique. First, I show how to identify a set of linear models that yield the same counterfactual equilibrium after a policy change in dynamic stochastic economies. This is achieved by imposing restrictions directly on structural equations characterizing the equilibrium. Second, I describe how to construct this counterfactual equilibrium using data under a benchmark policy. Thus, the methodology allows researchers and policy-makers to obtain quantitative predictions with respect to changes in systematic components of policy¹ with minimal *a-priori* structural assumptions, enhancing credibility of the analysis. This is the original motivation of Sims (1980) as well as more recent literature using *sufficient statistics*.

I apply this methodology to quantify how fiscal unions contribute to regional stabilization by redistributing resources between its members through the federal tax-and-transfer system.

¹As opposed to analysis with respect to exogenous components of policy, i.e, policy shocks.

This application provides much needed quantitative results that go beyond existing reduced-form calculations and calibration exercises in specific models,² and inform current discussions on European fiscal integration. In addition, I use this application to transparently illustrate the methodology. I focus on models where the federal government redistributes resources between the union members via a *transfer policy rule* that is a function of local variables. These transfers smooth the response to local, temporary shocks.³ Using US state-level data, I construct a counterfactual US economy without a transfer policy rule. This counterfactual is identical in a set of fiscal union models whose equilibrium equations satisfy a number of simple exclusion restrictions. The set is interesting because it encompasses models with rich features, such as nominal rigidities, adjustment costs, asset market incompleteness, and correlated exogenous processes. I find that, absent the transfer policy rule, employment differences across states in the Great Recession would have been significantly larger. This is because resources were redistributed from states that were doing relatively well to states that were doing relatively poorly. I conclude that US fiscal integration contributed to stabilization of regional economies.

The paper is divided into two parts. The first part is concerned with the application to fiscal unions and the second part with the generalization of the methodology to other applications. In the first part of the paper, I start by describing several models of an economy with many islands (or states) that are hit by idiosyncratic local shocks, and where asset markets are incomplete. For instance, I present models with and without nominal rigidities, or where leisure is a durable good. The goal is to show examples where the federal transfer policy rule can be evaluated, emphasizing their differences and commonalities. Moreover, I discuss the economic intuition behind the methodology using two of these models. The two models have very different economic mechanisms (i.e., leisure durability v. wage rigidities) but are observationally equivalent and produce identical counterfactual equilibria for alternative transfer policy rules.

Next, I present the methodology to quantify the contribution of fiscal unions to regional stabilization with minimal *a-priori* structural assumptions. I focus on fiscal union models that satisfy three properties. The first is that state-level economies in log-deviations from the aggregate union behave to a first-order approximation as if they were small open

²See Section 1.2 for examples and notable exceptions.

³As an analogy, a Taylor Rule for the nominal interest rate is a function of aggregate variables and smooths response to aggregate shocks.

economies—economies that are independent of others in the union and differ only in the realizations of purely idiosyncratic shocks that hit them.⁴ The second property is that a sufficient set of variables for characterizing the state-level equilibrium in log-deviation from aggregates includes employment, nominal wages, and assets, in addition to exogenous processes for the discount rate of households, wealth of households, and productivity of firms. The third is that the federal government gives lump-sum transfers that are a function of state-level employment, wages, and assets.⁵ This function is the transfer policy rule. Following this, I write a linear dynamic system of equations that characterizes equilibria in these models. It involves three equations: a log-linearized Euler equation, a sequential budget constraint, and an equation describing the labor market equilibrium. Moreover, I assume that these equilibria have a unique, stable structural vector autoregression (SVAR) representation.

Then, I present the main theoretical result that allows me to construct a counterfactual equilibrium without a transfer policy rule. I show that (i) knowledge of the SVAR representation of the equilibrium with a transfers policy rule and (ii) imposing a number exclusion restriction on the three equilibrium equations is sufficient to identify the set of fiscal union models that yield identical counterfactual SVAR representation of the equilibrium without a transfer policy rule. Furthermore, I show that the counterfactual SVAR representation can be constructed provided knowledge of the transfer policy rule. The result is useful because many rich and interesting models of fiscal unions (e.g., with nominal wage rigidities, durable leisure, asset market incompleteness, and correlated exogenous processes) are indeed consistent with the exclusion restrictions I impose.

I build the result on the following insights. Log-linearized fiscal union models can be described by the coefficients associated with state-level variables in the three equilibrium equations. These coefficients involve combinations of parameters resulting from model primitives (e.g., preferences and technologies). I note that models with different primitives or parameterizations but identical values for the coefficients have identical SVAR representation of the equilibrium. The standard approach to find such equilibrium is to describe the primitives, pick parameter values, and solve a *non-linear* system of equations for the SVAR given coefficient values.⁶ However, the system of equations is *linear* in the coefficients given

⁴This is an assumption about aggregation of linear economies with idiosyncratic shocks. The key conditions are that identical parameters associate with the variables in each system characterizing the equilibrium, and the idiosyncratic shocks “wash out” when aggregated. The special cases I described in the previous section satisfy these.

⁵In Section 1.6.7, I consider the case where taxes and transfers distort decisions at the margin.

⁶For example, by guessing a recursive equilibrium exists and then using the method of undetermined

values for the SVAR. This implies that imposing a number of exclusion restrictions (i.e., zero coefficients) makes the system exactly determined, thus identifying the set of models that are observationally equivalent under a given transfer policy and yield identical counterfactual equilibrium without it. For example, I impose that assets do not enter in the Euler equation or the discount rate process does not enter in the sequential budget constraint, which is true in many models of fiscal unions. Lastly, all identified coefficients in the Euler and labor market equations are invariant to policy changes because the transfer policy rule only appears in the sequential budget constraint in an additive fashion. This implies that, given knowledge of the rule in place, it is possible to construct the counterfactual equilibrium without the transfer policy rule.

I follow by constructing a counterfactual US economy without a transfer policy rule redistributing resources between states. I use annual state-level data on employment, wages, and assets between 2006 and 2011 to estimate all inputs required to implement the semi-structural methodology. I estimate a regional SVAR where I identify the structural shocks and impulse response matrix by using a set of theoretical restrictions on the joint response of variables in the SVAR to each type of shock. These restrictions are derived from the parameterized sequential budget constraint. This identification procedure extends results in Beraja et al. (2015b) to the larger class of models considered here.⁷ This is an ancillary contribution of this paper. Furthermore, I estimate the federal transfer policy rule elasticities with respect to state-level employment, wages, and assets. I use both OLS and, to deal with potential bias due to reverse causality, two-stage least squares. I instrument potentially endogenous variables using local house prices and the SVAR shocks that are orthogonal to the innovation in the policy rule but correlated with the endogenous variables.⁸ Results are similar in all specifications. I find that net transfers increase between 20 and 30 percent for every 1 percent decrease in nominal total wage income, and asset variation is not significant in explaining transfers variation once wages and employment are considered.⁹

I conclude the first part of the paper by presenting primary findings of the counter-coefficients.

⁷Identification of the impulse response matrix relies on a series of theoretical linear restrictions that are implied by these equations, and link the reduced form errors to the structural shocks.

⁸OLS estimates are consistent if the error term in the policy rule has no regional component (or if the error term is measurement error). However, the issue of endogeneity might arise because of reverse causality. When the innovation in the transfer policy rule causes regional variation in employment, wages, and assets, they both cause and are caused by transfers.

⁹The estimates are similar to those in Feyrer and Sacerdote (2013), who find a 0.25 decrease, and Bayoumi and Masson (1995), who find a 0.31 decrease.

factual exercise. The employment rate’s cross-state standard deviation in the US in 2010 was 2.6 percent. I find that the standard deviation would have been 3.5 percent without the transfer policy rule redistributing resources from well performing to poorly performing states. The results are similar for 2009 and 2011, and in the stationary distribution (which I construct using Monte Carlo simulations). To put these magnitudes in context, aggregate output volatility during the pre-Volcker period (1960:1 to 1979:2) was 1.8 whereas during the post-Volcker-disinflation period (1982:4 to 1996:4) it was 1.3.¹⁰ The stabilization role of fiscal integration seems to be in the same order of magnitude than this “Great Moderation”. Moreover, I assess how nominal and real rigidities interact with fiscal integration—the *rigidities channel*—by constructing a semi-structural transfer policy counterfactual in an economy without nominal or real rigidities, i.e., an economy where labor supply and demand are static. This allows me to conduct a theoretical difference-in-differences experiment to quantify the interaction between rigidities and the transfer policy rule. I find that rigidities amplified the gains from fiscal integration in poorly performing states during the Great Recession, accounting for one-third of their counterfactual employment gains, whereas it accounted for none of the counterfactual losses in the well performing states. Finally, I show that main quantitative findings are insensitive to alternative parameterizations and specifications of the policy rule. In particular, I allow for the possibility that taxes and transfers distort labor supply decisions.¹¹

In the second part of the paper, I generalize the methodology to a larger class of models that are commonplace in macroeconomics. I start by describing a class of linear dynamic systems that characterize (or approximate) stable equilibria in rational expectations models. I emphasize the distinction between the part of the system that is invariant to policy, i.e., the *structure* of the economy, and the part that is not, i.e., the *policy*. Within this class of models, I define a *semi-structural policy counterfactual* as a mapping from (i) the SVAR representation of an economy under a benchmark policy, (ii) a subset of the economy’s policy-invariant structure (i.e., a *semi-structure*), and (iii) the benchmark and alternative policies onto the SVAR representation of an economy under the same structure with the alternative policy. The alternative SVAR representation is identical in the set of models

¹⁰See Clarida et al. (2000) for details of this calculation and analysis of interest rate policy stabilization.

¹¹Findings under this alternative policy are essentially unchanged because the estimated transfer policy rule elasticities with respect to wages and employment are close to 1. What is important for determining the agents marginal decisions is the marginal effect of employment (or wages) in taxes and transfers. When coefficients are close to 1, the marginal effect is close to zero.

that are consistent with the semi-structure. I show sufficient conditions for existence and uniqueness of the mapping, as well as discuss implementation. Finally, although I do not pursue this application in this paper, I discuss how to use the methodology to falsify a set of models. Provided data before and after a policy change is available, it is possible to derive a simple *Wald χ^2 test* that compares the reduced-form VAR representation predicted by a set of models to the observed representation after the policy change. The set of models is rejected if these are “too far apart” in a statistical sense.

The rest of this paper is organized as follows. Section 1.2 discusses related literature. Section 1.3 and Section 1.4 introduce examples of fiscal union models. Section 1.5 illustrates the methodology using fiscal union models. Section 1.6 constructs a counterfactual US economy without fiscal integration and presents primary findings. Section 1.7 lays out the general methodology. Section 1.8 shows how to falsify a set of models via the methodology. Section 1.9 offers concluding remarks.

1.2 LITERATURE REVIEW

This paper relates to much literature. The original motivation of Sims (1980) seminal contribution on structural vector autoregressions (SVARs) was to do policy analysis in a way that was robust across many models. SVARs have been very successful for doing analysis with respect to the exogenous component of policy (i.e., policy shocks) in a relatively model-free way. However, because of Lucas critique, it has not been possible to use SVARs to do analysis with respect to the systematic component of policy (i.e., policy rules).¹² The semi-structural methodology allows their use by restricting the set of structural models that generate the SVAR representation of the equilibrium, while being immune to Lucas critique and retaining robustness to model selection within this restricted set.

Furthermore, since seminal research by Harberger (1964), sufficient statistics have been used to evaluate the consequences of policy changes, without requiring full knowledge of the underlying structural model. Chetty (2008) provides an excellent survey. More recent examples in trade literature include Arkolakis et al. (2012) and Blaum et al. (2015), and in macroeconomics, Alvarez et al. (2014a), Auclert (2014), and Atkeson and Burstein (2015). The non-parametric counterfactuals in Adao et al. (2015) are in the same spirit. The semi-

¹²In the application to fiscal unions, policy shocks would be exogenous, unanticipated transfers to a state. On the other hand, I am concerned with the systematic, anticipated component of tax-and-transfer policy as a known function of state-level variables—a transfer policy rule.

structural methodology shares their motivation and philosophy, using reduced-form statistics to construct counterfactuals that are identical in a given set of models.

Economic literature on robust control and estimation, pioneered by Hansen and Sargent (see Hansen and Sargent (2008) for a summary), is also concerned with model misspecification and its consequences for policy decision-making. Their analysis is different from the one in this paper. However, we both emphasize the study of these issues in dynamic, stochastic economies. Consequently, we use similar techniques from time series econometrics (e.g. Hamilton (1994); Stock and Watson (2001); DeJong and Dave (2011) and for solving such models (e.g. Blanchard and Kahn (1980); Uhlig (1995))

I extend the scheme proposed by Beraja et al. (2015b) for identification of structural shocks in SVARs. This identification scheme fits nicely with the semi-structural philosophy in the paper by using certain parts of the theoretical model about which we feel more confident. It is part of growing literature developing “hybrid” methods that, for example, construct optimal combinations of econometric and theoretical models (Carriero and Giacomini (2011); Del Negro and Schorfheide (2004)), or use the theoretical model to inform the econometric model’s parameter (An and Schorfheide (2007); Schorfheide (2000)). Other identification schemes that use restrictions derived from theory include sign restrictions as in Uhlig (2005), long run restrictions as in Blanchard and Quah (1989)), and short-run restrictions as in Clarida and Gertler (1997); Christiano et al. (1999).

Del Negro and Schorfheide (2009) are also concerned with doing policy analysis in misspecified models. In particular, they study changes in the nominal interest rate policy rule. They consider a benchmark DSGE model with sticky prices and a SVAR representation of the equilibrium in this model. Then, they consider an alternative “true” model that might be different from the benchmark and is unknown. They write the SVAR representation of the equilibrium in this “true” model as a sum of the benchmark SVAR and perturbation terms. They consider various assumptions about the perturbation terms concerning how they are affected by changes in policy. Taking a stance on this is necessary to do counterfactuals that are immune to Lucas critique within their framework. For example, they conduct the analysis assuming that the perturbation terms are invariant to changes in the policy rule and that they can be estimated jointly with the benchmark DSGE model parameters from past data. This approach is useful for considering the sensitivity of results to model misspecification around the benchmark model. The semi-structural methodology in this paper is rather different, although this difference is subtle. While I do not describe alternative mod-

els in terms of deviations from a benchmark model in the paper, it is possible to do so and useful for comparison with their paper. In this case, I would write the structural equations characterizing the equilibrium in the “true” model as the sum of a benchmark model and a perturbation term. Then, I would state which perturbation terms are policy invariant and which ones are not. The key difference is that my perturbation terms are placed directly in the structure of the model (e.g. Euler equations, labor demand, etc.), where it is easy to see which models are consistent with these terms being policy invariant, whereas Del Negro and Schorfheide (2009) place them directly on the SVAR representation, where most models would imply that the terms are not policy invariant.

The theory behind the stabilization benefits of fiscal integration, or more generally aggregate risk sharing arrangements, is well developed. Examples include Farhi and Werning (2014); Bucovetsky (1998)); Persson and Tabellini (1996a,b); Lockwood (1999). However, there is surprisingly little quantitative and empirical research on regional stabilization in fiscal unions. On the purely empirical side, Sala-i Martin and Sachs (1991), Feyrer and Sacerdote (2013)), and Bayoumi and Masson (1995) focus on reduced form estimates of properties of the tax-and-transfer system in fiscal unions, and in some cases, present back-of-the-envelope counterfactuals. Asdrubali et al. (1996) quantify the amount of risk sharing among states in the United States by decomposing cross-sectional variance in output. Atkeson and Bayoumi (1993) examine evidence for private insurance of regional risk in the United States and Europe.

Regarding quantitative research, Evers (2015) is closest to the application in this paper. He provides a quantitative analysis of federal fiscal rules in monetary unions by using a fully structural DSGE model, which builds on Chari and Kehoe (2007) and Kollmann (2001). The model has two countries, nominal wage, and price rigidities, and incomplete markets. It includes productivity and government spending shocks alone. He calibrates the model and finds very small differences in nearly every relevant business cycle statistic when comparing a monetary union with and without a central fiscal authority. I hypothesize that there are primarily two reasons he does not find quantitatively significant benefits from fiscal integration, while the opposite is true for this paper. First, his model does not include shifters in the Euler equation (e.g., discount rate shocks or borrowing constraints). In contrast, I find that most of the regional employment stabilization because of fiscal integration is because the tax-and-transfer system stabilizes “discount rate” shocks, and not “productivity” shocks. Second, his is a calibration exercise in a fully specified model, and so it fails to account for

several statistical properties of the joint distribution of the variables in the data (e.g., the autocorrelation of consumption or cross correlation between employment and output). Oh and Reis (2012) and McKay and Reis (2013) quantify the impact of targeted transfers and automatic stabilizers in an incomplete markets model, with heterogeneous agents and nominal rigidities. This paper shares their interest in the study of real-world tax-and-transfer systems and their macroeconomic consequences. However, they are largely concerned with implications for aggregate business-cycle stabilization, whereas I am concerned with implications for regional stabilization.

Moreover, the paper contributes to the recent surge in papers that exploit regional variation to highlight mechanisms of importance for economic fluctuations. Nakamura and Steinsson (2014) use sub-national US variation to inform the size of local government spending multipliers. Hurst et al. (2015) study regional redistribution through the US mortgage market. Mian and Sufi (2014), Mian et al. (2013), and Midrigan and Philippon (2011) exploit regional variation in the United States to explore the extent to which household leverage contributed to the Great Recession. Blanchard et al. (1992), Autor et al. (2013), and Charles et al. (2014) use regional variation to measure the responsiveness of labor markets to labor demand shocks.

Finally, the paper relates to econometrics literature on “partial identification” of structural models and how to perform statistical tests in these cases. Canay and Shaikh (2015) provides an extensive summary.

1.3 A SIMPLE FISCAL UNION MODEL

Consider an economy comprised of many islands, inhabited by a representative household and firm. The only other agent in the economy is a federal government. Households consume, work, and save/borrow in a non-state-contingent asset—a nominal bond in zero net supply. Firms produce final consumption goods using labor and intermediate goods. By assumption, the final consumption good is non-tradable, intermediate goods are tradable, and labor is not mobile across islands. Finally, each island has an exogenous endowment of intermediate goods. The federal government sets the nominal interest rate on the nominal bond, and gives lump-sum transfers to the islands. Assume that the nominal interest rate follows an endogenous rule that is a function of only aggregate variables (together with a fixed nominal exchange rate, this implies that the islands are part of a monetary union). Also, assume

that federal transfers are a function of island-level variables alone. Throughout, I assume that parameters governing preferences and production are identical across islands and the islands only differ, potentially, in the shocks that hit them—these shocks include a shifter of the households discount rate, a productivity shifter in the production function of final goods, and the exogenous endowment of tradable intermediate goods. Finally, I assume that all labor, goods and asset markets are competitive.

1.3.1 FIRMS AND HOUSEHOLDS

Final goods producers use labor N_{kt}^y and intermediates X_{kt} in island k at time t and face prices P_{kt} , wages W_{kt} , and intermediate prices Q_t (equalized across all islands because of assumed tradability). Their profits are

$$\max_{N_{kt}^y, X_{kt}} P_{kt} e^{z_{kt}} (N_{kt}^y)^\alpha (X_{kt})^\beta - W_{kt} N_{kt}^y - Q_t X_{kt}$$

where z_{kt} is a productivity shock and $(\alpha, \beta) : \alpha + \beta < 1$ are the labor and intermediates shares. Unlike the tradable goods prices, final good prices (P_{kt}) vary across islands.¹³

Households preferences are given by

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} e^{-\rho_{kt} - \delta_{kt}} \frac{(C_{kt} - \frac{\phi}{1+\phi} N_{kt}^{\frac{1+\phi}{\phi}})^{1-\sigma}}{1-\sigma} \right]$$

where C_{kt} is consumption of the final good, N_{kt} is labor, δ_{kt} is an exogenous processes driving the household's discount rate.

Households are able to spend their labor income $W_{kt} N_{kt}$ plus profits accruing from firms Π_{kt} and exogenous endowment of tradable goods $Q_t e^{\eta_{kt}}$, financial income $B_{kt} i_t$ and transfers from the government S_{kt} , where B_{kt} are nominal bond holdings at the beginning of the period and i_t is the nominal interest (equalized across islands given our assumption of a monetary union where the bonds are freely traded) on consumption goods (C_{kt}) and savings

¹³It is worth noting that all model shocks will generate endogenous variation in markups given assumed decreasing returns to scale. Additionally, what I call a “productivity shock” is isomorphic to any shifter of unit labor costs and, hence, labor demand schedules. I will not attempt to distinguish between the different interpretations of this shock in this paper.

$(B_{kt+1} - B_{kt})$. Thus, they face the period-by-period budget constraint

$$P_{kt}C_{kt} + B_{kt+1} \leq B_{kt}(1 + i_t) + W_{kt}N_{kt} + \Pi_{kt} + S_{kt} + Q_t e^{\eta_{kt}}$$

A well known issue in the international macroeconomics literature is that under market incompleteness of the type we just described there is no stationary distribution for bond holdings across islands in the log-linearized economy; and all other island variables in the model have unit roots. This is problematic for reasons both theoretical (I will like to study log-deviations from a deterministic steady state) and empirical (regional data for the US does not suggest the presence of such unit roots). I follow Schmitt-Grohe and Uribe (2003) and let ρ_{kt} be the endogenous component of the discount factor that satisfies $\rho_{kt+1} = \rho_{kt} + \Phi(\cdot)$ for some function $\Phi(\cdot)$ of the average per capita variables in an island. As such, agents do not internalize this dependence when making their choices. This modification induces stationarity for an appropriately chosen function $\Phi(\cdot)$. Schmitt-Grohe and Uribe (2003) show that alternative stationary inducing modifications (a specification with internalization, a debt-elastic interest rate or convex portfolio adjustment costs) all deliver similar quantitative results in the context of a small open economy real business cycle model.

1.3.2 FEDERAL GOVERNMENT

The federal government budget constraint is

$$B_t^g + \sum_k S_{kt} + Q_t G = B_{t-1}^g (1 + i_t)$$

where G is some exogenous level of government spending in intermediate goods. The key feature of a fiscally integrated economy is that the federal government has the ability to redistribute resources across islands via transfers S_{kt} . If the islands were fiscally independent such transfers would not be possible.

I assume that the federal government announces a nominal interest rate rule $i_t = i(\cdot)$ as a function of aggregate variables in the economy alone. Moreover, it announces a transfer policy rule as a function of per-capita employment, wages and assets in an island

$$S_{kt} = \bar{S}(\tilde{W}_{kt})^{\vartheta_w} (\tilde{N}_{kt})^{\vartheta_n} (\tilde{B}_{kt-1})^{\vartheta_b}$$

Again, agents do not internalize this dependence when making their choices.

1.3.3 EXOGENOUS SHOCKS AND PROCESSES

I assume the exogenous processes are AR(1) processes, with an identical autoregressive coefficient across islands, and that the innovations are iid, mean zero, random variables with an aggregate and island specific component. First, define $\gamma_{kt} \equiv \delta_{kt} - \delta_{kt-1}$. Then,

$$\begin{aligned} z_{kt} &= \rho_z z_{kt-1} + \tilde{\sigma}_z v_t^z + \sigma_z u_{kt}^z \\ \gamma_{kt} &= \rho_\gamma \gamma_{kt-1} + \tilde{\sigma}_\gamma v_t^\gamma + \sigma_\gamma u_{kt}^\gamma \\ \eta_{kt} &= \rho_\eta \eta_{kt-1} + \tilde{\sigma}_\eta v_t^\eta + \sigma_\eta u_{kt}^\eta \end{aligned}$$

with $\sum_k u_{kt}^z = \sum_k u_{kt}^\gamma = \sum_k u_{kt}^\eta = 0$. By assumption, I assume the average of the regional shocks sum to zero in all periods.

The “discount rate” process γ_{kt} is a shifter of a household’s discount rate, but it can be viewed as a proxy for the tightening of household borrowing limits. Such shocks have been proposed by Eggertsson and Krugman (2012), Guerrieri and Lorenzoni (2011), and Mian and Sufi (2014) as an explanation of the 2008 recession. Beraja et al. (2015b) find that these broad types of shocks explain most regional employment variation across states in the United States during the Great Recession and its aftermath. The “productivity” process z_{kt} can be interpreted as actual productivity, or a shifter of firm’s demand for labor or firm’s mark-ups. As an example, credit supply shocks to firms such as those proposed by Gilchrist et al. (2014) fall in this category. Finally, “wealth” process η_{kt} is modeled as an endowment of intermediate goods but can be interpreted as shifters of the budget constraint that agents face such as exogenous changes in household wealth.

1.3.4 EQUILIBRIUM

An equilibrium is a collection of prices $\{P_{kt}, W_{kt}, Q_t\}$ and quantities $\{C_{kt}, N_{kt}, B_{kt}, N_{kt}^y, X_{kt}\}$ for each island k and time t such that, for an interest rate rule $i_t = i(\cdot)$ and given exogenous processes $\{z_{kt}, \eta_{kt}, \gamma_{kt}\}$, they are consistent with household utility maximization and firm

profit maximization and such that the following market clearing conditions hold:

$$\begin{aligned}
C_{kt} &= e^{z_{kt}} (N_{kt}^y)^\alpha (X_{kt})^\beta \\
N_{kt} &= N_{kt}^y \\
G + \sum_k X_{kt} &= \sum_k e^{\eta_{kt}} \\
0 &= \sum_k B_{kt} + B_t^g
\end{aligned}$$

1.3.5 AGGREGATION

The first important assumption for aggregation is that all islands are identical with respect to their underlying production and utility parameters.¹⁴ The second assumption is that the joint distribution of island-specific shocks is such that its cross-sectional summation is zero. If K , the number of islands, is large this holds in the limit because of the law of large numbers. I log-linearize the model around this steady state and show that it aggregates up to a representative economy where all aggregate variables are independent of any cross-sectional considerations to a first order approximation.¹⁵ I denote with lowercase letters an island variable's log-deviation from the aggregate union equilibrium. Lowercase letters with a tilde denote deviations from the steady state. For example, $n_{kt} \equiv \tilde{n}_{kt} - \tilde{n}_t$ and $\tilde{n}_t \equiv \sum_k \frac{1}{K} \tilde{n}_{kt} = \sum_k \frac{1}{K} \log(N_{kt}/\bar{N})$. I assume that the monetary authority announces the nominal interest rate rule in log-linearized form: $\tilde{i}_{t+1} = \varphi_\pi \mathbb{E}_t[\tilde{\pi}_{t+1}]$ where $\tilde{\pi}_t$ is the aggregate inflation rate. Finally, I assume that the endogenous component of the discount factor is $\Phi(\cdot) = \Phi_0 n_{kt} + \Phi_1 n_{kt-1}$.¹⁶

The following lemma present the aggregation result and shows that we can write the

¹⁴Given that the broad industrial composition at the state level does not differ much across states, the assumption that productivity parameters are roughly similar across states is not dramatically at odds with the data.

¹⁵The model we presented has many islands subject to idiosyncratic shocks that cannot be fully hedged because asset markets are incomplete. By log-linearizing the equilibrium we gain in tractability, but ignore these considerations and the aggregate consequences of heterogeneity. As usual, the approximation will be a good one as long as the underlying volatility of the idiosyncratic shocks is not too large. If our unit of study was an individual, as for example in the precautionary savings literature with incomplete markets, the use of linear approximations would likely not be appropriate. However, since our unit of study is an island the size of a state I believe this is not too egregious of an assumption. The volatilities of key economic variables of interest at the state level are orders of magnitude smaller than the corresponding variables at the individual level.

¹⁶When $\Phi_0 > 0$ this will be enough to induce stationarity of island level variables in log-deviations from the aggregate. At the same time, since $\Phi(\cdot)$ depends only on these deviations, the aggregate equilibrium will feature a constant endogenous discount factor ρ .

island level equilibrium in deviations from these aggregates.

Lemma 1 *For given $\{z_{kt}, \gamma_{kt}, \eta_{kt}\}$, the behavior of $\{w_{kt}, n_{kt}, b_{kt}, p_{kt}, c_{kt}, x_{kt}\}$ in the log-linearized economy for each island in log-deviations from aggregates is identical to that of a small open economy where the price of intermediates and the nominal interest rate are at their steady state levels, i.e. $\tilde{q}_t = \tilde{i}_t = 0 \forall t$.*

Proof. See Appendix A.1 for a proof. ■

1.4 TRANSFERS POLICY IN ALTERNATIVE MODELS OF FISCAL UNIONS

The previous section described a model where many islands (I will call them states from now on) were fiscally integrated. I showed that to a first order approximation we can study each of these states separately as if they were small open economies that receive transfers from abroad. In this section, I ask: how would the equilibrium change if the transfer policy rule changed? In particular, the case when transfers are not a function of local variables $\vartheta_n = \vartheta_w = \vartheta_b = 0$ is of interest because it correspond to the case when the states are not fiscally integrated. By comparing the equilibrium of the small open economy with and without fiscal integration, it is possible to analyze the contribution of fiscal unions to regional stabilization. As a reminder, Lemma 1 implies that the following equations are sufficient to characterize the equilibrium dynamics of wages employment and assets $\{w_t, n_t, b_t\}$ of the states in log-deviations from the aggregate. For simplicity, I will drop the k notation indicating a given island from now on.

$$0 = \mathbb{E}_t \left[\frac{\sigma\beta}{\frac{\alpha\phi}{1+\phi} - 1} (w_{t+1} + n_{t+1} - w_t - n_t) + (\alpha + \beta - 1)(n_{t+1} - n_t) \right. \\ \left. + (\beta - 1)(w_{t+1} - w_t) + \left(1 + \frac{\sigma}{\frac{\alpha\phi}{1+\phi} - 1}\right)(z_{t+1} - z_t) + \Phi_0 n_{kt} + \Phi_1 n_{kt-1} - \gamma_{t+1} \right] \quad (\text{Euler})$$

$$0 = \left(\frac{1+\phi}{\phi} - (\alpha + \beta) \right) n_t - \beta w_t - z_t \quad (\text{Labor Market (1)})$$

$$0 = -\frac{B}{S} b_t + \frac{B}{S} (1+r) b_{t-1} - \frac{X}{S} (w_t + n_t) + \vartheta_n n_t + \vartheta_w w_t + \vartheta_b b_{t-1} + \frac{1}{S} \eta_t \quad (\text{SB})$$

Together with the exogenous processes

$$z_t = \rho_z z_{t-1} + \sigma_z v_t$$

$$\gamma_t = \rho_\gamma \gamma_{t-1} + \sigma_\gamma u_t^\gamma$$

$$\eta_t = \rho_\eta \eta_{t-1} + \sigma_\eta u_t^\eta$$

If parameters in this system of equations are calibrated or estimated, and the system is solved for the equilibrium, it is straightforward to construct alternative equilibria for different transfers rules simply by changing $\{\vartheta_n, \vartheta_w, \vartheta_b\}$. However, this would give an answer to the contribution of fiscal integration to regional stabilization for this particular model I described. One could imagine alternative models of fiscal unions that are reasonable a priori. For example, one could entertain the possibility that nominal wages are rigid and use a partial-adjustment model where a fraction λ of the gap between the actual and frictionless wage is closed every period. Formally, this implies a different (Labor Market) equation

$$\beta \frac{1-\lambda}{\lambda} (w_t - w_{t-1}) = \left(\frac{1+\phi}{\phi} - (\alpha + \beta) \right) n_t - \beta w_t - z_t \quad (\text{Labor Market (2)})$$

A similar specification has been used by Shimer (2010) and, more recently, by Beraja et al. (2015b) and Midrigan and Philippon (2011). Shimer (2010) argues that in labor market search models there is typically an interval of wages that both the workers are willing to accept and firms willing to pay. To resolve this wage indeterminacy he considers a wage setting rule that is a weighted average of a target wage and the past wage. The target wage in this case is the value of the marginal rate of substitution.

Alternatively, one could entertain the possibility that employment is durable as in Kydland and Prescott (1982). Formally, this implies a different (Labor Market) equation

$$\nu n_{t-1} = \left(\frac{1+\phi}{\phi} - (\alpha + \beta) \right) n_t - \beta w_t - z_t \quad (\text{Labor Market (3)})$$

where parameter ν governs how much one-period lagged employment choices affect current utility. Moreover, this implies a different Euler equation as well because the marginal utility

of consumption is affected.

$$0 = \mathbb{E}_t \left[\frac{\sigma\beta}{\frac{\alpha\phi}{1+\phi} - 1} (w_{t+1} + n_{t+1} - w_t - n_t) + (\alpha + \beta - 1)(n_{t+1} - n_t) + (\beta - 1)(w_{t+1} - w_t) \right. \\ \left. + \left(1 + \frac{\sigma}{\frac{\alpha\phi}{1+\phi} - 1}\right)(z_{t+1} - z_t) + \left(\Phi_0 - \frac{\sigma\alpha\frac{\phi}{1+\phi}}{(\alpha\frac{\phi}{1+\phi} - 1)}\right) n_{kt} + \left(\Phi_1 + \frac{\sigma\alpha\frac{\phi}{1+\phi}}{(\alpha\frac{\phi}{1+\phi} - 1)}\right) n_{kt-1} - \gamma_{t+1} \right] \\ \text{(Euler (3))}$$

Note that one could entertain other models with ever increasing complexity as well. To name a few, models with alternative utility functions (e.g., GHH preferences, CRRA or CARA preferences, with and without habits or durability in consumption and employment, etc.), adjustment costs for assets, debt-elastic interest rates, price and wage rigidities both forward- and backward- looking, other shocks (e.g., labor wedge shocks). The last two alternative models are particularly interesting because they both amplify the response of employment to shocks relative to response of wages. This allows to match the large employment volatility relative to wage volatility that we observe in actual economies. Hence, as a pure exercise in matching observed data generated by one of these models under a given transfer policy rule they could both do well. In this sense, some of their positive implications are observationally equivalent. I will exploit this lack of identification when constructing counterfactual equilibria for alternative transfer policy rules. I will argue that these and other models belong to a set that is indeed identified by their vector autoregression representation to the equilibrium, even if particular models are not, and that counterfactual equilibria are identical in all these models. The next section presents an example in this spirit. However, it is important to keep in mind that normative implications may be quite different, for example in models (2) and (3), because of the presence of nominal wage rigidities in the former. This cautions against the use of models in this set—that are observationally equivalent—for making normative statements, and it teaches us about when more data is needed to separately identify models in order to make such statements.

1.4.1 DIFFERENT MECHANISMS, SAME COUNTERFACTUAL...

The purpose of this section is to show an example of fiscal union models with very different economic mechanisms that: (i) are observationally equivalent given data on employment,

wages and assets alone, and (ii) produce identical counterfactual equilibrium for alternative transfer policy rule. This example is useful to gain some economic intuition before proceeding to the next section, which is more technical in nature and presents the main theoretical result for constructing semi-structural transfer counterfactuals.

I consider a model with both wage rigidities and durability in employment—combining models (2) and (3) from the previous section. For simplicity, I focus on the special case where the exogenous processes are not persistent ($\rho_\gamma = \rho_z = \rho_\eta = 0$). First, consider a benchmark parameterization $\xi^0 \equiv \{\lambda^0, \nu^0, \dots\}$ and transfer policy $\{\vartheta_n^0, \vartheta_w^0, \vartheta_b^0\}$. The equilibrium $\{n_t^0, w_t^0, b_t^0\}$ can be written in recursive form as,

$$\begin{bmatrix} n_t^0 \\ w_t^0 \\ b_t^0 \end{bmatrix} = \rho^0 \begin{bmatrix} n_{t-1}^0 \\ w_{t-1}^0 \\ b_{t-1}^0 \end{bmatrix} + \Lambda^0 \begin{bmatrix} u_t^\gamma \\ u_t^z \\ u_t^\eta \end{bmatrix}$$

where ρ^0, Λ^0 are functions of structural parameters and the transfer policy rule. Second, consider an alternative parameterization ξ^1 . Is it possible to find such parameterization that is consistent with $\{\rho^0, \Lambda^0\}$? The answer is yes, as the following claim shows, and it implies that the parameters are not identified given the recursive representation of the equilibrium. In other words, there are many models that are observationally equivalent.

Claim 1 *Given $\{\rho^0, \Lambda^0\}$, there is a parameterization ξ^1 such that $\rho^1 = \rho^0$ and $\Lambda^1 = \Lambda^0$. Moreover, $\{\lambda^1, \nu^1\}$ are such that $\frac{\lambda^1}{1-\lambda^1}\nu^1 = \frac{\lambda^0}{1-\lambda^0}\nu^0$ for any ξ^0 .*

Proof. See Appendix A.2. ■

The intuition for this observational equivalence result is the following. Higher employment durability increases the elasticity of employment to shocks relative to the response of wages. Analogously, higher wage rigidity decreases the elasticity of wages to shocks relative to the response of employment. Claim 1 formalizes this intuition by showing that models with high rigidity and low durability are observationally equivalent to models with low rigidity and high durability. However, the economic mechanisms that bring about variation in wages, employment and assets are very different.

Finally, note that the degrees of durability and rigidity only affect marginal decisions (i.e., the (Euler) and (Labor Market) equations) and not available resources (i.e., the (SB) equation). On the contrary, the transfer policy rule only affect available resources and not marginal decisions. This implies that models with different degrees of durability and rigid-

ity will have identical (Euler), (Labor Market) and (SB) equations for any policy and any parameterization consistent with Claim 1. Hence, given a benchmark policy and parameterization that produce the “observed” benchmark equilibrium, these observationally equivalent models will have identical counterfactual equilibrium for any alternative transfer policy rule.

1.5 SEMI-STRUCTURAL TRANSFER COUNTERFACTUAL IN FISCAL UNION MODELS

1.5.1 A SET OF LINEAR FISCAL UNION MODELS

I start by describing models of fiscal unions in broader terms than what I have been doing so far. I will focus on models that satisfy three properties. Examples of models with some of these properties include Farhi and Werning (2014), Nakamura and Steinsson (2014), Beraja et al. (2015b), and Evers (2015). In particular, the models in the previous sections satisfy them all.

Assumption 1 *Models of fiscal unions satisfy the following three properties:*

1. **Transfer policy rule:** *Tax-and-transfer system can be summarized as federal lump-sum transfers that are a function of state-level economic variables.*
2. **Linear aggregation:** *State-level economies in log-deviations from the aggregate union behave to a first-order approximation as if they were small open economies—independent of other states.*
3. **3 by 3:** *Employment n_t , nominal wages w_t , and assets b_t ; and exogenous processes $\{\gamma_t, \eta_t, z_t\}$ are sufficient variables for characterizing the state-level equilibrium in log-deviation from aggregates.*

Property 1 excludes models where the tax-and-transfers system affects decisions at the margin, as it would be the case with distortionary taxation. In Section 1.6.7, I relax this assumption and discuss robustness of results. Property 2 excludes from the analysis models where member states are inherently different because of industrial composition or household preferences, for example, and/or exogenous processes correlation structures across states are such that idiosyncratic shocks do not average out. In property 3, assets might encompass both non-state-contingent nominal bonds and certain types of tradable physical capital. What is important for the property to hold is that no other variables that are necessary to

describe the equilibrium are left out (e.g., other endogenous or exogenous state variables in the model).

Given the description of models of fiscal unions above, the system of matrix equations that characterizes the equilibrium in one island in log-deviations from the aggregate is written below. Without loss of generality, I will say that the first equation is the (Euler) equation in these models, the second is the (Labor Market) equation, and the third is the sequential budget constraint (SB) equation.

$$\begin{aligned}
0 &= (F + \Theta_f)\mathbb{E}_t \begin{bmatrix} n_{t+1} \\ w_{t+1} \\ b_{t+1} \end{bmatrix} + (G + \Theta_c) \begin{bmatrix} n_t \\ w_t \\ b_t \end{bmatrix} + (H + \Theta_p) \begin{bmatrix} n_{t-1} \\ w_{t-1} \\ b_{t-1} \end{bmatrix} \\
&+ L\mathbb{E}_t \begin{bmatrix} \gamma_{t+1} \\ z_{t+1} \\ \eta_{t+1} \end{bmatrix} + M \begin{bmatrix} \gamma_t \\ z_t \\ \eta_t \end{bmatrix} \\
0 &= - \begin{bmatrix} \gamma_t \\ z_t \\ \eta_t \end{bmatrix} + N \begin{bmatrix} \gamma_{t-1} \\ z_{t-1} \\ \eta_{t-1} \end{bmatrix} + \Sigma \begin{bmatrix} u_t^\gamma \\ u_t^z \\ u_t^\eta \end{bmatrix} \tag{FiscalSME}
\end{aligned}$$

I will say that a particular model of fiscal unions is characterized by the *structure* ξ and the *policy* Θ , where $\xi \equiv \{F, G, H, L, M, N, \Sigma\}$ are policy-invariant matrices, and $\Theta \equiv \{\Theta_f, \Theta_c, \Theta_p\}$ are matrices that characterize endogenous transfer policy rule. Θ_f contains the policy parameters associated with future expected variables, Θ_c with contemporaneous variables, and Θ_p past variables. The system evaluated at $\Theta = 0$ corresponds to a benchmark economy without a transfer policy in place. Elements in ξ involve combinations of subsets of structural parameters, e.g., preference and technology parameters. These combinations typically lack direct economic interpretation in terms of model primitives.

Given property 1, I will assume that the transfer policy is the same as in the models of fiscal unions from the previous section. Since the third equation is the sequential budget constraint, this implies:

Assumption 2 *The policy Θ is: $\Theta_c = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \vartheta_n & \vartheta_w & 0 \end{bmatrix}$, $\Theta_p = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \vartheta_b \end{bmatrix}$, and $\Theta_f = 0_{3,3}$*

1.5.2 EQUILIBRIUM REPRESENTATIONS

The description of models directly in terms of their equilibrium conditions is already a step forward toward construction of semi-structural transfers counterfactuals. The primitives or micro-foundations of such models need not be specified. Two models with different primitives but that have the same structure and policy in (FiscalSME) are equivalent for my purposes.

Assumption 3 ξ, Θ are such that the system (FiscalSME) is stabilizable.

Under Assumption 3, it is easy to derive, using the method of undetermined coefficients, a stable recursive solution to (FiscalSME) that can be written:

$$\begin{bmatrix} n_t \\ w_t \\ b_t \end{bmatrix} = P(\xi, \Theta) \begin{bmatrix} n_{t-1} \\ w_{t-1} \\ b_{t-1} \end{bmatrix} + Q(\xi, \Theta) \begin{bmatrix} \gamma_t \\ z_t \\ \eta_t \end{bmatrix} \quad (\text{FiscalRR})$$

Assumption 4 ξ, Θ are such that $Q(\xi, \Theta)$ is a non-singular square matrix.

Claim 2 If Assumptions 3 and 4 hold, there is a structural vector autoregression (SVAR) representation to the solution (RR) of the form:

$$\begin{bmatrix} n_t \\ w_t \\ b_t \end{bmatrix} = \rho_1(\xi, \Theta) \begin{bmatrix} n_{t-1} \\ w_{t-1} \\ b_{t-1} \end{bmatrix} + \rho_2(\xi, \Theta) \begin{bmatrix} n_{t-2} \\ w_{t-2} \\ b_{t-2} \end{bmatrix} + \Lambda(\xi, \Theta) \begin{bmatrix} u_t^\gamma \\ u_t^z \\ u_t^\eta \end{bmatrix} \quad (\text{FiscalSVAR})$$

where $\rho_1(\xi, \Theta) \equiv P(\xi, \Theta) + Q(\xi, \Theta)NQ(\xi, \Theta)^{-1}$; $\rho_2(\xi, \Theta) \equiv (P(\xi, \Theta) - \rho_1(\xi, \Theta))P(\xi, \Theta)$ and $V(\xi, \Theta) \equiv \text{Var}(\Lambda(\xi, \Theta)u_t) = \Lambda(\xi, \Theta)\Lambda(\xi, \Theta)'$, where $\Lambda(\xi, \Theta) \equiv Q(\xi, \Theta)\Sigma$.

Proof. We can write past exogenous processes as functions of past endogenous variables by inverting $Q(\xi, \Theta)$. Next, replace this and the law of motion for the exogenous states into (FiscalRR) to obtain (FiscalSVAR) representation. ■

1.5.3 SEMI-STRUCTURAL TRANSFER POLICY COUNTERFACTUAL

This section describes conditions under which it is possible to construct an equilibrium with $\vartheta_n = \vartheta_w = \vartheta_b = 0$ in a set of models of fiscal unions without the need to fully specify any particular model in this set. I begin by considering models of fiscal unions that can

be written as in (FiscalSME) and satisfy Assumptions 2 to 4. These models are completely described by structure ξ and policy Θ .

Next, I say that a set of the models is described by *semi-structure* $\xi^s \subseteq \xi$. In this application to fiscal unions, I will consider only semi-structures that take the form of exclusion restrictions, i.e., semi-structures that are described by certain elements in ξ being zeros. For example, one such exclusion restriction is that assets b_t do not enter in the (Labor Market) equation. Section 1.7 generalizes the semi-structures to linear restrictions on the elements of ξ .

Then, I ask: what are the semi-structures (i.e., models of fiscal unions) that are (i) consistent with a (FiscalSVAR) generated by a given $\{\xi, \Theta\}$, and (ii) deliver the same counterfactual (FiscalSVAR') if the policy changes from Θ to Θ' ? In particular, I am interested in $\Theta' = 0$ which corresponds to an economy without fiscal integration.

Proposition 1 *Given $\{\rho_1, \rho_2, \Lambda\}$ and Θ , if models of fiscal unions can be written as in (FiscalSME), satisfy Assumptions 2 to 4, and:*

1. *have identical (FiscalSVAR) representation $\{\rho_1, \rho_2, \Lambda\}$ when policy is Θ .*
2. *have identical semi-structure ξ^s , specified as at least 5 zeros in each equation in (FiscalSME).*

then counterfactual (FiscalSVAR') representation $\{\rho_1(\xi, \Theta'), \rho_2(\xi, \Theta'), \Lambda(\xi, \Theta')\}$ is identical for all these models. Moreover, this counterfactual can be constructed only with knowledge of $\{\rho_1, \rho_2, \Lambda, \Theta, \Theta'\}$.

Proof. See Appendix A.3. ■

The proposition identifies a set of models of fiscal unions that have identical counterfactual SVAR representation for alternative policies. Moreover, it shows that in order to construct a counterfactual economy without a transfers policy rule (i.e., $\Theta' = 0$) it is not necessary to fully specify and parameterize a particular model of fiscal unions that is consistent with (FiscalSME). Instead, suppose that both the structural vector autoregression representation of the equilibrium in a fiscal union, and the transfer policy Θ that generated it are known. Then, describing *sufficient semi-structure* ξ^s in the proposition is enough in order to construct a counterfactual equilibrium in the set of models that are consistent with this

semi-structure and are observationally equivalent in terms of their reduced-form equilibrium representation. I call this construction a *semi-structural transfer counterfactual*.¹⁷

The proposition is not obvious. This is because the structural vector autoregression representation is found as a solution to a highly non-linear system of equations given a structure and policy $\{\xi, \Theta\}$. Specifying a particular model, parameterizing it, and solving this system for the equilibrium under alternative policies is the standard way of evaluating them. However, fixing a SVAR representation, the mapping is linear back to the structure ξ that generated it. This is the key insight that allows identification of a set of models that yield identical counterfactuals by specifying only 5 zeros per equation in (FiscalSME). The reason why 5 zeros are sufficient is because there are 14 unknown elements per equation in structure ξ and there are 9 known elements per equation in structural vector autoregression representation (FiscalSVAR). Since the mapping is linear, specifying 5 of these unknown elements is enough to infer the remaining 9 unknown elements in structure ξ given the 9 known elements in (FiscalSVAR) and Θ . In Section 1.7, I show a general version of this proposition in linear dynamic stochastic models.

1.6 A COUNTERFACTUAL UNITED STATES ECONOMY WITHOUT FISCAL INTEGRATION

I start this section by describing all inputs to construct the semi-structural transfers counterfactual using Proposition 1. I need to: (i) specify a sufficient semi-structure ξ^s describing a set of fiscal union models, (ii) estimate a regional VAR using data for a fiscal union, (iii) identify the impulse response matrix, and (iv) estimate the transfer policy rule. Next, I present primary findings from the counterfactual exercise. I conclude the section by discussing the sensitivity of these findings to alternative ways of specifying the transfer policy rule.

1.6.1 THE SUFFICIENT SEMI-STRUCTURE ξ^s

The zeros in the matrices below describe a set of fiscal union models by specifying a semi-structure ξ^s that is sufficient in the sense of Proposition 1.

¹⁷The next section describes how to build all these inputs, in particular, how to identify the impulse response matrix in the SVAR using similar exclusion restrictions. From a methodological perspective, however, other SVAR identification schemes could be used as long as they are consistent the sufficient semi-structure.

Assumption 5 *A set of fiscal union models that can be written as in (FiscalSME) is described by:*

$$F = \begin{bmatrix} f_{11} & f_{12} & 0 \\ f_{21} & f_{22} & 0 \\ 0 & 0 & 0 \end{bmatrix}; G = \begin{bmatrix} g_{11} & g_{12} & 0 \\ g_{21} & g_{22} & 0 \\ g_{31} & g_{32} & g_{33} \end{bmatrix}; H = \begin{bmatrix} h_{11} & 0 & 0 \\ h_{21} & h_{22} & 0 \\ 0 & 0 & h_{33} \end{bmatrix};$$

$$L = \begin{bmatrix} 1 & l_{12} & l_{13} \\ 0 & 1 & l_{23} \\ 0 & 0 & 0 \end{bmatrix}; M = \begin{bmatrix} 0 & m_{12} & m_{13} \\ 0 & m_{22} & m_{23} \\ 0 & 0 & m_{33} \end{bmatrix}; N = \begin{bmatrix} n_{11} & n_{12} & n_{13} \\ 0 & n_{22} & n_{23} \\ 0 & n_{32} & n_{33} \end{bmatrix}$$

The absence of expected and lagged terms beside assets makes the third equation consistent with most log-linearized, incomplete market models that include a sequential budget constraint (SB). Moreover, I assume that the only exogenous shifter in the sequential budget constraint is “wealth” process η_t . The other two exogenous processes do not appear in the sequential budget constraint, which further restricts the set models that I analyze. Finally, it is assumed that the only exogenous shifters are “wealth” shocks, and the other two exogenous processes do not appear in the sequential budget constraint. In terms of the (Euler) and (Labor Market) equations—the first and second equations—I assume that (i) assets (future, contemporaneous, or lagged) $\{b_{t+1}, b_t, b_{t-1}\}$ do not appear in them, (ii) lagged wages w_{t-1} does not appear in the first equation, (iii) contemporaneous “discount rate” process γ_t do not shift these equations, and (iv) future “discount rate” process γ_{t+1} do not appear in the second equation. Finally, I assume that past “discount rate” process γ_{t-1} does not cause movements in “productivity” z_t and “wealth” η_t as is evidenced by autoregressive matrix N . These assumptions are consistent with all models of fiscal unions described in Section 1.4 as well as many others. The key feature of models described by this semi-structure is that the (Labor Market) and (Euler) equations dependence on future and lagged variables is relatively unconstrained, as well as the exogenous processes and their correlation structure. This is important for the question of regional stabilization in fiscal unions because it means that this set encompasses many models with rich features.

1.6.2 IDENTIFICATION OF $\Lambda(\xi, \Theta)$ AND SHOCKS

An important input in the construction of the semi-structural policy counterfactual is impulse response matrix $\Lambda(\xi, \Theta)$. The literature proposes myriad ways to identify it, ranging from

simple ordering assumptions to more sophisticated sign and long-run restrictions. I follow Lemma 3, which extends the methodology proposed in Beraja et al. (2015b). In essence, the procedure uses elements of theory to identify underlying shocks in a VAR. These theoretical restrictions imply a series of particular linear restrictions linking the reduced form errors to the structural shocks. I use the sequential budget constraint (SB) to generate these theoretical restrictions.¹⁸

The first step in the procedure consists of estimating the reduced form (FiscalSVAR) to obtain the autoregressive matrices $\{\rho_1, \rho_2\}$, and the reduced form errors covariance matrix V . Second, derive identification restrictions that will allow us to infer Λ and the shocks.¹⁹ Applying the conditional expectation operator $\mathbb{E}_{t-1}(\cdot)$ on both sides of the (SB) and constructing the reduced form expectational errors, we obtain:

$$0 = \begin{bmatrix} g_{31} + \vartheta_n & g_{32} + \vartheta_w & g_{33} \end{bmatrix} \Lambda \begin{bmatrix} u_t^\gamma \\ u_t^z \\ u_t^\eta \end{bmatrix} + m_{33}\sigma_{33}u_t^\eta \quad (\text{Id1})$$

This equation must hold for all realizations of the shocks. Whenever there is an innovation to u_t^γ or u_t^z and $u_t^\eta = 0$, employment, wages, and debt must co-move on impact in a way that satisfies this linear relationship. Hence, it gives us two linear restrictions in the second and third columns' elements of Λ for a given parameterization of (SB) when there are either contemporaneous u_t^γ or u_t^z shocks. These linear restrictions are the same as those implied by equation (6) in Section 1.7.1.

Similarly, constructing $\mathbb{E}_{t-1}(\cdot) - \mathbb{E}_{t-2}(\cdot)$, we obtain:

$$0 = \left(\begin{bmatrix} g_{31} + \vartheta_n & g_{32} + \vartheta_w & g_{33} \end{bmatrix} \rho_1 + \begin{bmatrix} 0 & 0 & h_{33} + \vartheta_b \end{bmatrix} \right) \Lambda \begin{bmatrix} u_{t-1}^\gamma \\ u_{t-1}^z \\ u_{t-1}^\eta \end{bmatrix} \\ + m_{33}\sigma_{33}n_{33}u_{t-1}^\eta + m_{33}\sigma_{33}n_{32}u_{t-1}^z \quad (\text{Id2})$$

This gives us one extra linear restriction in the second column's elements of Λ for a given parameterization of (SB) when there are u_{t-1}^γ shocks. These three linear restrictions,

¹⁸In Beraja et al. (2015b), we used a different equation to identify aggregate shocks—a wage-setting equation.

¹⁹The shocks are identified simply by multiplying the matrix of reduced form errors from the estimated VAR by the inverse of the matrix Λ .

combined with six non-linear restrictions coming from the orthogonalization of the shocks, are sufficient to identify all nine elements in Λ . Intuitively, equation (Id1) separates the “wealth” shock from the other two shocks. If the unexpected component of employment, wages, and assets does not co-move in the linear way implied by equation (Id1), when $u_t^\eta = 0$, a “wealth” shock must have occurred. Analogously, equation (Id2) separates “discount rate” and “productivity” shocks. If the unexpected component of employment, wages, and assets does not co-move in the linear way implied by equation (Id2), when $u_{t-1}^z = u_{t-1}^\eta = 0$, a “discount rate” shock occurred. For completeness, matrix Λ solves the system:

$$\begin{aligned} & \begin{bmatrix} G_{31} + \vartheta_n & G_{32} + \vartheta_w & G_{33} \end{bmatrix} \Lambda \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 \end{bmatrix} \\ & \left(\begin{bmatrix} G_{31} + \vartheta_n & G_{32} + \vartheta_w & G_{33} \end{bmatrix} \rho_1 + \begin{bmatrix} 0 & 0 & H_{33} + \vartheta_b \end{bmatrix} \right) \Lambda \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = 0 \\ & V = \Lambda \Lambda' \end{aligned}$$

This is the same system implied by Lemma 3 for this particular semi-structure.

1.6.3 DATA DESCRIPTION

I exclude Alaska, District of Columbia, and Hawaii from analysis, leaving 48 observations (one for each remaining state) per year, and 6 years (2006-2011) of data.

To make state-level nominal wages indices, I use data from the 2000 US Census and the 2001-2012 American Community Surveys (ACS).²⁰

The 2000 Census includes 5 percent of the US population. The 2001-2012 ACS’s include approximately 600,000 respondents between 2001-2004, and about 2 million after 2004. The large sample sizes allow detailed labor market information at the state level. I begin by using the data to make individual hourly nominal wages. I restrict the sample to only individuals who are employed, who report usually working at least 30 hours per week, and who worked at least 48 weeks during the prior 12 months. For each individual, I divide total labor income

²⁰I access the data through the IPUMS-USA website <https://usa.ipums.org/usa/>. See Ruggles et al. (1997).

earned during the prior 12 months by a measure of annual hours worked during the prior 12 months.²¹ The composition of workers differs across states and within a state over time, which might explain some variation in nominal wages across states over time. To account for this, I run the following regression:

$$\ln(w_{itk}) = K_t + \Gamma_t X_{itk} + u_{itk}$$

where $\ln(w_{itk})$ is log-nominal wages for household i in period t residing in state k , and X_{itk} is a vector of household specific controls. The vector of controls include a series of dummy variables for usual hours worked (30-39, 50-59, and 60+), a series of five-year age dummies (with 40-44 being the omitted group), 4 educational attainment dummies (with some college being the omitted group), three citizenship dummies (with native born being the omitted group), and a series of race dummies (with white being the omitted group). I run these regressions separately for each year such that both constant K_t and the vector of coefficients on the controls, Γ_t , can differ for each year. I then take the residuals from these regressions for each individual, u_{itk} , and add back constant K_t . Adding back the constant from the regression preserves differences over time in average log wages. To compute average log wages within a state, holding composition fixed, I average $u_{itk} + K_t$ across all individuals in state k . I refer to this measure as the demographically adjusted, log-nominal wage in time t in state k .

The measure of employment at the state level is the employment rate for each state, calculated using data from the US Bureau of Labor Statistics. The BLS reports annual employment counts and population numbers for each state and year. I divide employment counts by population to make an annual employment rate measure for each state.

Data on federal transfers net of taxes paid come from the Bureau of Economic Analysis.²² Transfers include retirement and disability insurance benefits, medical benefits, income maintenance benefits, unemployment insurance compensation, veterans benefits, federal education and training assistance, and other transfer receipts of individuals from governments. Federal taxes are the sum of personal income taxes that are withheld, usually by employers,

²¹Total labor income during the prior 12 months is the sum of both wage and salary earnings and business earnings. Total hours worked during the previous 12 months is a multiple of total weeks worked during the prior 12 months and the respondents' reports of their usual hours worked per week. For some years, bracketed reports are provided for weeks worked during the prior 12 months, and the usual hours per week worked. In those cases, I take the midpoint of the brackets.

²²I access the data through the BEA website on regional GDP and personal income: <http://www.bea.gov/iTable/index.regional.cfm>

from wages and salaries, quarterly payments of estimated taxes on income that is usually not subject to withholding, and final settlements, which are additional tax payments made when tax returns for a year are filed, or as a result of audits by the Federal Government.²³

Given the unavailability of official state-level data on asset positions, I construct a measure of state-level assets as the sum of physical and financial assets. From national account identities, we can derive the law of motion for assets B_t in a given state as:

$$B_t = B_{t-1}(1 + r_t) + Y_t - C_t + S_t - G_t^{local} + v_t$$

where Y_t is nominal gross domestic product, C_t is private consumption expenditures, S_t are net transfers (i.e., expenditures minus taxes) from the federal government, G_t^{local} are expenditures from the local government, and r_{t-1} captures the change in asset valuation between $t-1$ and t . Finally, error term v_t includes income receipts from abroad minus income payments to foreigners, federal government expenditures not counted as federal transfers (e.g., salaries and wages), and differences in returns between physical and financial assets for which no data are available.²⁴ I obtain Y_t and C_t directly from the Bureau of Economic Analysis website. $S_t - G_t^{local}$ also comes from several variables in the BEA. I calculate it as (personal current transfers receipts) - (personal current taxes paid + taxes on production and imports net of subsidies).²⁵ The revaluation of assets term r_t is obtained residually to ensure that the growth rate of the sum of local assets across states is consistent with the growth rate of aggregate net worth in the US economy. Having all components in the law of motion for B_t , I calculate assets at each point in time for each state simply by iterating forward with 2006 as the initial observation. I obtain initial assets in 2006 by aggregating at the state level, the zip code total net worth data from Mian et al. (2013). In order to construct financial assets at the zip code level Mian et al. (2013) they use data on dividends and interest income from the IRS Statistics of Income (SOI). They assume that households hold identical shares of stocks and bonds (they hold the market index portfolio). Given the share of total dividends and interest income received by a zip code they can construct the share of total stocks and bonds held by that zip code. Then, they total financial assets

²³Excise, Medicare and Social security federal taxes are not included in this measure.

²⁴Error term v_t accounts for most of the “wealth” exogenous process. The remainder is the error term e_t in the difference between observed net transfers and estimated policy rule in equation Section 1.6.4.

²⁵As long as local government expenditures plus transfers are close enough to local tax revenues (i.e., local governments have a nearly balanced budget), the calculation is accurate. If not, the difference is absorbed by error term e_t .

from the Federal Reserve’s Flow of Funds data to zip codes based on these shares. For the value of nominal debt owed by households they use data based on information from Equifax Predictive Services. Then they match the Federal Reserve Flow of Funds data by using the share of Equifax total debt in a zip code to allocate Flow of Funds debt. The final component of the asset measure is the value of housing wealth which they estimate using the 2000 Decennial Census data. They construct total home value as of 2000 in a zip code as the product of the number of home owners and the median home value. Then, they project it forward into later years using the CoreLogic zip code level house price index and an aggregate estimate of the change in homeownership and population growth.

1.6.4 ESTIMATION OF TRANSFER POLICY RULE Θ

Figure 1.1 shows a scatter plot of net federal transfers growth (direct federal transfers minus federal taxes growth) between 2006 and 2010 against nominal wage income growth (wage plus employment growth) between 2006 and 2010 in the United States.²⁶ There is a very strong, negative relationship between the two. In particular, a one percentage point increase in nominal wage income associates with a 1.25 percentage point decrease in net transfers. If the tax-and-transfer system helps stabilize regional economies, it is because a state whose economy temporarily worsens relative to the average receives some temporary relief through transfer payments or lower tax payments to the federal government—a type of insurance against temporary negative shocks. The opposite is true for states whose economies are in a relative boom. The negative relationship is a consequence of both the progressivity of the tax system in the United States and the presence of automatic stabilizers like unemployment insurance, and particularly during this period, federal emergency unemployment compensation and food stamps.²⁷

As a reminder, the transfer policy rule is:

$$s_t = \vartheta_n n_t + \vartheta_w w_t + \vartheta_b b_{t-1} + e_t$$

For regional data to be used to estimate Θ , one of the following must hold: (1) the innovations to the policy rule have no regional component ($e_t = 0$)—in which case, a simple OLS

²⁶See Section 1.6.3 for a detailed description of the construction of these data

²⁷If agents understand that federal emergency unemployment compensation would always be passed in situations of economic malaise, it is correct to include them as part of the implicit net federal transfer policy rule in place in a fiscal union.

Figure 1.1: Net Transfers growth v. Wage Income growth (2006-2010)



Note: Figure shows a simple scatterplot of the log-growth of net transfers in a state between 2006 and 2010 against nominal wage income log-growth during the same period. See text for details of variable construction. The population size in 2006 of the underlying state is represented by the size of the circle in the figure. The line represents a weighted regression line from the bivariate regression, where the weights are the population in 2006.

regression produces consistent estimates—or (2) valid instruments can be found that isolate movements in n_t, w_t, b_{t-1} that are orthogonal to e_t . The issue of endogeneity arises because of reverse causality. When the innovation in the policy rule is part of the “wealth” shock u_t^η , employment and wages both cause and are caused by net transfers in the equation above. To deal with the endogeneity of n_t, w_t, b_{t-1} , I proceed variously. First, I estimate a regression of net transfers onto nominal wage income alone (assuming $\theta_w = \theta_n$) using house price growth between 2006 and 2010 as an instrument. This accords with many recent papers, including Mian and Sufi (2014). Contemporaneous housing price growth strongly predicts contemporaneous nominal wage income growth. The instrument is valid as long as local housing prices are orthogonal to the transfer policy rule shock, which appears plausible. In the second approach, I use “discount rate” and “productivity” shocks in 2008, estimated from (FiscalSVAR), as instruments for wages and employment. They are linear combinations of wages, employment, and assets in 2008 that are orthogonal to the “wealth” shock,

and hence e_t , by construction.

Table 1.1: Policy rule baseline estimates

	ϑ_n	ϑ_w	ϑ_b	ϑ_{w+n}	R^2
OLS	-1.6** (0.5)	-0.9* (0.7)	-0.03 (0.02)	.	0.41
IV w/ shocks (1)	-1.3* (0.7)	-1.4* (0.8)	-0.02 (0.02)	.	0.41
IV w/ house prices (1)	.	.	-0.03 (0.02)	-1.1** (0.4)	0.42
IV w/ house prices and shocks (1)	-1.4* (0.6)	-1.2* (0.6)	-0.02 (0.02)	.	0.43
(2)	-1.3* (0.7)	-1.4* (0.7)	0.01 (0.08)	.	0.38

Note: Numbers in parenthesis are OLS (or second stage) standard errors. Variables with ‘**’ are significant at a 5% level. Variables with ‘***’ are significant at 1%. All variables are state log-growth rates between 2006 and 2010. b_{t-1} is exogenous in (1) and endogenous in (2)

Table 1.1 presents results for several specifications. The dependent variable is the log-growth rate of transfers minus the growth rate of taxes between 2006 and 2010 for each state. The independent variables are the log-growth rate of nominal wages between 2006 and 2010 and the log-growth rate of employment between 2006 and 2010 in the first two columns, and the log-growth rate of assets between 2006 and 2009 in the third column. In the fourth column, the independent variable is the sum of wage and employment growth. The first line is a simple OLS regression. The second presents two-stage, least-squares results using the “discount rate” and “productivity” shocks in 2008 $u_{2008}^\gamma, u_{2008}^z$. The third uses house price log-growth between 2006 and 2008 as an instrument instead. The fourth uses all three instruments. For all specifications, and when possible, I consider case (1) when b_{t-1} is not endogenous, and case (2) when b_{t-1} might be endogenous.


I find that the policy rule estimates have the expected sign and are significant in all specifications. They are also similar in magnitude, ranging from -1.3 to -1.6 for ϑ_n and -0.9 to -1.4 for ϑ_w . Lagged assets have nearly no independent explanatory power for net transfers across all specifications. To give a sense of the magnitudes involved, when net transfers increase by 30 percent for every 1 percent decrease in nominal wage income, and the average

income tax rate is 0.17, for every 1 dollar decrease in nominal wage income, a state receives 0.22 dollars in federal transfers. This result is similar to findings by Feyrer and Sacerdote (2013), who find a 0.25 decrease, and Bayoumi and Masson (1995), who find a 0.31 decrease.

1.6.5 MAIN FINDINGS

I estimate the vector autoregression (FiscalSVAR) via weighted OLS where the weights are the 2006 population in the state, using data described in subsection 1.6.3. For each variable and year, I take the cumulative log-growth between 2006 and 2011 and express it in log-deviations from the average across states. I pool all data between 2006 and 2011, leaving 240 observations (5 years*48 states), and estimate common autoregressive coefficients ρ_1, ρ_2 and reduced form errors U covariance matrix for all states $V = \frac{UU'}{240-3*2}$. I follow 1.6.2 to identify Λ , setting $g_{33} \equiv \frac{\bar{B}}{\bar{S}} = 2.25$ to match the median net worth to revenues ratio across states in the United States in 2006 and $\vartheta_n = -1.6, \vartheta_w = -0.9, \vartheta_b = -0.03$, which correspond to the OLS policy rule estimates in Table 1.1. Following Proposition 1, I construct a semi-structural transfers counterfactual for $\vartheta_n = \vartheta_w = \vartheta_b = 0$ which corresponds to an economy without a federal transfer policy rule.

Figure 1.2: Impulse Response Function to a Demand shock



Note: Figure shows the response of employment, wages and assets to a one-standard-deviation “discount rate” shock; both in the actual and counterfactual economy without transfers. The units are a state’s percentage deviation from the aggregate.

Figure 1.2 shows the impulse response functions of employment, wages, and assets to a one-standard-deviation “discount rate” shock γ , both in the actual and the counterfactual economy without transfers.²⁸ I find that employment and wages both decrease on impact, whereas assets increase in response to a “discount rate” shock. This accords with the theoretical impulse response functions in models in Section 1.4. As for the effects of fiscal integration, I find that amplification and persistence of “discount rate” shocks are mitigated by the transfer policy rule—the employment response (after two years) is -1.2 percent in the actual economy, whereas it is -2.1 percent in the counterfactual economy without transfers.

Table 1.2: Employment statistics: Fiscal Integration v. Fiscal Autonomy

		γ	(γ, z)	(γ, z, η)
σ_n^{2010}	$\Theta \neq 0$	2.3	2.5	2.6
	$\Theta = 0$	3	3.4	3.5
sk_n^{2010}	$\Theta \neq 0$	-0.26	-0.29	-0.2
	$\Theta = 0$	-0.39	-0.37	-0.3
$\bar{\sigma}_n$	$\Theta \neq 0$	2.3	2.7	3.5
	$\Theta = 0$	3.9	4.5	4.9
\overline{sk}_n	$\Theta \neq 0$	-0.03	-0.09	0.02
	$\Theta = 0$	-0.02	-0.07	-0.01
$\sqrt{s_n(0)}$	$\Theta \neq 0$	1.5	6.3	7.8
	$\Theta = 0$	5.5	9.7	10.7

Note: σ_n^{2010} is the standard deviation of the distribution of employment n_t across states in 2010 in percentages. sk_n^{2010} is the skewness of the distribution in 2010. $\bar{\sigma}_n$ and \overline{sk}_n are the standard deviation and skewness in the stationary distribution. $s_n(0)$ is the spectrum at zero frequency (the long-run variance of n_t). Line $\Theta \neq 0$ corresponds to the fiscal union economy, and line $\Theta = 0$ to results from the semi-structural counterfactual. Column γ corresponds to the case with only “discount rate” shocks, and column (γ, z) to the case with “discount rate” and “supply shocks. Column (γ, z, η) corresponds to the case with all shocks.


Table 1.2 presents moments of the employment distribution in the actual and counterfactual economies without transfers. The cross-state employment standard deviation in the US data in 2010, σ_n^{2010} , was 2.6 percent (this corresponds to the first line and third column in the table, when all shocks are present). I then consider the following thought experiment. At the end of 2007, it is announced that from 2008 onwards the United States federal government will cease to give transfers and collect taxes to/from the states in the union. How

²⁸Mian and Sufi (2014) and Beraja et al. (2015b), among others, argue that this type of shocks were key drivers of regional business cycles during the Great Recession in the United States.

would employment, wages, and assets have evolved were regional economies hit by the same sequence of shocks? I find that, absent a federal transfer policy rule, the standard deviation of employment in 2010 would have been 3.5 percent. The skewness of the distribution, sk_n^{2010} , would have been -0.3 instead of -0.2. To give some context to these numbers, aggregate output volatility during the pre-Volcker period (1960:1 to 1979:2) was 2.7, whereas during the post-Volcker-disinflation period (1982:4 to 1996:4) volatility was 2.06.²⁹ Much literature examines causes of this “Great Moderation”. The consequences of the US federal tax-and-transfer system are in the same order of magnitude.

Results above imply that the federal tax-and-transfer system helped stabilize regional economies by redistributing resources from regions that were doing relatively well to regions that were doing relatively poorly. Figure 1.3 elaborates on this point, showing the employment gain (or loss) from fiscal integration for each state in 2010, where states are sorted according to their employment in 2010 from lowest to highest. We observe that fiscal integration increased employment in states with the worst employment outcomes, and the opposite is true for states with the best employment outcomes.

Figure 1.3: Employment Gains from Fiscal Integration by State in 2010



Note: For each state, the figure shows the employment difference between the counterfactual economy without a federal transfer policy rule, constructed using the semi-structural methodology, and the actual economy in 2010. The states were sorted according to their actual employment in 2010 in ascending order. To the left, the states with the worst employment realizations; and, to the right, the states with best.

In the first and second columns of the table, I calculate the same statistics if regional

²⁹These numbers come from Clarida et al. (2000).

economies had been hit by only “discount rate” shocks or both “discount rate” and “productivity” shocks. Comparing the first and last columns in the first line, I find that most of the employment variation across states in 2010 is accounted for by “discount rate” shocks (approximately 90 percent). This accords with findings in Beraja et al. (2015b), in which data, and particularly the identification strategy, are different. Similarly, the federal transfer policy rule reduced employment dispersion primarily by stabilizing regional “discount rate” shocks. This is evidenced by comparing the first column in the second line of the table to the other columns. Of the 0.9 volatility reduction, 0.7 is achieved because of stabilization of “discount rate” shocks, and only 0.2 because of “productivity” shocks.³⁰

In the lower half of Table 1.2, I present Monte Carlo estimates of the standard deviation $\bar{\sigma}_n$ and skewness \bar{sk}_n of employment in the stationary distribution. I construct them by sampling with replacement 1,000,000 observations from the empirical distribution of shocks, feeding them to the (FiscalSVAR) and calculating the corresponding statistic. In the last line of the table, I present an estimate of the square root of the long-run variance of employment. The long-run variance is constructed as the diagonal element corresponding to employment in the spectrum at frequency zero of the multivariate (FiscalSVAR). Its square root is a common measure of the persistence in the series, as proposed by Pesaran et al. (1993). It associates with the speed at which the impulse response function to a given set of shocks decays, and hence their persistence. Results for the reduction in stationary employment volatility are qualitatively similar to the ones during the Great Recession. Finally, I find that the persistence of employment increases from 7.8 to 10.9 in the counterfactual economy without a federal transfer policy rule. Interestingly, most of the persistence in employment in the fiscally integrated economy is due to “productivity” shocks. In the counterfactual economy, “discount rate” shocks generate much larger persistence in employment.

1.6.6 CHANNELS OF REGIONAL STABILIZATION

To elucidate channels behind measured benefits from fiscal integration, I examine the contribution of what I define as a *rigidity channel*, which includes nominal and real rigidities in comparison to other broad frictions like asset market incompleteness and/or labor mobility costs. The idea is to construct an economy without nominal and/or real rigidities using the

³⁰The volatility reduction of 0.9 is the difference between the actual and counterfactual volatilities of 3.5 and 2.6. Analogously, 0.7 is the difference between 3 and 2.3, and 0.2 is the difference between 3.4 minus 2.5 and 0.7.

semi-structural methodology, and conduct the transfer policy counterfactual in this alternative economy. Then, follow a diff-in-diff approach to evaluate contributions of nominal/real rigidities to the benefits of fiscal integration by using differences across all four combinations; that is, an economy with rigidities and transfers (r, s) , rigidities and no transfers (r, ns) , no rigidities and transfers (nr, s) , and no rigidities and no transfers (nr, ns) . Any observation ω can be written as:

$$\omega = \beta_{nr,ns} + \beta_r \mathbb{I}_r + \beta_s \mathbb{I}_s + \beta_{r,s} \mathbb{I}_r \mathbb{I}_s$$

where \mathbb{I}_r is an indicator variable that is turned on when rigidities are present, and \mathbb{I}_s is an indicator variable turned on when transfers are present. For example, in the previous section, I identify $\beta_s + \beta_{r,s}$ for the employment distribution. This corresponds to the total difference between economies with and without transfers (the transfers counterfactual). The purpose of this section is to identify $\beta_{r,s}$, which measures the contribution of nominal/real rigidities to the total difference in the transfers counterfactual $\beta_s + \beta_{r,s}$ (i.e., the *rigidity channel*). Following diff-in-diff logic, we can identify $\beta_{r,s}$ by comparing the transfers counterfactual in economies with and without rigidities. That is:

$$\beta_{r,s} = (\omega_{r,s} - \omega_{r,ns}) - (\omega_{nr,s} - \omega_{nr,ns})$$

To construct counterfactual economies without nominal or real rigidities, I follow the semi-structural methodology in this paper. Given the semi-structure specified in Section 1.6.1, the second equation in (FiscalSME) is a combination of labor demand, labor supply, wage-setting, and/or price-setting schedules. This point was made in the aforementioned section. For the purpose of this counterfactual exercise, I further assume that parameters related to the degree of forward- and backward-lookingness in labor demand, price-setting, labor supply, or wage-setting schedules in the presence of nominal and real rigidities affect only the second equation in (FiscalSME). Then, the semi-structural methodology is also useful for defining a counterfactual economy in which there are no nominal and real rigidities. This counterfactual economy is consistent with models in which labor supply and demand are static. The way to construct such counterfactual economy is by imposing counterfactual policy matrices that cancel the forward- and backward-looking terms in the equation. Given the way we specified the semi-structure ξ^s , this is achieved by setting

$$\Theta'_{f,21} = -f_{21}; \Theta'_{f,22} = -f_{22}; \Theta'_{p,22} = -h_{22}; \Theta'_{p,21} = -h_{21}.$$

Table 1.3 presents results for the 20th, 50th, and 80th percentiles of the distribution of employment growth between 2008 and 2010. Figure A.1 presents kernel density estimates for all four distributions used in the construction of Table 1.3, and summary statistics.

Table 1.3: Employment in the Great Recession: Channels Decomposition

Percentiles	20 th	50 th	80 th
<i>Rigidity channel</i> (β_{rs})	0.4	-0.2	0
<i>Transfers counterfactual</i> ($\beta_s + \beta_{rs}$)	1.3	0	-1.1

Findings for the *transfers counterfactual* are as mentioned in the previous section. The increase in employment (1.3 percent) because of fiscal integration in the bottom part of the employment distribution (below the 20th percentile) is about the same as the decrease (-1.1 percent) in the upper part (above the 80th percentile). The degree to which nominal and real rigidities amplify (or attenuate) the consequences of fiscal integration attenuates, the *rigidity channel*, is not the same across the employment distribution. I find that about one-third of the 1.3 percent employment increase because of fiscal integration in states with the worst employment outcomes during the Great Recession was due to this channel. On the other hand, the channel was irrelevant for states with the best employment outcomes, meaning that at least for these states, the employment reduction because of fiscal integration involved frictions other than nominal and real rigidities.

1.6.7 ROBUSTNESS TO ALTERNATIVE POLICY SPECIFICATIONS

Results from the previous section are based on a particular transfer policy specification. As a reminder, the benchmark policy matrices were:

$$\Theta_f = 0; \Theta_c = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \vartheta_n & \vartheta_w & 0 \end{bmatrix}; \Theta_p = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \vartheta_b \end{bmatrix}$$

where $\vartheta_n = -1.6$, $\vartheta_w = -0.9$, $\vartheta_b = 0.03$, which corresponds to the OLS estimates of the transfer policy rule in Table 1.1.

First, we would like to evaluate the sensitivity of results to estimates of the transfer

policy rule other than the benchmark OLS estimates. Second, the specification is consistent with households in each state receiving lump-sum transfers from the federal government because the only affected equation that characterizes the equilibrium is the sequential budget constraint (the third equation in this case). In practice, lump-sum transfers are rare, and instead, the federal government uses distortionary taxes. This section explores the sensitivity of the counterfactual exercise to both of these issues.

Table A.1 shows the counterfactual employment standard deviation in an economy without transfers, both in 2010 and in the stationary distribution, for alternative initial policy parameterizations corresponding to the instrumental variable estimates in Table 1.1. Results are qualitatively similar to those reported in the previous section for the benchmark policy rule. Although quantitatively reduced somewhat for some of the parameterizations, gains from fiscal integration in terms of the reduction of the dispersion in employment across states remain large. The largest quantitative difference is for the case in which I restrict coefficients on employment and wages in the policy rule to be identical ($\vartheta_n = \vartheta_w = -1.1$). As the last column in the table shows, the counterfactual employment standard deviation in 2010 would have been 3.1 percent (instead of 2.6 percent in the data), and the counterfactual standard deviation in the stationary distribution is 4.5 percent (instead of 3.5). The counterfactual using the benchmark policy estimates instead resulted in counterfactual standard deviations of 3.5 and 4.9 percent, respectively.

Evaluating sensitivity to an alternative policy that accounts for the potentially distortionary effects of taxes is less straightforward. For simplicity, consider the transfer policy without lagged assets in log-deviations from the aggregate.³¹

$$s_t = \vartheta_n n_t + \vartheta_w w_t$$

This implies that tax rate τ_t per unit of nominal labor income $w_t + n_t$ in log-deviations from the aggregate can be written as:

$$\tau_t = -(1 + \vartheta_n)n_t - (1 + \vartheta_w)w_t$$

The potential labor supply (or wage-setting) tax distortion relates to τ_t , not total transfers s_t for which we estimated elasticities ϑ_w, ϑ_n . If the federal tax-and-transfer system affects

³¹Since the estimated coefficient on lagged assets is so small, results are identical whether we include it.

equilibrium equations beyond the sequential budget constraint, $(1+\theta_w)$, $(1+\theta_n)$ would appear in these equations, not θ_w, θ_n .

I consider the case in which the second equation in (FiscalSME) is a wage-setting equation, as in Section 1.4. If the federal tax-and-transfer system is distortionary, and given semi-structure ξ_s from Section 1.6.1, we can write it as:

$$0 = f_{21}\mathbb{E}_t[n_{t+1}] + f_{22}\mathbb{E}_t[w_{t+1}] + \left(g_{21} + \frac{\bar{\tau}\vartheta_n}{1 - \bar{\tau}\vartheta_n}(1 + \vartheta_n)\right)n_t + \left(g_{22} + \frac{\bar{\tau}\vartheta_n}{1 - \bar{\tau}\vartheta_n}(1 + \vartheta_w)\right)w_t \\ + h_{21}n_{t-1} + h_{22}w_{t-1} + \mathbb{E}_t[z_{t+1}] + l_{23}\mathbb{E}_t[\eta_{t+1}] + m_{22}z_t + m_{23}\eta_t$$

The equation above accords with the tax rate that affects the target wage in the wage-setting equation by distorting the marginal rate of substitution. The distortion is given by terms $1 + \theta_n$ and $1 + \theta_w$. For example, the case $\theta_w = \theta_n = -1$ is such that the tax schedule is flat (i.e., a proportional labor income tax). Due to the lack of curvature, it would not affect the island's log-deviations of the marginal rate of substitution from the aggregate.

I construct a semi-structural policy counterfactual using this alternative policy specification, in which I set $\bar{\tau} = 0.17$ to match the average tax rate in the US economy. Results from the previous section are unchanged because estimates of transfer policy rule ϑ_n, ϑ_w are close to -1 . Thus, policy-related terms that distort this equation are very small in absolute magnitude, in comparison with terms in the policy-invariant structure. To see this, consider the case in which the second equation is interpreted as a static labor supply equation, and the policy rule depends on employment alone: $w_t = \left(-g_{21} - \frac{\bar{\tau}\vartheta_n}{1 - \bar{\tau}\vartheta_n}(1 + \vartheta_n)\right)n_t$. Plausible calibrations of labor supply Frisch elasticity $-g_{21}$ are in the range 0.5 to 4.³² The policy-related term for the case when $\bar{\tau} = 0.17, \vartheta_n = -1.6$ is -0.08, which is an order of magnitude smaller than the Frisch elasticity.

1.7 SEMI-STRUCTURAL POLICY COUNTERFACTUALS: A GENERAL METHODOLOGY

This section describes the construction of semi-structural policy counterfactuals in a large class of linear models of dynamic stochastic economies. First, I define the class of models I consider, and develop notation and language that I use throughout the section. Second, I offer a formal definition for a *semi-structural policy counterfactual* in this class of models. I

³²See Chetty et al. (2011).

conclude by showing how to construct it and what is needed to do so depending on what we are willing to assume about the underlying structure of the economy.

The class of linear models I consider are such that equilibrium is characterized by the system of matrix equations,

$$\begin{aligned}
0 &= (F + \Theta_f)\mathbb{E}_t[x_{t+1}] + (G + \Theta_c)x_t + (H + \Theta_p)x_{t-1} + L\mathbb{E}_t[z_{t+1}] + Mz_t \\
0 &= -z_t + Nz_{t-1} + \Sigma u_t; \quad \text{iid } u_t \text{ with } \mathbb{E}[u_t] = 0, \text{Var}(u_t) = 1 \\
0 &= y_t + R_1x_t + R_2x_{t-1} + R_3z_t \\
&(H + \Theta_p)x_0; Mz_0 \text{ given}
\end{aligned} \tag{SME1}$$

where x_t is a column vector that includes all endogenous state variables, and could include endogenous control variables too; z_t is a column vector of exogenous state variables; and y_t a column vector of other endogenous variables not included in x_t . In the application of the methodology to fiscal unions x_t included log-deviations from the aggregate union of employment, wages and assets at the state-level, and z_t included the discount rate, productivity and wealth processes. Alternatively, in the context of Hansen (1985) real business cycles model, x_t includes log-deviations from the steady state of capital and potentially employment, z_t includes productivity, and y_t consumption.

For initial conditions, I say that a particular economy is characterized by the *structure* $\{\xi, \xi^a\}$ and the *policy* Θ , where $\xi \equiv \{F, G, H, L, M, N, \Sigma\}$, and $\xi^a \equiv \{R_1, R_2, R_3\}$ are policy-invariant matrices, and $\Theta \equiv \{\Theta_f, \Theta_c, \Theta_p\}$ are matrices that characterize endogenous policy rules. Θ_f contains the policy parameters associated with future expected variables, Θ_c with contemporaneous variables, and Θ_p past variables. The system (SME1) evaluated at $\Theta = 0$ corresponds to a benchmark economy without a policy in place.³³ Elements in $\{\xi, \xi^a\}$ typically involve combinations of subsets of parameters derived from a fully specified model. They typically lack direct economic interpretation in terms of the primitives of such a model. Elements in Θ generally have an economic interpretation. For example, the former could include non-linear combinations of labor supply and demand elasticities (as in the application to fiscal unions), and the latter, the transfers rule elasticities or the elasticity of a Taylor rule for nominal interest rates with respect to inflation in New Keynesian models. The description of models directly, in terms of its equilibrium conditions, is already a step forward toward

³³Uhlig (1995), from whom I borrow some notation, studies a very similar system of equations, and provides a computational toolkit for finding recursive solutions.

construction of semi-structural policy counterfactuals. The primitives or micro-foundations of such models need not be specified. Two models with different primitives but that have the same representation (SME1) are equivalent for the purposes of this section.

Assumption 6 ξ, Θ are such that the system (SME1) is stabilizable.

Under Assumption 6, it is easy to derive, using the method of undetermined coefficients, a stable recursive solution to (SME1) that can be written:

$$\begin{aligned} x_t &= P(\xi, \Theta)x_{t-1} + Q(\xi, \Theta)z_t \\ z_t &= Nz_{t-1} + \Sigma u_t \end{aligned} \tag{RR}$$

Assumption 7 ξ, Θ are such that $Q(\xi, \Theta)$ is a non-singular square matrix.

Claim 3 If Assumptions 6 and 7 hold, there is a structural vector autoregression (SVAR) representation to the solution (RR) of the form:

$$x_t = \rho_1(\xi, \Theta)x_{t-1} + \rho_2(\xi, \Theta)x_{t-2} + \Lambda(\xi, \Theta)u_t \tag{SVAR}$$

where $\rho_1(\xi, \Theta) \equiv P(\xi, \Theta) + Q(\xi, \Theta)NQ(\xi, \Theta)^{-1}$; $\rho_2(\xi, \Theta) \equiv (P(\xi, \Theta) - \rho_1(\xi, \Theta))P(\xi, \Theta)$ and $V(\xi, \Theta) \equiv \text{Var}(\Lambda(\xi, \Theta)u_t) = \Lambda(\xi, \Theta)\Lambda(\xi, \Theta)'$, where $\Lambda(\xi, \Theta) \equiv Q(\xi, \Theta)\Sigma$.

Proof. We can write $z_{t-1} = Q(\xi, \Theta)^{-1}(x_{t-1} - P(\xi, \Theta)x_{t-2})$ and replace it and the law of motion for the exogenous states into the law of motion for the endogenous variables to obtain the VAR(2) representation ■.

To summarize, linear models that can be written as in (SME1), where ξ, Θ satisfy Assumptions 6 and 7, formalize the general class of models on which semi-structural policy counterfactuals can be constructed.

The counterfactual answers how the structural vector autoregression representation of equilibrium (SVAR) changes when the policy changes from Θ to Θ' .

Definition 1 For $\xi^s \subseteq \xi$, a semi-structural policy counterfactual is a mapping Ω such that: $\Omega : \{\rho_1(\xi, \Theta), \rho_2(\xi, \Theta), \Lambda(\xi, \Theta), \Theta, \Theta', \xi^s\} \rightarrow \{\rho_1(\xi, \Theta'), \rho_2(\xi, \Theta'), \Lambda(\xi, \Theta')\}$

The mapping Ω requires knowledge of the matrices in (SVAR), the policy Θ that generated the (SVAR) and counterfactual policy Θ' . It also requires knowledge of a subset ξ^s

regarding the structure of economy ξ . I refer to this subset as a *semi-structure*. The definition nests the fully structural counterfactual when $\xi^s = \xi$, where one could directly solve the model (SME1) given the counterfactual policy. The fact that only a semi-structure ξ^s is required is what gives substance to the definition.

Assumption 8 *Given $\{\rho_1(\xi, \Theta), \rho_2(\xi, \Theta)\}$, the equation $\rho_2(\xi, \Theta) = (P - \rho_1(\xi, \Theta))P$ has a unique solution P with all eigenvalues inside the unit circle.*

Proposition 2 *If Assumption 8 holds, a semi-structural policy counterfactual Ω exists and is unique, and knowledge of $\{F, G\}$ is sufficient to construct Ω , i.e. $\xi^s \subset \{F, G\}$.*

Proof. See Appendix A.4.

The key to proving the proposition is to notice that $\{\rho_1(\xi, \Theta), \rho_2(\xi, \Theta), \Lambda(\xi, \Theta)\}$ are the solution to a linear system of matrix equations for all ξ, Θ , and that two such systems corresponding to Θ and Θ' are linear translations of each other. Since these systems are linear, we can construct the semi-structural policy counterfactual simply by subtracting one from the other. Knowledge of the matrices $\{F, G\}$ is sufficient because only these elements in ξ enter in a particular non-linear fashion in the systems (i.e., multiplying future and contemporaneous endogenous variable).

Construction of the semi-structural policy counterfactual in Proposition 2 might be difficult to implement in practice. In many applications, we might not want to, or be able to, specify $\{F, G\}$ entirely because, for example, the elements they contain might be the most controversial ones. In the the application to fiscal unions, for example, this includes parameters in the wage-setting equation governing how forward-looking wages are. Moreover, the construction requires knowledge of the impulse response matrix $\Lambda(\xi, \Theta)$. The next section demonstrates that both of these issues can be resolved by specifying an alternative, sufficient semi-structure ξ^s .

1.7.1 IMPLEMENTATION: HOW TO DESCRIBE THE SEMI-STRUCTURE ξ^s

This section discusses specification of the semi-structure ξ^s in a way that is useful to applications of the semi-structural methodology—one that allows identification of $\{F, G, \Lambda(\xi, \Theta)\}$ given $\{\rho_1(\xi, \Theta), \rho_2(\xi, \Theta), V(\xi, \Theta), \Theta\}$. I proceed in two steps. First, I describe identification of $\{F, G\}$ for a fixed matrix $\Lambda(\xi, \Theta)$. Second, I describe identification of $\Lambda(\xi, \Theta)$ by extending

the methodology in Beraja et al. (2015b) to a more general class of linear models considered in this paper.³⁴

I note that the matrices in the vector autoregression representation (SVAR) are found as a solution to a non-linear system of equations, given a structure and policy. The system of equations (1)-(4) is found in the proof of Proposition 2 in Appendix A.4.³⁵ This is the typical way to proceed when solving fully specified, explicit models of the economy. Notice, however, that we can think of this same system as being linear in the structural matrices, given a policy and the vector autoregression representation matrices. In general, the number of unknown elements in the full structure is larger than the number of equations in the system. Thus, the linear system is underdetermined, and it is not possible to infer the full structure with knowledge of the policy and vector autoregression representation matrices alone. The sufficient semi-structure is described by a minimum number of the elements in the full structure that, once specified, make the linear system in the remaining unspecified elements of the structure, exactly determined. I proceed to formalize this logic in what follows.

Given matrices $\{\rho_1(\xi, \Theta), \rho_2(\xi, \Theta), \Lambda(\xi, \Theta), \Theta\}$, the structure ξ satisfies the linear system of equations,

$$(FP + G + \Theta_c)P + H + \Theta_p = 0 \quad (5)$$

$$(LN + M)\Sigma + (F\rho_1 + G + \Theta_c)\Lambda = 0 \quad (6)$$

$$(LN + M)\Sigma\Sigma^{-1}N\Sigma + ((F\rho_1 + G + \Theta_c)\rho_1 + F\rho_2 + H + \Theta_p)\Lambda = 0 \quad (7)$$

$$(P - \rho_1)P - \rho_2 = 0 \quad (8)$$

This new system can be obtained from the system of equations (1)-(4). Per Assumption 7, all matrices in the system are $k \times k$ square matrices, where k is the number of variables in x_t . Suppose we are interested in identifying the first line in matrices $\{F, G\}$ (identification of other lines follows identical logic). We can count the number of equations involving the elements in the first line of $\{F, G\}$, and the number of unknown elements in the structure in this line. Then, we specify enough elements in the structure ξ such that the elements in the first line of $\{F, G\}$ are exactly determined (i.e., there are the same number of equations

³⁴Beraja et al. (2015b) discuss identification in a special case of this class of models.

³⁵See Uhlig (1995) and Blanchard and Kahn (1980) for a formal treatment.

and unspecified unknown elements). The following lemma formalizes this.

Lemma 2 *Given matrices $\{\rho_1(\xi, \Theta), \rho_2(\xi, \Theta), \Lambda(\xi, \Theta), \Theta\}$, the elements in line 'l' of the $k \times k$ matrices $\{F, G\}$ are uniquely identified if the following are specified:*

1. j_l elements in line 'l' of $(LN + M)\Sigma [I \quad \Sigma^{-1}N\Sigma]$.
2. n_l elements in line 'l' of $\{F, G, H\}$, where $n_l \equiv 2k - j_l - (1 - \text{rank}(\Theta_l))$.

Proof. See Appendix A.5. ■

The description above is not unique but useful since it separates required knowledge into two dimensions. First is knowledge about what exogenous variables enter in what equations, and their autoregressive matrix N . Second is knowledge about what endogenous variables enter in what equations, and how to parameterize certain elements in $\{F, G, H\}$. In the application to fiscal unions these were exclusion restrictions implying certain elements are zeros in the structure. In general, however, specifying a number of linear restrictions on the elements in the structure would suffice.

If $\Lambda(\xi, \Theta)$ were known, we could follow Lemma 2 for all lines 'l' from 1 to k . This would specify a sufficient semi-structure ξ^s to identify all lines of $\{F, G\}$. But what if $\Lambda(\xi, \Theta)$ is unknown? Identification of the impulse response matrix in vector autoregression models requires assumptions—ordering of shocks, sign restrictions, long-run restrictions, etc. These represent several routes that could be followed. Alternatively, the primary methodological insight in Beraja et al. (2015b) is that the impulse response matrix can be identified through specification of one or more structural equations, or in the terminology used in this paper, a semi-structure.

The following lemma describes this semi-structural identification scheme. It is an extension of the methodology in Beraja et al. (2015b) to the general class of linear models described in the previous section.

Lemma 3 *Take as given the matrices $\{\rho_1(\xi, \Theta), \rho_2(\xi, \Theta), V(\xi, \Theta)\}$ and all the elements in 's' lines of*

$\{F, G, H, \Theta\}$, where $s < k$. Let \mathbb{I}_s be a selection matrix for the 's' lines, and $A(s) \equiv \begin{bmatrix} \mathbb{I}_s(LN + M)\Sigma \\ \mathbb{I}_s(LN + M)N\Sigma \end{bmatrix}$ and $B(s) \equiv \begin{bmatrix} \mathbb{I}_s(F\rho_1 + G + \Theta_c) \\ \mathbb{I}_s((F\rho_1 + G + \Theta_c)\rho_1 + F\rho_2 + H + \Theta_p) \end{bmatrix}$. The impulse response matrix $\Lambda(\xi, \Theta)$ is

identified uniquely if the following are satisfied:

1. $B(s)$ has full rank.
2. There are $r(k) = k(k - 1)/2$ elements specified in $A(s)$.
3. There is at least one specified element in r_k in each of ' $k - 1$ ' columns of $A(s)$.
4. $\Lambda(\xi, \Theta)\Lambda(\xi, \Theta)' = V(\xi, \Theta)$.

Proof. See Appendix A.6. ■

Conditions 1 through 3 in the lemma imply a set of linear restrictions linking the reduced form error from (SVAR) to the structural shocks. They impose a linear co-movement that the unexpected components in the variables in (SVAR) (i.e., the columns in the impulse response matrix) must satisfy. These restrictions, combined with orthogonalization conditions in 4, are sufficient to identify all elements in the matrix $\Lambda(\xi, \Theta)$. The application to fiscal unions in this paper is a special case, with three endogenous variables and one specified structural equation (i.e., $k = 3$ and $s = 1$). Section 1.6.2 describes construction of the linear restrictions implied by Lemma 3 corresponding to this application, and provides intuition. Beraja et al. (2015b) is another example.

I conclude this section by noting that $r(k)$ elements in the ' s ' lines from Lemma 3, combined with $\{j_l, n_l\}$ elements in lines $l \notin s$ from Lemma 2, describe a sufficient semi-structure ξ^s for identifying $\{F, G, \Lambda(\xi, \Theta)\}$ uniquely.

1.8 FALSIFYING A SET OF MODELS VIA THE SEMI-STRUCTURAL METHODOLOGY

So far I have discussed how to construct a counterfactual economy after a policy change given observations of the economy before the policy change. There are many instances, however, where observations before and after the policy change are available. For instance, in the context of the application to stabilization in fiscal unions, we could entertain the possibility that in the future Catalunya decided to secede from the fiscal union that is Spain or that Europe would fiscally integrate. Or that only country-level monetary policy rules changed before and after the European Monetary Union, going from being country-specific to union-wide determined. In cases like this, I argue that we could compare the predicted VAR representation from a set of models and the actual VAR representation that we observe after the policy change. If these are “too far apart” in a statistical sense, we can

reject the null hypothesis that this set of models generated the observed data. Moreover, under normality of the shocks, the OLS estimates of the VAR are the maximum likelihood estimators. Hence, the statistical definition of “too far part” is given by a simple Wald test. Thus, the methodology can further be used to falsify a set of models in a way that is reminiscent to relatively new literature in econometrics on “partially identified” models.

alallalala

Formally, we wish to test the null hypothesis

$$\mathbf{H}_0 : h(\rho_1, \rho_2, V, \rho'_1, \rho'_2, V', \Theta, \Theta') \equiv \{\rho_1(\xi, \Theta'), \rho_2(\xi, \Theta'), V(\xi, \Theta')\} - \{\rho'_1, \rho'_2, V'\} = 0$$

where $\{\rho_1(\xi, \Theta'), \rho_2(\xi, \Theta'), V(\xi, \Theta')\}$ is the predicted VAR representation under the new policy Θ' by a set of models with structure ξ , and $\{\rho'_1, \rho'_2, V'\}$ is the observed VAR representation. The predicted representation is constructed via the semi-structural policy counterfactual mapping as a function of the observed VAR representation $\{\rho_1, \rho_2, V\}$ under the original policy Θ , the new policy Θ' and the semi-structure ξ^s . Hence, the test is about the null hypothesis that semi-structure ξ^s generated observed VAR representation after the change in policy.

Assumption 9 *The function $h(\rho_1, \rho_2, V, \rho'_1, \rho'_2, V', \Theta, \Theta')$:*

1. *is continuously differentiable on a neighborhood of $\{\tilde{\rho}_1, \tilde{\rho}_2, \tilde{V}, \tilde{\rho}'_1, \tilde{\rho}'_2, \tilde{V}', \tilde{\Theta}, \tilde{\Theta}'\}$.*
2. *has full rank first derivative matrix.*
3. *has positive definite second derivative matrix.*

Finally, given OLS consistent estimators $\{\tilde{\rho}_1, \tilde{\rho}_2, \tilde{V}, \tilde{\rho}'_1, \tilde{\rho}'_2, \tilde{V}', \tilde{\Theta}, \tilde{\Theta}'\}$, we can construct the Wald statistic that has asymptotic χ^2 distribution under the null hypothesis H_0 . The set of models described by semi-structure ξ^s is falsified if the null hypothesis is rejected.

1.9 CONCLUSIONS

In this paper, I propose a methodology to conduct quantitative counterfactual analysis in widely used models in macroeconomics with minimal *a-priori* structural assumptions. This semi-structural methodology allows the use of vector autoregression models to construct counterfactuals with respect to changes in policy rules. This analysis is robust across many

models while being immune to Lucas critique, in the spirit of Sims (1980). While I have emphasized benefits and feasibility of the semi-structural methodology, it also has some limitations when compared to fully structural or “sufficient statistics” approaches. The most important is that it is an inherently *linear* methodology. Counterfactual questions that involve non-linearities are not easily, or at all, accommodated.³⁶ A further limitation is that the methodology does not directly lend itself to welfare comparisons or optimal policy analysis.

I apply the methodology to quantify how fiscal unions, through their federal tax-and-transfer system, contribute to regional stabilization. This quantitative question has received surprisingly little attention in the literature beyond reduced-form calculations and simple calibration exercises in specific models, despite existing theoretical work on fiscal unions and its relevance for current discussions about European fiscal integration. My primary quantitative finding is that during the Great Recession fiscal integration significantly reduced cross-state employment differences by redistributing resources from states that were doing relatively well to states that were doing relatively poorly. On the one hand, this finding may overstate regional stabilization benefits because fiscal integration could partially displace existing private risk-sharing arrangements. On the other hand, measured gains from fiscal integrations would be larger if the reduction in state-level volatility reduces within-state individual risk exposure by more than it reduces per-capita (or average) risk exposure of the state’s “representative agent”.

In future research, the semi-structural methodology could help answering other questions in macroeconomics where stochasticity and dynamics are important and linear models provide a good approximation—for example, questions related to stabilization via monetary policy, financial integration, tax reform, exchange rate dynamics. Moreover, the methodology could be used in applications in other fields, like industrial organization and labor economics, where models typically involve simultaneous systems of equations of first order conditions that need to be estimated. Finally, I discussed how to use the methodology to falsify a set of models provided data before and after a policy change are available.

³⁶For example, in Beraja et al. (2015a), we study monetary policy in an environment where the pass-through from posted interest rates to effective interest rates is influenced by collateralized lending, a type of lending that naturally introduces non-linearities.

Chapter 2

THE AGGREGATE IMPLICATIONS OF REGIONAL BUSINESS CYCLES (WITH ERIK HURST AND JUAN OSPINA)

2.1 INTRODUCTION

A large and growing literature is exploiting regional variation to learn about the determinants of aggregate economic variables.¹ However, we argue that making inferences about the aggregate economy using only regional variation is complicated by two issues. First, we show that, in a monetary union model, local and aggregate elasticities to the same type of shock are quantitatively different both because of factor mobility and general equilibrium forces. This discrepancy makes it problematic to use local shock elasticities estimated from regional data to ascertain the importance of a given aggregate shock. Second, purely aggregate shocks get differenced out when using cross-region variation. As a result, it is not possible to learn anything about these aggregate shocks by exploiting variation across regions. Furthermore, we provide evidence of both these issues by contrasting the behavior of US aggregate time-series and cross-state patterns during the Great Recession. We document a strong relationship across states between local employment growth, and local nominal and real wage growth. These relationships are much weaker in US aggregates. In summary, we cannot expect to understand the joint evolution of aggregate variables by using cross-regional variation alone.

Therefore, we present a methodology that uses regional data along with aggregate data in order to identify aggregate shocks driving business cycles. The methodology exploits theoretical restrictions implied by a wage setting equation that hold in many monetary union

¹For recent examples, see Autor et al (2013), Charles et al (2015), Hagedorn et al (2015), Mehrotra and Sergeyev (2015), Mian and Sufi (2014) and Mondragon (2015).

models with wage stickiness. In turn, the extent to which aggregate wages are sticky is a key restriction in identifying the type of shocks driving aggregate fluctuations (e.g., “demand” vis a vis “supply” shocks)². Under certain conditions, we show how to use cross-region variation in wages, prices, and employment to estimate this wage setting equation—thus parameterizing the theoretical restrictions and linking regional business cycles to shock decompositions of aggregate business cycles.

Using household and retail scanner data for the US, we construct state-level wage and price indices as well as a measure of employment. Given the strong comovement of wages and employment across states, our estimates of the wage setting equation suggest that wages are relatively flexible—thus limiting the contribution of “demand” shocks to aggregate employment decline during the Great Recession. Instead, we find that a combination of “demand” and other shocks are necessary to account for the joint dynamics of aggregate prices, wages and employment during the 2007–2012 period. In particular, the relative stability of aggregate wages in the time-series compared to state-level wages is not caused by wage stickiness, but because different aggregate shocks have relatively offsetting effects on aggregate wages. We conclude that the wage stickiness necessary for demand shocks to be the primary cause of aggregate employment decline during the Great Recession is inconsistent with the flexibility of wages estimated from cross-state variation.

The paper is organized as follows. In Sections 2.2 and 2.3, we begin by documenting a series of new facts about the variation in nominal and real wages across US states during the Great Recession. Using data from the *2000 US Census* and the 2000 - 2012 *American Community Surveys* (ACS), we construct state-level nominal wage indices during the 2000 to 2012 period. We restrict our sample to full time workers with a strong attachment to the labor force. We adjust our wage measures to cleanse them from observable changes in labor force composition over the business cycle. In order to construct a measure of real wages we deflate our nominal wage indices with state-level price indices created using data from Nielsen’s *Retail Scanner Database*. The *Retail Scanner Database* includes weekly prices and quantities for given UPC codes at over 40,000 stores from 2006 through 2012. While the price indices we create from this data are based mostly on consumer packaged goods,

²We refer to a “demand” shock as a shock that moves employment and real wages in opposite directions and moves employment and prices in the same direction. In the model of the monetary union we develop below, these shocks can be formalized as shocks to the household’s discount rate or as shocks to the aggregate nominal interest rate rule. Our model also allows for a productivity/markup shock and a shock to household preference for leisure.

we show how under certain assumptions the indices can be scaled to be representative of a composite local consumption good. Furthermore, we show that an aggregate price index created with the retail scanner data matches the BLS's Food CPI nearly identically.

Using our indices, we show that states that experienced larger employment declines between 2007 and 2010 had significantly lower nominal and real wage growth during the same time period. These cross-state patterns stand in sharp contrast with the well documented aggregate time-series trends for prices and wages during the same time period. As both aggregate output and employment contracted sharply in the US during the 2007-2012 period, aggregate nominal wage growth remained robust and real wage growth did not break trend.³ In sum, while aggregate wages appear to be sticky during the Great Recession, state-level wages do not.

In Section 2.4, we present a monetary union model that we use for two purposes. First, a calibrated version of the model allows us to sign the elasticities to a given shock and quantify the differences between aggregate and local elasticities. Second, the model makes explicit assumptions that are sufficient to estimate the parameters in an aggregate wage setting equation using cross-state variation in employment, wages and prices. As we highlight below, these parameters help us identify the underlying aggregate drivers of the joint dynamics of employment, wages and prices.

The model has many islands linked by trade in intermediate goods which are used in the production of a non-tradable final consumption good. The only asset the economy is a one-period, non-state contingent nominal bond. The nominal interest rate on this asset follows a rule that endogenously responds to aggregate variables and is set at the union level. Labor is the only other input in production, which is not mobile across islands. We assume that nominal wages are only partially flexible. This is the only nominal rigidity in the model. Finally, the model includes multiple shocks: a shock to the household's discount rate, shocks to non-tradable and tradable productivity/markup, a shock to the household's preference for leisure, and a monetary policy shock. Aside from the monetary policy shock, all shocks have both local and aggregate components. By definition the weighted average of the local shocks sum to zero. We show that, under relatively few assumptions, the log-linearized economy aggregates. This allows us to study the aggregate and local behavior separately,

³The robust growth in nominal wages during the recession is viewed as a puzzle for those that believe that the lack of aggregate demand was the primary cause of the Great Recession. For example, this point was made by Krugman in a recent New York Times article ("Wages, Yellen and Intellectual Honesty", NYTimes 8/25/14).

a property that we will exploit when estimating the aggregate and regional shocks through our methodology.

Using a calibrated version of the model, we show that local employment elasticities to a local discount rate shock are two to three times larger than the aggregate employment elasticity to a similarly sized aggregate discount rate shock. This implies that elasticities often estimated for “demand” shocks (i.e., our discount rate shock) using cross-region variation are likely to dramatically overstate the elasticities of aggregate variables to “demand” shocks in the aggregate. The key general equilibrium forces in the model that may dampen aggregate elasticities are the endogenous response of nominal interest rates to aggregate variables and trade in the intermediate input. We show that local and aggregate elasticities get much closer together when the interest rate does not endogenously respond to changes in aggregate prices or employment (as when the economy is close to the zero lower bound).⁴

In Section 2.5, we turn to estimation of aggregate shocks. We present a procedure that allow us estimate the shocks in a larger class of monetary union models than the benchmark model outlined above, thus imposing less a-priori structure and making the analysis more persuasive. In particular, we consider models where the aggregate and local equilibria can be represented as a structural vector autoregression (SVAR) in price inflation, nominal wage inflation, and employment with three shocks. We refer to the three shocks as the discount rate shock (which is a combination of the discount rate and monetary policy shock), the productivity/markup shock (which is a combination of the productivity/markup shocks in the tradable and non tradable sectors) and the leisure shock (which is the shock to leisure preference). In order to identify the aggregate shocks, we estimate a SVAR and impose certain properties of our benchmark monetary union model. Our results will be consistent with monetary union models that satisfy all of these. First, we use the aggregate wage setting equation to derive a series of linear restrictions linking the reduced form errors to the underlying structural shocks. Second, we use the sign of the joint-response of employment, wages and prices (on impact) to a discount rate and a productivity/markup shock.⁵ These

⁴A similar point is made in Nakamura and Steinsson(2014) with respect to local estimates of fiscal multipliers.

⁵We view this methodology as an additional contribution of our paper. Beraja (2015) presents an extension of this scheme to a more general class of models. These are part of a growing literature developing “hybrid” methods that, for instance, constructs optimal combinations of econometric and theoretical models (Carriero and Giacomini (2011), Del Negro and Schorfheide (2004)) or uses the theoretical model to inform the econometric model’s parameter (An and Schorfheide (2007), Schorfheide(2000)). Our procedure is closest in spirit to the procedure recently developed in Baumeister and Hamilton (2015).

two, together with the usual shock-orthogonality conditions, are sufficient to identify the structural shocks.

The methodology requires parameterizing the structural wage setting equation. We use state-level data on prices, wages and employment during the 2006-2012 period to estimate the two parameters in our base specification, i.e., the Frisch elasticity of labor supply and the degree of wage stickiness. Across a variety of specifications and identification procedures, including instrumenting for local labor demand shocks, we estimate only a modest degree of wage stickiness. These estimates are much smaller than estimates of wage stickiness obtained using only aggregate time-series data.

With the parameterized aggregate wage setting equation, we use the SVAR identification procedure described above to estimate the shocks driving aggregate employment, prices, and wages during the Great Recession. Our results suggest that during the early part of the recession (2008-2009) roughly 30 percent of the aggregate employment decline can be attributed to the discount rate shock (i.e., the “demand” shock). The leisure shock explains roughly 30 percent of the decline in aggregate employment while the productivity/markup shock explains the remaining 40 percent. Over a longer period (2008-2012), however, the discount rate shock cannot explain any of the persistence in employment decline. Instead, it is the productivity/markup and labor supply shocks that explain why employment remained low from 2010-2012. In sum, while “demand” shocks may have been important in the early part of the recession, they cannot explain the persistently low levels of employment in the US after 2009.⁶ Furthermore, we find that the aggregate leisure shock - not sticky wages - explains why aggregate wages did not fall during the Great Recession.

Our paper contributes to many literatures. First, our work contributes to the recent surge in papers that have exploited regional variation to highlight mechanisms of importance to aggregate fluctuations. For example, Mian and Sufi (2011 and 2014), Mian, Rao, and Sufi (2013) and Midrigan and Philippon (2011) have exploited regional variation within the US to explore the extent to which household leverage has contributed to the Great Recession.⁷

⁶Christiano et al (2015a) estimate a New Keynesian model using data from the recent recession. Although their model and identification are different from ours, they also conclude that something akin to a supply shock is needed to explain the joint aggregate dynamics of prices and employment during the Great Recession. Likewise, Vavra (2014) and Berger and Vavra (2015) document that prices were very flexible during the Great Recession. They also conclude that something more than a demand shock is needed to explain aggregate employment dynamics given the missing aggregate disinflation.

⁷There has been an explosion of papers using regional data to better understand aggregate dynamics during the Great Recession. Some recent papers include: Giroud and Mueller (2015), Hagedorn et al. (2015), Mehrotra and Sergeyev (2015), and Mondragon (2015).

Nakamura and Steinsson (2014) use sub-national US variation to inform the size of local government spending multipliers. Blanchard and Katz (1991), Autor et al. (2013), and Charles et al. (2015) use regional variation to measure the responsiveness of labor markets to labor demand shocks. Our work contributes to this literature on two fronts. First, we show that local wages also respond to local changes in economic conditions at business cycle frequencies. Second, we provide a procedure where local variation can be combined with aggregate data to learn about the nature and importance of certain mechanisms for aggregate fluctuations. With respect to the latter innovation, our paper is similar in spirit to Nakamura and Steinsson (2014).

Second, our paper contributes to the recent literature trying to determine the causes of the Great Recession. In many respects, our model is more stylized than others in this literature in that we include a broad set of shocks without trying to uncover the underlying micro-foundations for these shocks. However, the shocks we chose to focus on were designed to proxy for many of the popular theories about the drivers of the Great Recession. For example, our discount rate shock can be thought of as reduced form representation of tightening of household borrowing limits. For example, such shocks have been proposed by Eggertsson and Krugman (2012), Guerrieri and Lorenzoni (2011) and Mian and Sufi (2014) as an explanation of the 2008 recession. Likewise, our productivity/markup shock can be interpreted as anything that changes firms' demand for labor. In a reduced form sense, credit supply shocks to firms, such as those proposed by Gilchrist et al (2014), would be similar to our productivity/markup shock. Finally, our leisure shock can be seen as a proxy for increased distortions in the labor market due to changes in government policy (e.g., Mulligan (2012) or as a reduced form representation of a skill mismatch story within the labor market (e.g., Charles et al. (2013, 2015)).

2.2 CREATING STATE-LEVEL PRICE AND WAGE INDICES

2.2.1 STATE-LEVEL WAGE INDEX

To construct nominal wage indices at the state level, we use data from the 2000 *Census* and the 2001-2012 *American Community Surveys* (ACS). The 2000 *Census* includes 5 percent of the US population while the 2001-2012 ACS's includes around 600,000 respondents per year between 2001 and 2004 and around 2 million respondents per year between 2005 and 2012. The large coverage allows us to compute detailed labor market statistics at the state level.

For each year of the *Census/ACS* data, we calculate hourly nominal wages for prime-age males with a strong attachment to the labor force. In particular, we restrict our sample to only males between the ages of 21 and 55, who were employed at the time of the Census, who reported usually working at least 30 hours per week, and who worked at least 48 weeks during the prior 12 months. Then, for each individual in the resulting sample, we divide total labor income earned during the prior 12 months by a measure of annual hours worked during prior 12 months.⁸

Despite our restriction to prime-age males with a strong attachment to the labor force, the composition of workers on other dimensions may still differ across states and within a state over time. The changing composition of workers could be explaining some of the variation in nominal wages across states over time. To cleanse our wage indices from these compositional issues, we create a composition adjusted wage measure (at least based on observables) by running the following regression on the ACS data:

$$\ln(w_{itk}) = \gamma_t + \Gamma_t X_{it} + \eta_{itk}$$

where $\ln(w_{itk})$ is log nominal wages for household i in period t residing in state k and X_{it} is a vector of household specific controls. The vector of controls include a series of dummy variables for usual hours worked (with “40-49 hours per week” being the omitted group), a series of five year age dummies (with “40-44” being the omitted group), four educational attainment dummies (with “some college” being the omitted group), three citizenship dummies (with “native born” being the omitted group), and a series of race dummies (with “white” being the omitted group). We run these regressions separately for each year so that both the constant, γ_t , and the vector of coefficients on the controls, Γ_t , can differ for each year. Then, we take the residuals from these regressions, η_{itk} , and add back the constant, γ_t . Adding back the constant from the regression preserves differences over time in average log-wages. To compute average wages in a state holding composition fixed, we average $e^{\eta_{itk} + \gamma_t}$ across all individuals in state k . We refer to this measure as the “adjusted nominal wage index” in time t in state k . This is the series we use to exploit cross-state variation in wages during the Great Recession.

The benefit of the *Census/ACS* data is that it is large enough to compute detailed labor

⁸Total labor income during the prior 12 months is the sum of both wage and salary earnings and business earnings. Total hours worked during the previous 12 month is the product of total weeks worked during the prior 12 months and the respondents report of their usual hours worked per week.

market statistics at state levels. However, one drawback of the *Census/ACS* data is that it not available at an annual frequency prior to 2000. To complement our analysis, we use data from the March Supplement of the *Current Population Survey (CPS)* to examine longer run aggregate trends in both nominal and real wages. These longer run trends are an input into our aggregate shock decomposition procedure discussed below. We compute the wage indices using the *CPS* data analogously to the way we computed the wage indices within the *Census/ACS* data.⁹ For the remainder of the paper, we use the *Census/ACS* data to explore regional wage variation and the *CPS* data to examine aggregate time series wage variation. However, for the 2000-2012 period, we can compare the time-series variation in aggregate wages using the *Census/ACS* data with the time series variation in aggregate wages using the *CPS* data. The two series have a correlation of 0.99 during this time period.

2.2.2 STATE-LEVEL PRICE INDEX

2.2.2.1 PRICE DATA

State-level price indices are necessary to measure state-level real wages. In order to construct state-level price indices we use the *Retail Scanner Database* collected by AC Nielsen and made available at The University of Chicago Booth School of Business.¹⁰ The Retail Scanner data consists of weekly pricing, volume, and store environment information generated by point-of-sale systems for about 90 participating retail chains across all US markets between January 2006 and December 2012. As a result, the database includes roughly 40,000 individual stores selling, for the most part, food, drugs and mass merchandise.

For each store, the database records the weekly quantities and the average transaction price for roughly 1.4 million distinct products. Each of these products is uniquely identified by a 12-digit number called Universal Product Code (UPC). To summarize, one entry in

⁹In particular, we compute hourly wages for men 21-55 with a strong attachment to the labor force (those currently working at least 30 hours a week and those who worked at least 48 weeks during the prior year). Again, like for the ACS data, we adjust the wages to account for a changing vector of observables over time. A full discussion of our methodology to compute composition adjusted wages in the *CPS* can be found in the Online Appendix that accompanies the paper.

¹⁰The data is made available through the Marketing Data Center at the University of Chicago Booth School of Business. Information on availability and access to the data can be found at <http://research.chicagobooth.edu/nielsen/>. Contemporaneously, Coibion et al. (2015), Kaplan and Menzio (2015) and Stroebel and Vavra (2014) also use local scanner data/household price data to estimate that local prices vary with local economic conditions at business cycle frequencies. Our paper complements this literature by actually making price indices using the Nielsen scanner data for each state at the monthly frequency and using those price indices to estimate structural parameters of the local wage setting equation.

the database contains the number of units sold of a given UPC and the weighted average price of the corresponding transactions, at a given store during a given week. The database only includes items with strictly positive sales in a store-week and excludes certain products such as random-weight meat, fruits, and vegetables since they do not have a UPC assigned. Nielsen sorts the different UPCs into over one thousand narrowly defined “categories”. For example, sugar can be of 5 categories: sugar granulated, sugar powdered, sugar remaining, sugar brown, and sugar substitutes. We use these categories when defining our price indices.

Finally, the geographic coverage of the database is outstanding and is one of its most attractive features. It includes stores from all states except for Alaska and Hawaii. Likewise, it covers stores from 361 Metropolitan Statistical Areas (MSA) and 2,500 counties. The data comes with both zip code and FIPS codes for the store’s county, MSA, and state. Over the seven year period, the data set includes total sales across all retail establishments worth over \$1.5 trillion. In this paper, we aggregate data to the level of US states and compute state-level retail scanner data price indices. Online Appendix Table R1 shows summary statistics for the retail scanner data for each year between 2006 and 2012 and for the sample as a whole.¹¹

2.2.2.2 A RETAIL SCANNER DATA PRICE INDEX

In order to construct state-level price indices we follow the BLS construction of the CPI as closely as possible.¹² While we briefly outline the price index construction in this sub-section, the full details of the procedure are discussed in the Online Appendix that accompanies our paper. Our retail scanner price indices are built in two stages. In the first stage, we aggregate the prices of goods within the roughly 1,000 categories described above. For our base indices, a good is a given store-UPC pair such that a UPC in store A is treated as a different good than the same UPC sold in store B. This allows for the possibility that prices may change as households substitute from a high cost store (that provides a different shopping experience) to a low cost store when local economic conditions deteriorate. Then, we compute, for

¹¹The Online Appendix is available at:
http://faculty.chicagobooth.edu/erik.hurst/research/regional_online_appendix.pdf

¹²There is a large literature discussing the construction of price indices. See, for example, Diewert (1976). Cage et al (2003) discuss the reasons behind the introduction of the BLS’s Chained Consumer Price Index. Melser (2011) discuss problems that arise with the construction of price indices with scanner data. In particular, if the quantity weights are updated too frequently the price index will exhibit “chain drift”. This concern motivated us to follow the BLS procedure and keep the quantity weights fixed for a year when computing the first stage of our indices rather than updating the quantities every month. Such problems are further discussed in Diewert et al. (2011).

each good, the average price and total quantity sold in a given month and state. Next, we construct the quantity weighted average price for all goods in each detailed category in a given month and state. We aggregate our index to the monthly level to reduce the number of missing values.¹³

Specifically, for each category, we compute:

$$P_{j,t,y,k} = P_{j,t-1,y,k} \times \frac{\sum_{i \in j} p_{i,t,k} \bar{q}_{i,t-1,k}}{\sum_{i \in j} p_{i,t-1,k} \bar{q}_{i,t-1,k}} \quad (2.1)$$

where $P_{j,t,y,k}$ is category level price index for category j , in period t , with a base year y , in state k . $p_{i,t,k}$ is the price at time t of the specific good i (from category j) in state k and $\bar{q}_{i,t-1,k}$ is the average monthly quantity sold of good i in the prior year in state k . By fixing quantities at their prior year's level, we are holding fixed household's consumption patterns as prices change. We update the basket of goods each year and produce the chained index for each category in each state.

In the second stage of our construction we aggregate the category-level price indices into an aggregate index for each state k . The inputs are the category-level prices and the total expenditures of each category. Specifically, for each state we compute:

$$\frac{P_{t,k}}{P_{t-1,k}} = \prod_{j=1}^N \left(\frac{P_{j,t,y,k}}{P_{j,t-1,y,k}} \right)^{\frac{\bar{S}_{j,k}^t + \bar{S}_{j,k}^{t-1}}{2}} \quad (2.2)$$

where $\bar{S}_{j,k}^t$ is the share of expenditure of category j in month t in state k averaged over the year.

Finally, as a consistency check, we compare our retail scanner price index for the aggregate US to the BLS's CPI for food. We choose the BLS Food CPI as a benchmark given that approximately 60 percent of the goods in our database can be classified as food.¹⁴ Figure B.1 shows that our retail scanner aggregate price index matches nearly exactly the BLS's Food

¹³One issue discussed in greater depth in the Online Appendix is how we deal with missing data when computing the price indices. Monthly prices may be missing, for instance, in the case of seasonal goods, the introduction of new goods, and the phasing out of existing goods. When computing our price indices, we restrict our sample to only include (1) goods that had positive sales in the prior year and (2) goods that had positive sales in every month of the current year. Online Appendix Table R1 shows the share of sales included in the price index for each sample year.

¹⁴The non-food goods in our sample include health and beauty products (13 percent), alcoholic beverages (6 percent), and paper products and household cleaning supplies (13 percent). The remaining items includes batteries, cutlery, pots and pans, candles, cameras, small consumer electronics, office supplies, and small household appliances.

CPI at the monthly level between 2006 and 2012.

2.2.2.3 A STATE-LEVEL PRICE INDEX FROM THE RETAIL SCANNER PRICE INDEX

The previous subsection described the construction of a state-level price index for goods sold in retail grocery and mass merchandising stores. However, our goal is to construct state-level price indices that are representative of the composite basket of consumer goods and services. In this subsection, we describe conditions under which our retail scanner price index and a composite local price index differ only by a scaling factor. We then propose to estimate this scaling factor using available data from the BLS. Nonetheless, as we highlight throughout, using this scaling factor (as opposed to using our retail scanner price indices directly) has little effect on the quantitative results of the paper.

Most goods in our sample are produced outside a local market and are simultaneously sold to many local markets. These intermediate production costs represent the traded portion of local retail prices. If there were no additional local distribution and/or trade costs, one would expect little variation in retail prices across states; the law of one price would hold. However, these “non-tradable” costs do exist, including the wages of workers in the retail establishments, the rent of the retail facility, and expenses associated with local warehousing and transportation.¹⁵

Assuming that the shares of these non-tradable costs are constant across states and identical for all firms in the retail industries, we can express local retail scanner prices, P^r , in region k during period t as:

$$P_{t,k}^r = (P_t^T)^{1-\kappa_r} (P_{t,k}^{NT})^{\kappa_r}$$

where P_t^T is the tradable component of local retail scanner prices in period t (which does not vary across states) and $P_{t,k}^{NT}$ is the non-tradable component of local retail prices in period t (which potentially does vary across states). κ_r represents the share of non-tradable costs in the total price for the retail scanner goods in our sample.

Analogously, we can express local prices in other sectors for which we do not have data as:

¹⁵Burstein et al (2003) document that distribution costs represent more than 40 percent of retail prices in the US.

$$P_{t,k}^{nr} = (P_t^T)^{1-\kappa_{nr}} (P_{t,k}^{NT})^{\kappa_{nr}}$$

where $P_{t,k}^{nr}$ is local prices in these sectors outside of the grocery/mass-merchandising sector and κ_{nr} is the share of non-tradable costs in the total price for these other sectors.¹⁶

Next, assume that the price of household's composite basket of goods and services in a state can be expressed as a composite of the prices in the retail scanner sectors ($P_{t,k}^r$) and prices in the other sectors (P_t^{nr}):

$$P_{t,k} = (P_t^{nr})^{1-s} (P_{t,k}^r)^s \equiv (P_t^T)^{1-\bar{\kappa}} (P_{t,k}^{NT})^{\bar{\kappa}}$$

where s is expenditure share of grocery/mass-merchandising goods in an individual's consumption bundle and $\bar{\kappa} \equiv (1-s)\kappa_{nr} + s\kappa_r$ is the non-tradable share in the aggregate consumption good, constant across all states.

Given these assumptions, we can transform the variation in retail scanner prices across states into variation in the broader consumption basket across states. Taking logs of the above equations and differencing across states we get that the variation in log-prices of the composite good between two states k and k' , $\Delta \ln P_{t,k,k'}$, is proportional to the variation in log-retail scanner prices across those same states, $\Delta \ln P_{t,k,k'}^r$. Formally,

$$\Delta \ln P_{t,k,k'} = \left(\frac{\bar{\kappa}}{\kappa_r} \right) \Delta \ln P_{t,k,k'}^r$$

If $\frac{\bar{\kappa}}{\kappa_r} > 1$, the local grocery/mass-merchandising sector will use a lower share of non-tradables in production than the composite local consumption good. In order to construct the scaling factor $\frac{\bar{\kappa}}{\kappa_r}$, it would be useful to have local indices for both grocery/mass-merchandising goods and for a composite local consumption good. While we do not have such indices for every US state, we can compare the relationship between local food inflation and local total inflation using BLS metro area price indices. These indices are only available for 27 MSAs at varying degrees of frequency (monthly, bi-monthly, semi-annually).¹⁷ As a result, they are not overly useful in measuring prices for a broad set of local areas. However, for the

¹⁶The grocery/mass-merchandising sector is only one sector within a household's local consumption bundle. For example, there are other sectors where the non-tradable share may differ from those in our retail-scanner data. For example, many local services primarily use local labor and local land in their production (e.g., dry-cleaners, hair salons, schools, and restaurants). Conversely, in other retail sectors, the traded component of costs could be large relative to the local factors used to sell the good (e.g., auto dealerships).

¹⁷In the online appendix that accompanies this paper, we discuss the BLS local price indices in greater depth.

MSAs covered, the BLS creates both a local food price index and a price index for the total local consumption basket. One approach to estimate $\frac{\bar{\kappa}}{\kappa_r}$, therefore, would be to estimate a regression of local food inflation on local total inflation using data for these 27 MSAs. However, the BLS cautions against such a regression because they report that the local price indices contain a substantial amount of measurement error.¹⁸ Such measurement error will bias our estimate of $\frac{\bar{\kappa}}{\kappa_r}$ towards zero.

To get around the measurement error problem, we follow the lead of Fitzgerald and Nicolini (2014) and regress food (total) inflation on some measure of local economic activity that is measured with relatively more precision. Taking a ratio of the coefficients from these two separate regressions can yield an estimate of $\frac{\bar{\kappa}}{\kappa_r}$. Specifically, we regress the 3-year inflation rate (either for food or total CPI) at the MSA level on the 3 year change in the unemployment rate during the 2007-2010 period. Within the BLS data, we find that a 1 percentage point increase in the local unemployment rate is associated with a 0.34 percentage point decline in the local food inflation rate (standard error = 0.22) and 0.47 percentage point decline the local composite inflation rate (standard error = 0.15). These estimates are very similar to those reported by Fitzgerald and Nicolini (2014) who use data over a longer time period. The fact that the coefficient on the change in unemployment rate is smaller in the food inflation regression than the total inflation regression is consistent with our belief that the tradable share of food is higher than the tradable share of the local composite consumption good. Given these coefficients, the BLS data suggests a measure of $\frac{\bar{\kappa}}{\kappa_r}$ of 1.4 (-0.47/-0.34). We will use this as our base adjustment factor throughout the paper. However, our main decompositions later in the paper are robust to any scaling factor between 1.0 and 2.0.


2.3 COMPARING CROSS-STATE PATTERNS TO AGGREGATE TIME-SERIES PATTERNS

The goal of this section is to contrast the strong co-movement of wages and economic activity at the local level to the relatively weaker co-movement at the aggregate level, during the Great Recession.

The left hand panel of Figure 2.1 shows the log-change in our demographic adjusted nominal wage indices between 2007 and 2010 across states against the log-change in the

¹⁸See, for example, <http://www.bls.gov/opub/btn/volume-1/pdf/consumer-price-index-data-quality-how-accurate-is-the-us-cpi.pdf>

Figure 2.1: State Employment Growth vs. State Nominal and Real Wage Growth, 2007-2010



Note: Figure shows a simple scatter plot of the percent growth in the state employment rate between 2007 and 2010 against nominal wage growth (left panel) and real wage growth (right panel) during the same period. The state employment rate comes dividing state employment from the BLS by total state population from the BLS. Nominal wages are computed from the ACS and are adjusted for the changing labor market composition of workers within each state over time. We restrict wage measures to a sample of men between the ages of 21 and 55 with a strong attachment to the labor market. Our composition adjustment controls for age, education, race, nativity and usual hours worked. See text for details. To compute real wages, we adjust our nominal wage measures by our local price indices created using the retail scanner data. The size of the underlying state is represented by the size of the circle in the figure. The line represents a weighted regression line from the bi-variate regression.

employment rate. As seen from the figure, nominal growth was strongly, positively correlated with employment growth in the 2007-2010 period. A simple linear regression through the data (weighted by the state’s 2006 labor force) suggests that a 1 percent change in a state’s employment rate is associated with a 0.62 percent change in nominal wages (standard error = 0.10). These findings are consistent with the extensive literature in labor economics and public finance showing that local labor demand shocks cause both employment and wages to vary together in the short to medium run. For example, Blanchard and Katz (1991), Autor, Dorn and Hanson (2013) and Charles, Hurst and Notowidigdo (2013) all find that negative local labor demand shocks cause substantial declines in local wages over the three to five year horizon. Our results further suggest that wages are fairly flexible in response to labor demand shocks at the local level. However, we illustrate the patterns at business cycle frequencies.¹⁹

¹⁹The patterns we document in Figure 2.1 also show up in other wage series. While there are no government data sets that produce broad based composition adjusted wage series at the local level, the Bureau of Labor Statistics’s Quarterly Census of Employment and Wages (QEW) collects firm level data on employment counts and total payroll at local levels. In Online Appendix Figures R1 and R2 we present results using local wage indices constructed from the QEW data instead. In these data, a one percent increase in a state’s

The right hand panel of Figure 2.1 shows similar patterns for real wage variation. We compute local real wages by deflating local nominal wage growth with the growth in the prices of a composite local consumption good ($P_{t,k}$).²⁰ A simple linear regression through the data (weighted by the state’s 2006 labor force) suggests that a 1 percent change in a state’s employment rate is associated with a 0.52 percent change in real wages (standard error = 0.15). Growth in local nominal and real wages were highly correlated with changes in many other measures of state economic activity during the 2007-2010 period as well. Although not shown, lower GDP growth, lower unemployment growth, lower hours growth and lower house price growth were all strongly correlated with lower nominal and real wage growth during the recent recession.

Figure B.2 shows our composition adjusted aggregate wage indices for the 2000 to 2012 period calculated using CPS data. To construct aggregate composition adjusted real wages, we deflate the aggregate nominal adjusted wages from the CPS by the aggregate June CPI-U with 2000 as the base year. Between 2007 and 2010, average composition adjusted nominal wages in the US increased by roughly 4 percent despite aggregate employment falling substantially. The patterns in our data replicate the aggregate nominal wage growth patterns documented by many others in the literature.²¹ Given that consumer prices increased by 5 percent during the same period, aggregate real wages in the US fell by roughly 1 percent between 2007 and 2010. This was similar to the trend in real wages prior to the start of the recent recession. As seen from Figure B.2 nominal wages increased slightly and real wage growth did not seem to break trend during the Great Recession. The “puzzle” is why aggregate wages did not decline relative to trend despite the very weak aggregate labor market. Wage stickiness is one potential explanation. However, as seen from Figure 2.1, local nominal and real wages moved quite a bit with changes in local employment during the same time period.

Table 2.1 compares these cross-state elasticities with the corresponding aggregate time-series elasticities during the Great Recession.²² The top panel displays the local wage elas-

employment growth between 2007 and 2010 was associated with a roughly 0.5 increase in the state’s nominal per capita earnings growth during the same time period.

²⁰As discussed in the previous section, we scale the growth in the retail scanner price index by a factor of 1.4 to account for the fact that grocery/mass merchandising goods have a higher tradable share than the composite local consumption good.

²¹See, for example, Daly and Hobijn (2015).

²²We thank Bob Hall for giving us the idea for this table. We base it on the analysis he did as part of his discussion of our paper at the 2015 NBER summer EFG program meeting.

Table 2.1: Comparison of Cross-State and Time-Series Estimates of Wage Elasticities During the Great Recession

	Employment Rate	Nominal Wage	Real Wage
Cross-State			
Cross-State Wage Elasticity With Respect to Employment, 2007-2010		0.62 (0.10)	0.52 (0.15)
Aggregate			
Actual Aggregate Growth, 2007-2010	-7.7 percent	3.8 percent	-0.9 percent
Expected Aggregate Growth, 2007-2010 (Based on 2000-2007 Trend)	-0.9 percent	5.5 percent	-2.1 percent
Aggregate Deviation from Expected Growth, 2007-2010	-6.8 percent	-1.7 percent	1.2 percent
Aggregate Wage Elasticity With Respect to Employment, 2007-2010		0.25	-0.17

Note: Table compares the wage elasticity to a one percent change in the employment rate estimated off of cross-state data (top panel) to a similarly defined wage elasticity estimated off of aggregate time-series data during the 2007 to 2010 period (bottom panel). The cross-state elasticities come from the simple scatter plots shown in Figure 3. Standard errors from the regression line in the scatter plots are shown in parentheses. The aggregate time-series elasticity is computed using aggregate data. For the aggregate nominal wage data, we use the adjusted wage series we created using data from the CPS. See text for details. For aggregate real wages, we adjust the nominal wage data by the June CPI-U. To get predicted nominal and real wage growth between 2007 and 2010, we take a simple linear prediction of the corresponding nominal and real growth between the 2000 and 2007 period. Once we get the deviation between actual wage growth and predicted wage growth between 2007 and 2010, we divide that difference by -6.8 percent. -6.8 percent is the decline in the aggregate employment rate between 2007 and 2010 above and beyond what would have been predicted from changes in the employment rate between 2000 and 2007. We use aggregate data from the BLS to compute the employment rate in 2000, 2007 and 2010.

ticities from the simple scatter plots shown in Figure 2.1. The bottom panel provides an estimate of similar elasticities over the same time period at the aggregate level. In particular, the last row shows the aggregate nominal (and real) wage elasticity with respect to changes in employment between 2007 and 2010. To construct these elasticities we use our adjusted nominal wage measure from the *CPS* (in the case of real wages, we deflate them with June CPI-U) and the aggregate employment-to-population ratio from the BLS. We de-trend all variables by estimating a linear trend between 2000 and 2007. The de-trended employment decline between 2007 and 2010 was 6.8 percent whereas the de-trended nominal wage decline was 1.7 percent. De-trended real wages actually increased by 1.2 percent during the 2007-2010 time period. Therefore, the implied aggregate wage elasticities with respect to

employment during the Great Recession are 0.25 for nominal wages (-1.7/-6.8) and -0.17 for real wages (1.2/-6.8).

Our main empirical finding comes from comparing the cross-state wage elasticities with the aggregate wage elasticities. The response of wages to changes in employment were *much* stronger at the state level during the Great Recession than at the aggregate level. For example, the local nominal wage elasticity with respect to employment changes was over twice as big as the aggregate elasticity (0.62 vs. 0.25). It is these differences in the relationships between wages and employment at the local level and at the aggregate level that forms the basis of the remainder of this paper. Why did local wages adjust so much when local employment conditions deteriorated during the Great Recession while aggregate wages hardly responded at all despite a sharp deterioration in aggregate employment conditions? Can aggregate wages be sticky when local wages adjust so much? We turn to answering these questions next.

2.4 A MONETARY UNION MODEL

In this section we present a monetary union model with several goals in mind. First, the model allows us to discuss the patterns we documented in the previous section in a formal environment where local economies aggregate. Second, the model makes explicit our assumptions on how wages are set. The nominal wage stickiness we specify will be essential to our identification strategy in later parts of the paper. Third, a calibrated version of the model allows us to quantify differences in aggregate vis a vis local elasticities to a variety of different shocks. While the theoretical possibility of these differences are known, much less is known about their magnitudes. The calibration exercise provides guidance to researchers who want to take an estimated local elasticity to a given shock and apply it to the aggregate economy. Fourth, the model provides an example of an economy that is encompassed by our SVAR procedure in Section 5 of the paper. The SVAR approach will allow us to estimate shocks for a larger set of models than the one we write down in this section. Finally, the model provides us with theoretical co-movements between variables that help us identify the shocks in the SVAR as well as give them an economic interpretation.

Formally, our model economy is composed of many islands inhabited by infinitely lived households and firms in two distinct sectors that produce a final consumption good and intermediates that go into its production. The only asset in the economy is a one-period

nominal bond in zero net supply where the nominal interest rate is set by a monetary authority. We assume intermediate goods can be traded across islands but the consumption good is non-tradable.²³ Finally, we assume labor is mobile across sectors but not across islands.²⁴ Throughout we assume that parameters governing preferences and production are identical across islands and that islands only differ, potentially, in the shocks that hit them.

2.4.1 FIRMS AND HOUSEHOLDS

Producers of tradable intermediate goods x in island k use local labor N_k^x and face nominal wages W_k (equalized across sectors) and prices Q (equalized across islands k). The time subscripts are omitted for clarity. Their profits are

$$\max_{N_k^x} Q e^{z_k^x} (N_k^x)^\theta - W_k N_k^x$$

where z_k^x is a tradable productivity shock in island k and $\theta < 1$ is the labor share in the production of tradables. Final (retail) goods y producers face prices P_k and obtain profits

$$\max_{N_k^y, X_k} P_k e^{z_k^y} (N_k^y)^\alpha (X_k)^\beta - W_k N_k^y - Q X_k$$

where z_k^y is a final good (retail) productivity shock and $(\alpha, \beta) : \alpha + \beta < 1$ are the labor and intermediates shares. Unlike the tradable goods prices, final good prices (P_k) vary across islands.²⁵

Households preferences are given by

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} e^{-\rho_{kt} - \delta_{kt}} \frac{(C_{kt} - e^{\epsilon_{kt}} \frac{\phi}{1+\phi} N_{kt}^{\frac{1+\phi}{\phi}})^{1-\sigma}}{1-\sigma} \right]$$

where C_{kt} is consumption of the final good, N_{kt} is labor, and δ_{kt} and ϵ_{kt} are exogenous processes driving the household's discount factor and the disutility of labor, respectively.

²³The final good can be thought of as being retail goods and services purchased in places such as: restaurants, barbershops and stores; and the intermediate sector providing physical goods such as: food ingredients, scissors and cellphones.

²⁴We explore the issue of labor mobility during the Great Recession when we take the model to the data.

²⁵It is worth noting that all model shocks will generate endogenous variation in markups given our assumption of decreasing returns to scale. Additionally, what we call a "productivity shock" is isomorphic to any shifter of unit labor costs and, hence, labor demand schedules. Later we will refer to it as the productivity/markup shock. We do not attempt to distinguish between the different interpretations of this shock in this paper.

Our base preferences abstract from income effects on labor supply. However, we show in section 7.4 that relaxing this assumption does not quantitatively change the conclusions of the paper.

Households are able to spend their labor income $W_{kt}N_{kt}$, profits accruing from firms Π_{kt} , financial income $B_{kt}i_t$, and transfers from the government T_t , where B_{kt} are nominal bond holdings at the beginning of the period and i_t is the nominal interest (equalized across islands given our assumption of a monetary union where the bonds are freely traded). Thus, they face the period-by-period budget constraint

$$P_{kt}C_{kt} + B_{kt+1} \leq B_{kt}(1 + i_t) + W_{kt}N_{kt} + \Pi_{kt} + T_{kt}$$

A well known issue in the international macroeconomics literature is that under market incompleteness of the type we just described there is no stationary distribution for bond holdings across islands in the log-linearized economy; and all other island variables in the model have unit roots. We follow Schmitt-Grohe and Uribe (2003) and let ρ_{kt} be the endogenous component of the discount factor that satisfies $\rho_{kt+1} = \rho_{kt} + \Phi(\cdot)$ for some function $\Phi(\cdot)$ of the average per capita variables in an island. As such, agents do not internalize this dependence when making their choices. This modification induces stationarity for an appropriately chosen function $\Phi(\cdot)$. Schmitt-Grohe and Uribe (2003) show that alternative stationary inducing modifications (a specification with internalization, a debt-elastic interest rate or convex portfolio adjustment costs) all deliver similar quantitative results in the context of a small open economy real business cycle model.

2.4.2 STICKY WAGES

We allow for the possibility that nominal wages are sticky and use a partial-adjustment model where a fraction λ of the gap between the actual and frictionless wage is closed every period. Formally:

$$W_{kt} = (P_{kt}e^{\epsilon_{kt}}(N_{kt})^{\frac{1}{\phi}})^{\lambda}(W_{kt-1})^{1-\lambda}$$

Given our assumption on household preferences, $P_{kt}e^{\epsilon_{kt}}(N_{kt})^{\frac{1}{\phi}}$ corresponds to the marginal rate of substitution between labor and consumption and the parameter λ measures the degree of nominal wage stickiness. In particular, when $\lambda = 1$ wages are fully flexible and

when $\lambda = 0$ they are fixed. This implies that workers will be off their labor supply curves whenever $\lambda < 1$. A similar specification has been used by Shimer (2010) and, more recently, by Midrigan and Philippon (2011). Shimer (2010) argues that in labor market search models there is typically an interval of wages that both the workers are willing to accept and firms willing to pay. To resolve this wage indeterminacy he considers a wage setting rule that is a weighted average of a target wage and the past wage. The target wage in our case is the value of the marginal rate of substitution.

Popular alternatives in the literature include the wage bargaining model in the spirit of Hall and Milgrom (2008) as in Christiano, Eichenbaum and Trabandt (2015b); and the monopsonistic competition model where unions representing workers set wages period by period as in Galí (2010). The key difference with the partial adjustment model is that both alternatives result in a forward looking component in the wage setting rule that is absent in our specification. In fact, this wage setting rule can be derived from the monopsonistic competition setup in the case where agents are myopic about the future; or from the labor market search setup in the special case where firms make take-it-or-leave-it offers and the probability of being employed in the future is independent of the current employment status.²⁶

2.4.3 EQUILIBRIUM

An equilibrium is a collection of prices $\{P_{kt}, W_{kt}, Q_t\}$ and quantities $\{C_{kt}, N_{kt}, B_{kt}, N_{kt}^x, N_{kt}^y, X_{kt}\}$ for each island k and time t such that, for an interest rate rule $i_t = i(\cdot)e^{\mu t}$ and given exogenous processes $\{z_{kt}^x, z_{kt}^y, \epsilon_{kt}, \delta_{kt}, \mu_t\}$, they are consistent with household utility maximization and firm profit maximization and such that the following market clearing conditions hold:

$$\begin{aligned} C_{kt} &= e^{z_{kt}^y} (N_{kt}^y)^\alpha X_{kt}^\beta \\ N_{kt} &= N_{kt}^y + N_{kt}^x \\ \sum_k X_{kt} &= \sum_k e^{z_{kt}^x} (N_{kt}^x)^\theta \end{aligned}$$

²⁶While there is no forward looking component in the reset wage in our base specification, we consider the implications of including forward looking behavior in Section 7.4.

2.4.4 SHOCKS

We assume exogenous processes are AR(1) processes, with an identical autoregressive coefficient across islands (and sectors in the case of productivity), and that the innovations (i.e., shocks) to these processes are iid, mean zero, random variables with an aggregate and island specific component. Let $\gamma_{kt} \equiv \delta_{kt} - \delta_{kt-1} - \mu_t$ be a combination of the discount rate exogenous growth and the monetary policy exogenous process that shows up as a wedge in the Euler equation. Then, we write the exogenous processes as:

$$\begin{aligned} z_{kt}^y &= \rho_z z_{kt-1}^y + \sigma_z u_t^y + \tilde{\sigma}_y v_{kt}^y \\ z_{kt}^x &= \rho_z z_{kt-1}^x + \sigma_z u_t^x + \tilde{\sigma}_x v_{kt}^x \\ \gamma_{kt} &= \rho_\gamma \gamma_{kt-1} + \sigma_\gamma u_t^\gamma + \tilde{\sigma}_\gamma v_{kt}^\gamma \\ \epsilon_{kt} &= \rho_\epsilon \epsilon_{kt-1} + \sigma_\epsilon u_t^\epsilon + \tilde{\sigma}_\epsilon v_{kt}^\epsilon \end{aligned}$$

with $\sum_k v_{kt}^y = \sum_k v_{kt}^x = \sum_k v_{kt}^\gamma = \sum_k v_{kt}^\epsilon = 0$. By assumption, we assume the weighted average of the island specific shocks sums to zero in all periods.

Let $u_t^z \equiv u_t^y + \beta u_t^x$ be a combination of productivity shocks in both sectors. We will call u_t^z , u_t^γ and u_t^ϵ the aggregate *Productivity/Markup*, *Discount rate* and *Leisure shocks* respectively. These are the shocks that the econometric procedure aims to identify. Analogously, v_{kt}^y , v_{kt}^x , v_{kt}^γ , v_{kt}^ϵ are the *Regional shocks*. The interpretation of the Leisure and Productivity/Markup shocks is relatively straightforward given our model environment. They are shifters of households and firms' labor supply (wage setting) and labor demand schedules, respectively. On the other hand, what we identify as a "discount rate shock" (γ_{kt}) is the combination of two more fundamental shocks. First, a shock to the marginal rate of substitution between consumption in consecutive periods. Second, a shock to the nominal interest rate rule set by the monetary authority. Our procedure is unable to distinguish between the two given that they both show up in the household's Euler equation, thus we treat them as a single shock.

2.4.5 AGGREGATION

Our first key assumption for aggregation is that all islands are identical with respect to their underlying production parameters (α , β , and θ), their underlying utility parameters (σ

and ϕ) and the degree of wage stickiness (λ).²⁷ Our second assumption is that islands are identical in the steady state and that price and wage inflation are zero. The last assumption is that the joint distribution of island-specific shocks is such that its cross-sectional sum is zero. If K , the number of islands, is large this holds in the limit because of the law of large numbers. We log-linearize the model around this steady state and show that it aggregates up to a representative economy where all aggregate variables are independent of any cross-sectional considerations to a first order approximation.²⁸ We denote with lowercase letters a variable's log-deviation from its steady state. Variables without a k subscript represent aggregates. For example, $n_{kt} \equiv \log\left(\frac{N_{kt}}{N}\right)$ and $n_t \equiv \sum_k \frac{1}{K} n_{kt}$. We assume that the monetary authority announces a nominal interest rate rule which is a function of aggregate variables. In log-linearized form, the rule is: $i_{t+1} = \varphi_\pi \mathbb{E}_t[\pi_{t+1}] + \varphi_y(y_t - y_t^*) + \mu_{t+1}$ where π_t is the aggregate inflation rate and $y_t - y_t^*$ is the output gap, defined as the difference between actual output and the flexible wage equilibrium output for the same realization of shocks. Finally, we assume that the endogenous component of the discount factor is $\Phi(\cdot) = \Phi_0(c_{kt} - c_t)$.²⁹

The following lemmas present a useful aggregation result and show that we can write the island-level equilibrium in log-deviation from the aggregate union equilibrium. Let w_t^r be real wage growth and π_t^w be nominal wage growth. Formally, $w_t^r \equiv \log\left(\frac{W_t/P_t}{W/P}\right)$ and $\pi_t^w = w_t^r - w_{t-1}^r + \pi_t$.

Lemma 4 *The behavior of π_t^w, w_t^r, n_t in the log-linearized economy is identical to that of a representative economy with only a final goods sector with labor share in production $\alpha + \theta\beta$, no endogenous discount factor, and only 3 exogenous processes $\{z_t, \epsilon_t, \gamma_t\}$.*

²⁷When implementing our procedure using data on US states, we discuss the plausibility of this assumption. Given that the broad industrial composition at the state level does not differ much across states, the assumption that productivity parameters and wage stickiness are roughly similar across states is not dramatically at odds with the data. As a robustness exercise, we estimate our key equations with industry fixed effects and show that our key cross section estimates are unchanged.

²⁸The model we presented has many islands subject to idiosyncratic shocks that cannot be fully hedged because asset markets are incomplete. By log-linearizing the equilibrium, we gain in tractability but ignore these considerations and the aggregate consequences of heterogeneity. The approximation will be good as long as the underlying volatility of the idiosyncratic shocks is not too large. If our unit of study was an individual, as for example in the precautionary savings literature with incomplete markets, the use of linear approximations would likely not be appropriate. However, since our unit of study is an island the size of a small country or a state, we believe this is not too egregious of an assumption. The volatilities of key economic variables of interest at the state or country level are orders of magnitude smaller than the corresponding variables at the individual level.

²⁹ $\Phi_0 > 0$ is enough to induce stationarity of island-level variables in log-deviations from the aggregate. Furthermore, since $\Phi(\cdot)$ depends only on these deviations, the aggregate equilibrium will feature a constant endogenous discount factor ρ .

Denote any variable $\tilde{x}_t \equiv x_{kt} - x_t$ as corresponding to island's k log-deviation from aggregates at time t , where the subscript k is dropped for notational simplicity.

Lemma 5 *For given $\{\tilde{z}_t^y, \tilde{z}_t^x, \tilde{\gamma}_t, \tilde{\epsilon}_t\}$, the behavior of $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t^y, \tilde{n}_t^x\}$ in the log-linearized economy for each island in deviations from aggregates is identical to that of a small open economy where the price of intermediates and the nominal interest rate are at their steady state levels, i.e. $q_t = i_t = 0 \forall t$.*

Proof. See Appendix B.1 for a proof of Lemma 1 and 2. ■

2.4.6 AGGREGATE VS. LOCAL SHOCK ELASTICITIES

Having described the model, we now explore the extent to which aggregate employment, price and wage elasticities to a given shock differ from local employment, price and wage elasticities to the same shock. Many researchers use clever identification strategies exploiting regional variation to estimate local elasticities to a given shock. For example, Mian and Sufi (2014) uses variation in debt across US metropolitan areas to isolate the extent to which a local “demand” shock (i.e., our discount rate shock) affects local employment. We show in this sub-section that the local employment (price, wage) elasticity to a given discount rate shock (productivity shock, leisure shock) is different, in general, from the aggregate elasticity to the same shock. Moreover, we calibrate the model in order to quantify the difference.

To gain some intuition as to the difference between local and aggregate elasticities in our model, we first consider the special case where there is an endowment of the tradable good and no labor is used in its production, i.e. $\theta = 0$. Focusing on a discount rate shock in this special case makes the comparison very transparent. We let $\xi_0^{agg} \equiv \frac{dn_0}{d\gamma_0}$ and $\xi_0^{reg} \equiv \frac{d\tilde{n}_0}{d\tilde{\gamma}_0}$ be the employment elasticities to the discount rate shock on impact. By solving for the recursive laws of motion in equilibrium we obtain,

$$\xi_0^{agg} = \frac{(1 - \lambda)}{(1 - \alpha + \frac{\lambda}{\phi})(\varphi_p - 1) + (\varphi_y \alpha - (\varphi_p - 1)(1 - \alpha)) \frac{1-\lambda}{\rho_\gamma}}$$

$$\xi_0^{reg} = \frac{d\tilde{n}_0}{d\tilde{\gamma}_0} = \frac{(1 - \lambda + \lambda\beta)}{\left(1 - \alpha + \frac{\lambda}{\phi} + \left(\frac{\sigma(1-\lambda(\alpha-\frac{1}{\phi}))}{1-\frac{\alpha\phi}{1+\phi}} - (1 + \frac{\lambda}{\phi})\right) \beta\right) \left(\frac{1+r}{\rho_\gamma} - 1\right)}$$

These expressions help understand the general equilibrium forces that make local and aggregate elasticities different. From the perspective of the closed economy, the endogenous

response of the nominal interest rate rule $\{\varphi_p$ and $\varphi_y\}$ reduces the aggregate employment impact elasticity to an unanticipated discount rate shock. A negative discount rate shock puts downward pressure on employment and prices. The monetary authority can lower interest rates to offset such a shock. The parameters of the interest rate rule are entirely absent in the expression for the regional elasticity. Therefore, the aggregate employment elasticity to a discount rate shock is typically smaller than the local employment elasticity to a local discount rate shock.

From the local perspective, since island level economies in deviations from the aggregate are small open economies, there are two extra margins of adjustment that are absent in the aggregate closed economy. First, the possibility to substitute labor for intermediate goods in the production of final consumption goods ($\beta > 0$) decreases the regional employment elasticity to the shock (as long as the term $\left(\frac{\sigma(1-\lambda(\alpha-\frac{1}{\phi}))}{1-\frac{\phi}{1+\phi}\alpha} - 1\right)$ is positive). Second, the possibility to transfer resources intertemporally through saving/borrowing at the interest rate r , as seen in the term $\left(\frac{1+r}{\rho_\gamma} - 1\right)$, decreases the regional employment elasticity. Theoretically, therefore, the aggregate employment elasticity to an aggregate discount rate shock can be either greater or smaller than the local employment elasticity to a local discount rate shock.

It is also interesting to compare how these discount rate elasticities change with the degree of nominal wage stickiness. Our identification procedure allows us to do this exercise when we estimate the impulse response to a discount rate shock. When $\varphi_p > 1$, both elasticities are decreasing in λ . In particular, employment does not respond to discount rate shocks at all in the limit when wages are perfectly flexible ($\lambda \rightarrow 1$).

While it is generally understood that local and aggregate elasticities can differ, there has been little quantitative work assessing the potential size of these differences. A parameterized version of our model allows us to directly compute the local and aggregate employment elasticities to different types of shocks. To this end, Table B.2 quantifies the employment impact elasticities to each of the shocks in the full model. Table B.1 presents and explains the parameterization of our model. Most of the parameters are standard from the literature or are chosen to match the labor share in the tradable and non-tradable sectors. The Online Appendix has an extended discussion of our baseline parameter choice. For our base specification, we use estimates of λ and ϕ of 2 and 0.7, respectively. These are the parameters that show up in the aggregate and local wage setting equations. The value of these parameters are the ones that we estimate using local variation in Section 6.

Column 1 of Table B.2 shows our base estimates of the local and aggregate employment elasticities. In columns 2 - 8 of Table B.2, we show how the elasticities change across alternate parameterization. Specifically, in column 2, we re-compute the elasticities reducing the Frisch elasticity of labor supply (ϕ) from 2 to 1. In column 3, we make wages more sticky by reducing λ from 0.7 to 0.5 (returning the Frisch elasticity to our base parameterization). In column 4, we set $\beta = 0$, thus shutting down the possibility to substitute labor for intermediate goods in the production of final goods. In the next two columns, we shut down the endogenous feedback in the nominal interest rate to changes in the employment gap such that φ_y is set to zero. In the first of those two columns, we leave the response of the nominal interest rate to the inflation target (φ_p) at its base parameterization. In the second of those two columns, we lower φ_p such that the local and aggregate responses to a discount rate shock are the same on impact. Finally, in the last two columns, we explore how the elasticities change as the persistence of the demand shock changes.

In our base specification, we find that the regional employment elasticity to a discount rate shock is 2.3 times larger than the aggregate employment elasticity to a discount rate shock. This implies that using cross-region variation to estimate local employment elasticities to demand shocks dramatically overstates employment responses when those local elasticities are applied to the aggregate. The conclusion remains unchanged across the different parameterizations of the wage setting rule, as shown in columns 2 and 3. Local employment elasticities to discount rate shocks are always two to three times larger than the aggregate employment elasticities. In columns 4 to 6, we see the importance of general equilibrium forces. As we shut down the ability to substitute labor for intermediate goods ($\beta = 0$), the gap between the regional and aggregate elasticities gets larger. The ability to trade intermediates across regions dampens the local employment elasticity to discount rate (demand) shocks. In columns 5 and 6, we see that the endogenous monetary policy response also dramatically dampens the aggregate response to a discount rate shock. This suggests that in periods where the economy is at the zero lower bound, aggregate and local employment elasticities to a demand shock are more similar, a point also made in Nakamura and Steinsson (2014). The last column explores the sensitivity to changes in the persistence of the discount rate shock. The less persistent is the local discount rate shock, the smaller the local employment elasticity because the regions can borrow from and lend to each other. Table B.2 also shows the local and aggregate employment response to local and aggregate productivity/markup and leisure shocks. For these two shocks, the local employment elas-

ticities are usually *smaller* than their aggregate counterparts. For the most part, this results from the particular specification of the nominal interest rate rule. To summarize, the quantitative difference between aggregate and local employment elasticities depends on the type underlying shock and can be quite large.

Tables B.3 and B.4 summarize the aggregate and regional impulse responses, respectively, for all variables and shocks in our benchmark calibration. We show results upon impact (the "short-run" elasticities) and after 5 years (the "long-run" elasticities). These tables allow us to assess the model's predictions. We use the same parameterization as in Table B.1. The short run responses in Columns 1 of Table B.3 and Table B.4 just restate the employment elasticities in column 1 of Table B.2. The remainder of the tables show the estimates for the price, nominal wage and real wage elasticities to all the underlying shocks in the model upon impact. As seen from Table B.3, an aggregate negative discount rate shock (households become less patient) lowers aggregate employment, lowers aggregate prices, and lowers (slightly) aggregate real wages. Conversely, an aggregate negative productivity shock lowers aggregate employment, raises aggregate prices, and raises aggregate real wages. We will use the sign of these impact elasticities to help identify the shocks in the SVAR in Section 2.5.

2.5 A PROCEDURE FOR IDENTIFYING AGGREGATE SHOCKS

In this section, we develop a procedure that allow us estimate the shocks in a larger class of monetary union models than the benchmark model outlined above, thus imposing less a-priori structure and making the analysis more persuasive.³⁰ Specifically, we consider models where aggregate equilibria can be represented as a structural vector autoregression (SVAR) in price inflation, nominal wage inflation, and employment with three shocks. In order to identify the shocks, we use three properties of our benchmark monetary union model: the wage setting equation, the sign of the impact elasticities to a discount rate and productiv-

³⁰We also performed a business cycle accounting exercise by solving the model and using the data to recover the exogenous stochastic processes. By doing so, we learned that different "wedges" are required to explain the joint dynamics of employment prices and wages. We find that the labor wedge is quite important in explaining variation in employment in the early stages of the Great Recession. Like our SVAR results, the Euler equation wedge only explained less than half of the employment decline during the early part of the recession and explained essentially none of the persistence. However, unlike the SVAR, the business cycle accounting does not allow us to recover the fundamental shocks and, therefore, we do not report more specific results here. As has been shown in Buera and Moll (2015), slight changes in model specification can alter the mapping between underlying structural shocks and corresponding aggregate wedges.

ity/markup shocks, and the orthogonality of shocks. Our results will be consistent with monetary union models that satisfy all of these. Beraja (2015) discusses this identification procedure in detail, as well as its application to more general SVARs and theoretical models than the ones in this paper.

We begin by noting that the recursive solution to the equilibrium system of equations in Lemma 1 can be written as a SVAR(∞) in $\{\pi_t, \pi_t^w, n_t\}$.³¹

$$(I - \rho(L)) \begin{bmatrix} \pi_t \\ \pi_t^w \\ n_t \end{bmatrix} = \Lambda \begin{bmatrix} u_t^\epsilon \\ u_t^z \\ u_t^\gamma \end{bmatrix}$$

Knowledge of $\rho(L)$ and an invertible matrix Λ together with aggregate data on prices, nominal wages and employment allow recovering the structural shocks.

The first step in our procedure consists of estimating the reduced form VAR to obtain the autoregressive matrix $\rho(L)$ and the reduced form errors covariance matrix V . In practice we will truncate $\rho(L)$ to be of finite order as it is typically done in the literature. The second step involves deriving a set of theoretical restrictions to identify the structural shocks from the reduced form errors.

As a reminder, the wage setting equation³² in log-linearized form is:

$$\pi_t^w = \lambda(\pi_t + \epsilon_t - \epsilon_{t-1}) + \frac{1}{\phi}(n_t - n_{t-1}) + (1 - \lambda)\pi_{t-1}^w$$

Applying the conditional expectation operator $\mathbb{E}_{t-1}(\cdot)$ on both sides and constructing expectational errors, we obtain:

$$\begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \Lambda \begin{bmatrix} u_t^\epsilon \\ u_t^z \\ u_t^\gamma \end{bmatrix} + \lambda \sigma_\epsilon u_t^\epsilon = 0 \quad (2.3)$$

³¹The exogenous processes are AR(1) and the system of equations characterizing the equilibrium is of first order. When written in matrix form it is easy to show that there is a representation as a SVAR(∞).

³²At the end of Section 7, we show the sensitivity of our estimation procedure to alternative wage setting equations.

Similarly, constructing $\mathbb{E}_{t-1}(\cdot) - \mathbb{E}_{t-2}(\cdot)$, we obtain:

$$\left(\begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \rho_1 + \begin{bmatrix} 0 & 1 - \lambda & 0 \end{bmatrix} \right) \Lambda \begin{bmatrix} u_{t-1}^\epsilon \\ u_{t-1}^z \\ u_{t-1}^\gamma \end{bmatrix} + \lambda(\rho_\epsilon - 1)\sigma_\epsilon u_{t-1}^\epsilon = 0 \quad (2.4)$$

where ρ_1 is the matrix collecting the first order autoregressive coefficients in the reduced form VAR.

The above equations (2.3) and (2.4) have to hold for all realizations of the shocks. In particular, equation (2.3) gives us two linear restrictions in the elements of Λ for given parameters in the wage setting equation when there are either contemporaneous discount rate or productivity/markup shocks. These two restrictions, together with the six restrictions coming from the orthogonalization of the shocks, are sufficient to identify the column in the impulse response matrix Λ corresponding to the leisure shock (u_t^ϵ). In order to identify the discount rate and productivity/markup shocks (u_t^γ, u_t^z) we proceed as follows. From equation (2.4), we obtain two extra linear restrictions that hold when there is a lagged discount rate shock or a lagged productivity/markup shock. However, these restrictions alone cannot “separate” the discount rate from the productivity/markup shocks because they are identical for both. Therefore, we use the sign of the impact elasticities from our model to a discount rate and productivity/markup shock (u_t^γ and u_t^y), respectively. Specifically, we search over all linear combinations $\psi \in [0, 1]$ of the independent restrictions coming from equation (2.4) such that a discount rate (productivity shock/markup) shock: (i) moves prices and employment in the same (opposite) direction on impact, and (ii) moves real wages and employment in opposite (same) direction on impact. If more than one linear combination of the restrictions satisfy these, we pick the one that is closer to giving equal weighting to both restrictions.

For completeness, the matrix Λ solves the system:

$$\begin{aligned} \begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \Lambda \begin{bmatrix} 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} &= \begin{bmatrix} 0 & 0 \end{bmatrix} \\ \left(\begin{bmatrix} \lambda & -1 & \frac{\lambda}{\phi} \end{bmatrix} \rho_1 + \begin{bmatrix} 0 & 1 - \lambda & 0 \end{bmatrix} \right) \Lambda \begin{bmatrix} 0 \\ \psi \\ 1 - \psi \end{bmatrix} &= 0 \\ \Lambda \Lambda' &= V \end{aligned}$$

It is worth noting that, when $\lambda = 1$, this procedure cannot identify all columns in the impulse response matrix Λ because the system above is underdetermined (i.e., (2.4) implies linear restrictions that are merely linear combinations of the restrictions implied by equation (2.3)). Therefore, some degree of wage stickiness is key for identification of the shocks through this procedure. The next section shows how to estimate λ using regional data—thus linking particular regional patterns to particular aggregate shock decompositions when combined with the procedure in this section.

2.6 ESTIMATING THE WAGE SETTING EQUATION USING REGIONAL DATA

In this section, we discuss how we estimate λ and ϕ which are necessary inputs in our shock identification procedure. Given the above assumptions, the aggregate and local wage setting equations can be expressed as:

$$\pi_t^w = \lambda(\pi_t + \frac{1}{\phi}(n_t - n_{t-1})) + (1 - \lambda)\pi_{t-1}^w + \lambda(u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1})$$

$$\pi_{kt}^w = \lambda(\pi_{kt} + \frac{1}{\phi}(n_{kt} - n_{kt-1})) + (1 - \lambda)\pi_{kt-1}^w + \lambda(u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1}) + \lambda v_{kt}^\epsilon$$

The aggregate and local wage setting curves are functions of the Frisch elasticity of labor supply (ϕ) and the wage stickiness parameter (λ). There is a literature on estimating micro and macro labor supply elasticities. However, it is hard to estimate the degree of wage stickiness using aggregate data given the small degrees of freedom inherent to aggregate

data and given that at the aggregate level it is hard to isolate movements in employment growth and price growth that are arguably uncorrelated with the aggregate leisure shock (u_t^ϵ). In some instances, regional data can be used to estimate these parameters.

In order for regional data to be used to estimate λ and ϕ , one of the following must hold: either (1) the leisure shock has no regional component ($v_{kt}^\epsilon = 0$) or (2) the regional component of the leisure shock must be uncorrelated with changes in local economic activity (i.e., $cov(v_{kt}^\epsilon, (n_{kt} - n_{kt-1})) = 0$ and $cov(v_{kt}^\epsilon, \pi_{kt}) = 0$). The latter condition holds if a valid instrument can be found that isolates movement in $n_{kt} - n_{kt-1}$ and π_{kt} that is orthogonal to v_{kt}^ϵ . In this section, we estimate λ and ϕ using regional data on prices, wages and employment growth during the Great Recession. We argue that state-level leisure shocks were small during the Great Recession, thus allowing us to estimate λ and ϕ by OLS. Additionally, we use state-level house price variation during the early part of the Great Recession as an instrument to isolate movements in $n_{kt} - n_{kt-1}$ and π_{kt} that are orthogonal to local leisure shocks. Both procedures yield estimates of λ and ϕ that are fairly similar.

2.6.1 ESTIMATING EQUATION AND IDENTIFICATION ASSUMPTIONS

Formally, we estimate the following specification using our state-level data:

$$\pi_{kt}^w = b_0 + b_1\pi_{kt} + b_2(n_{kt} - n_{kt-1}) + b_3\pi_{kt-1}^w + \Psi D_t + \Gamma X_k + e_{kt}$$

where $b_1 = \lambda$, $b_2 = \lambda/\phi$, $b_3 = (1 - \lambda)$, and $b_0 = \lambda(u_t^\epsilon - (1 - \rho_\epsilon)\epsilon_{t-1})$. Any aggregate leisure shocks are embedded in the constant term. The local error term includes λv_{kt}^ϵ as well as measurement error for the local economic variables. We estimate this equation pooling together all annual employment, price and wage data for years between 2007 and 2011. When estimating the above regression, we include year fixed effects (D_t). This ensures that we are only using the cross-state variation to estimate the parameters. We estimate this equation annually because we only have annual measures of wages at the state level. Our annual nominal wage measures at the state level are the composition adjusted nominal log wages computed from the American Community Survey discussed above. π_{kt}^w , therefore, is the log-growth rate in adjusted nominal wages in the state between year t and $t - 1$. Our measure of employment growth at the state level is calculated using data from the US Bureau of Labor Statistics. The BLS reports annual employment counts and population numbers for each state in each year. We divide employment counts by population to make an annual

employment rate measure for each state. $n_{kt} - n_{kt-1}$ is the log-change in the employment rate between year t and $t - 1$. π_{kt} is log-change in the average price index in each state i in year t . In our base specification, we use the retail scanner data local inflation rate scaled to account for the difference between the local non-tradable share in the retail sector and the composite consumption good. In alternative specifications, we use the raw inflation rate from the retail scanner data as our measure of local inflation. Finally, in some specifications, we include controls for the state’s industry mix in 2007. This allows for the possibility that local leisure shocks, to the extent that they exist, may be correlated with the state’s industry structure. Given that we have observations on 48 states for 4 years of growth rate data, our estimating equation includes 192 observations in our base specification. We also show results restricting our data to the period from 2007 to 2009, before the large changes in unemployment benefits extension starting in 2010.

Two additional comments are needed about our estimating equation. First, the theory developed above implies that $b_1 + b_3 = 1$. We impose this condition when estimating the cross-state regression. Second, we believe our local wage and price indices are measured with error. The measurement error, if classical, will attenuate our estimates of b_1 and b_3 . Additionally, because we are regressing wage growth on lagged wage growth, any classical measurement error in wages in year t will cause a negative relationship between wage growth today and lagged wage growth. We proceed as follows to deal with this issue. Given the large sample sizes on which our wage indices (price indices) are based, we split the sample in each year and compute two measures of wage indices (price indices) for each state in each year. For example, if we have 1 million observations in the 2007 American Community Survey, we split the sample into two distinct samples with 500,000 observations each. Within each sub-sample, we compute a wage measure for each state. Both these wage measures are measured with error. Then, we use the growth rates in wages in one sub-sample as an instrument for growth rate in wages in the other sub-sample. We discuss these procedure in detail in the Online Appendix that accompanies the paper. As we show in that appendix, the procedure corrects the attenuation bias from measurement error in our estimates.³³

In order to recover unbiased estimates of λ and ϕ via OLS, we must assume $v_{kt}^e = 0$. The assumption that there are no local leisure shocks cannot generically be true. However, in the Online Appendix, we provide some evidence suggesting that this assumption may be roughly

³³This procedure is similar to the split sample instrumental variable estimation in Angrist and Krueger (1995).

valid during the 2007-2011 period. We show that many potential leisure shocks highlighted in the literature to explain the Great Recession had large aggregate components but varied little in across US states. For example, the decline in routine jobs (Jaimovich and Siu (2014), Charles et al (2013, 2015)) was dramatic at the aggregate level during the 2007-2011 period, but occurred in all US states with roughly equal propensity. We show these results in Online Appendix Figure R7. Likewise, some have argued that the expansion of government policies acted like a leisure shock that discouraged work (Mulligan (2012)). We show that many of these government policies—such as the expansion of the Supplemental Nutrition Assistance Program (SNAP) —was large at the aggregate level but had little cross state variation. In Online Appendix Figure R4, we show that SNAP benefits per recipient increased by roughly 30 percent between 2007 and 2011. Because the increase in per recipient benefit occurred at the federal level, there was statutorily no variation in per recipient benefits across US states during this time period.³⁴

One policy that has received considerable attention in its potential to act as a labor supply shock is the differential extension of the duration of unemployment benefits across states during the Great Recession. By law in 2010, weeks of unemployment benefits were tied to the state’s unemployment rate. However, as of 2010, most US states met triggers that resulted in the duration of unemployment benefits being close to the maximum of 99 weeks. These states comprised the bulk of the US population. However, some smaller states, mostly in the Plains region of the US, had smaller employment declines and, as a result, had a smaller extension of unemployment benefits.³⁵ Despite the fact that there was little population-weighted variation across states in unemployment benefit extensions during the Great Recession, we still perform two additional robustness exercises to account for the fact that the small policy differences across states that did occur may have discouraged labor supply. First, when using our full time period, we exclude any state that had less than 85 weeks of unemployment benefit extensions leaving us with a sample of states that had essentially no remaining variation in unemployment benefit extensions. Given that the

³⁴In the Online Appendix, we also show that there was no systematic variation in state labor income tax rates during the 2007-2010 period. Additionally, we show that there was little state variation in Federal programs to help underwater homeowners (like HAMP) that occurred during the 2007-2010 period. The reason is that take up rates of the program during this time period were very low (with take up rates being essentially zero prior to 2010).

³⁵States also had some discretion as to whether they opted into the program. This explains why some states did not have the maximum weeks of unemployment benefits even when their unemployment rate was higher. We discuss these policies and how they varied across states in detail in the Online Appendix.

exclude states were small in population terms, such exclusion had essentially no effect on our estimates. Additionally, we re-estimate our key parameters using only data prior to 2010. Prior to 2010, the duration of extended unemployment benefits were the same across all states. We discuss these results below.

While we defend that OLS estimation of the above equation yields unbiased estimates of λ and ϕ using cross state variation during the Great Recession, it is impossible to completely rule out that leisure shocks are causing some of the variation in state business cycles during this period. To further explore the robustness of our results, we also estimate IV specifications of the above equation. Following the work of many recent papers, including Mian and Sufi (2014), we use contemporaneous and lagged variation in local house prices as our instruments for local employment and price growth. The argument is that local house price variation during the 2007-2011 period (in our base specification) or during the 2007-2009 period (in our restricted specification) is orthogonal to movements in local leisure shocks. This seems like a plausible assumption for the 2007-2009 period as state policy changes did not occur prior to 2009. In the Online Appendix, we discuss the IV procedure in detail. We also show that contemporaneous housing price growth strongly predicts contemporaneous employment growth and lagged measures of housing growth predicts price growth. We find that IV and OLS estimates are very similar.

Before turning to the estimation, it is also worth discussing the no cross-state migration assumption that we have imposed throughout. If individuals were more likely to migrate out of poor performing states and into better performing states, our estimated labor supply elasticities from the state regression may be larger than the aggregate labor supply elasticity. While theoretically interstate migration could be problematic for our results, empirically it is not the case. Using data from the 2010 American Community Survey, we compute migration flows to and from each state and, then, construct a net-migration rate for each state. As documented by others, we find that the net migration rate was very low during the Great Recession (Yagan 2014). This can be seen from Appendix Figure A1. Both the low level of inter-state migration and the fact that it is uncorrelated with employment growth during this period makes us confident that our estimated parameters of our local wage setting curve can be applied to the aggregate.

2.6.2 ESTIMATES OF λ AND ϕ

Column 1 of Table 2.2 shows the estimates of our base OLS specification where we use all data from 2007-2011 and do not include any additional controls. Our base estimates are $b_1 = 0.69$ (standard error = 0.13) and $b_2 = 0.31$ (standard error = 0.08). As noted above, b_1 is λ and b_2 is λ/ϕ . Given our base estimates, the cross sectional variation in prices and wages implies a labor supply elasticity of 2.2. Standard macro models imply a labor supply elasticity of 2 to 4 based on time-series variation. The estimates from the cross-section of states are in-line with these macro time-series estimates. Our base estimate of $\lambda = 0.69$ suggests only a modest amount of wage stickiness. Perfectly flexible wages imply $\lambda = 1$ while perfectly sticky wages imply $\lambda = 0$. In other words, lagged wage growth predicts current wage growth conditional on current employment and price growth, but the effect is much less than one for one (as would be implied by perfectly sticky wages). Below, we show that similar regressions run on aggregate data yield much smaller estimates of λ implying a greater amount of wage stickiness.

Table 2.2: Estimates of λ and $\frac{\lambda}{\phi}$ using Cross- Region Data

	Specification						
	OLS			IV			
	2007-2011		(3)	2007-2009		2007-2011	2007-2009
(1)	(2)	(5)		(6)	(7)	(8)	
λ	0.69 (0.13)	0.69 (0.13)	0.75 (0.13)	0.73 (0.17)	0.73 (0.17)	0.77 (0.13)	0.79 (0.18)
λ/ϕ	0.31 (0.08)	0.32 (0.08)	0.31 (0.07)	0.39 (0.09)	0.39 (0.10)	0.76 (0.17)	0.99 (0.25)
Implied ϕ	2.2	2.2	2.4	1.9	1.9	1.0	0.8
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry Controls	No	Yes	Yes	No	Yes	Yes	Yes
Scaling Factor of Prices	1.4	1.4	1.0	1.4	1.4	1.4	1.4

Note: Table shows the estimates of λ and $\frac{\lambda}{\phi}$ from our base wage setting specification using the regional data. Each observation in the regression is state-year pair. Each column shows the results from different regressions. The regressions differ in the years covered and additional control variables added. The first three columns show the OLS results using all local data between 2007 and 2011. Columns 4 and 5 show OLS results using only data from 2007 through 2009. The final two columns show IV results for the different time periods. In the IV specifications, we instrument contemporaneous employment and price growth with contemporaneous and lagged house price growth. We adjust for measurement error in wage growth, lagged wage growth, and price growth using the split sample methodology discussed in the Online Data Appendix. All regressions included year fixed effects. All standard errors are clustered at the state level.

Columns 2 and 3 of Table 2.2 show a variety of robustness checks for our base estimates. In column 2 we include industry controls. Specifically, we include the share of workers in 2007 working in manufacturing occupations or in routine occupations. This allows us to proxy for different degrees of wage stickiness or different potential leisure shocks that are correlated with industrial mix. In column 3, we use the actual retail scanner data price index as opposed to the scaled price index representative of the composite local consumption good. Neither the inclusion of controls for local industry mix nor changing the scaling on local retail price variation affect our estimates of λ and ϕ in any meaningful way. In columns 4 and 5, we re-estimate our base specification with and without industry controls using only data from 2007-2009 prior to the changes in national policy extending unemployment benefit duration. Again, our estimates λ and ϕ remain 0.73 and 1.9, respectively.³⁶ Finally, in columns 6 and 7, we show our IV estimates for the 2007-2011 period and the 2007-2009 period where we instrument local employment growth and local price growth with contemporaneous and one lag of local house price growth. Our estimates of λ and λ/ϕ are 0.77 (standard error = 0.13) and 0.76 (standard error = 0.17) implying an estimated Frisch elasticity of 1.0.

Regardless of our specification we estimate labor supply elasticities of between roughly 1.0 and 2.0. More importantly, all of our estimates imply a fair degree of wage flexibility with our estimates of λ ranging from about 0.7 to 0.8. These results are consistent with the patterns shown in Figure 2.1 where local wages co-moved strongly with local employment during the Great Recession. The estimation that wages are fairly flexible is a key insight that is important for our main results in the next section and has broader implications for the literature. In the context of the methodology we presented in the previous section as well as our monetary union model, it is hard to get aggregate “demand” shocks to be the primary shock driving economic conditions during the Great Recession if wages are fairly flexible. In other words, if wages were sticky enough in the aggregate to have “demand” shocks be the primary driver of aggregate employment decline during the recent recession, we would not have observed wages moving as much as they did in the cross-section of states during the same time period.

To show the stark difference between local and aggregate estimates of wage stickiness, we use aggregate data on prices, nominal wages, and employment between 1976 and 2012. This is the same data that we will use in our SVAR estimation in the next section. Given the

³⁶Additionally, we estimated our base specification excluding CA, NV, AZ, and FL. In both cases, our estimates were nearly identical to our base specification in column 1 of Table 3.

short time-series sample, power is an issue. However, across all specifications we explored, estimates of λ using aggregate data ranged from about 0.4 to 0.6. These estimates are below the estimates of 0.7 to 0.8 using local variation. If aggregate leisure shocks occur along with shocks that shift labor demand, wages will appear sticky in the aggregate time-series. The assumption of no aggregate leisure shocks is a common one when estimating wage stickiness using aggregate data (see, for example, Christiano et al. (2015b)).

2.7 THE US GREAT RECESSION: FROM REGIONAL TO AGGREGATE

The cross-regional facts presented above represent a puzzle. Aggregate nominal wages did not fall much (relative to trend) during the Great Recession. However, local nominal wages and employment were significantly, negatively correlated. Why did aggregate wages respond so little during the Great Recession while there was a strong relationship across states?

One potential explanation is that a series of shocks made aggregate employment fall. Some of these shocks put downward pressure on wages while others put upward pressure, thus making wages seem unresponsive. However, if the shocks putting upward pressure on wages were purely aggregate, they would be differenced out when considering variation across states—thus resulting in the observed negative correlation between employment and wages across states. Our methodology allows us to quantify the relative magnitudes of these shocks and to assess their contributions to the behavior of prices, wages and employment.

2.7.1 FINDINGS IN THE AGGREGATE

We follow the procedure described in Section 2.5. We first estimate the VAR with two lags in aggregate employment growth, price growth and nominal wage growth via OLS equation by equation using annual data from 1976 to 2012. We obtain sample estimators of the covariance matrix $\hat{V} = \frac{UU'}{Years - \#Variables - \#Lags}$ from reduced form errors U .

We construct aggregate variables that are comparable to our regional measures. Given that our cross-sectional equations are estimated using annual data, we analogously define our aggregate data at annual frequencies. We use data from the CPI-U to create our measure of aggregate prices. Specifically, we take log-change in the CPI's between the second quarter of year t and $t - 1$ for our measure of p_t . For n , we use BLS data on the aggregate employment to population rate of all males 25-54. We choose this age range so as to abstract from the downward trend in employment rates due to the aging of the population over the last 30

years.³⁷ Finally, we use data from the Current Population Survey (CPS), discussed above, to construct our aggregate composition adjusted wage measure. As with the CPI, we take the log-change in this wage measure between t and $t - 1$ for our measure of w_t . For all data, we use years between 1976 and 2012.

Figures B.3, B.4, and B.5 report the impulse response of aggregate employment, nominal wages and price growth to each of the shocks using our benchmark estimates for λ and ϕ reported in column 1 of Table 2.2 ($\lambda = 0.69$ and $\phi = 2.2$). Figure B.3 shows their behavior after a one-standard deviation discount rate (γ) shock. Qualitatively, after a discount rate shock both prices and employment increase sharply relative to trend while real wages decline slightly relative to trend. These results are identical to the theoretical predictions shown in Table B.3. Figure B.4 shows the impulse responses to a one-standard deviation productivity/markup (z) shock. Prices decrease on impact while employment increases sharply. Nominal wages, however, only decline slightly. Again, these predictions match the predictions of our simple theoretical model shown in Table B.3. While both γ and z shocks increase employment, the γ shock puts upward pressure on prices while the z shock puts downward pressure on prices. Figure B.5 shows the impulse response of employment, prices, and nominal wages to the leisure shock. Upon impact, the leisure shock reduces employment and prices while it increases wages. Again, these predictions match the predictions from the benchmark monetary union model.

We now turn to quantifying the contribution of each shock to explaining the behavior of the aggregate US economy during the Great Recession. To do so, we present the counterfactual cumulative response of each individual variable when we feed the VAR with the sequence of shock realizations between 2008 and 2012, one at a time.³⁸ Our analysis suggests that “demand” shocks cannot be solely responsible for the employment decline during the Great Recession. If “demand” shocks (i.e., discount rate shocks) were solely responsible, price and wage inflation would have been lower. Instead a combination of “supply” shocks (i.e., productivity/markup and leisure shocks) explain the “missing price and wage deflation”.³⁹

³⁷We detrend all data when estimating the VAR. Specifically, we allow for a linear trend in the employment to population ratio between 1978 and 2007. For the price inflation rate and the nominal wage inflation rate, we use an HP filter (with a smoothing parameter of 100). Given that we detrend the data, our results are essentially unchanged when we use the employment to population ratio for all individuals as opposed to using it just for prime age males.

³⁸For the interested reader, the actual realizations of the shocks we estimate can be seen in Figure B.6.

³⁹The robust growth in consumer prices during the recession is also viewed as a puzzle for those that believe that the lack of aggregate demand was the primary cause of the Great Recession. For discussions of the “missing deflation”, see Hall (2011), Ball and Mazumder (2011), Stock and Watson (2012), and Del

Figure 2.2: Counterfactual Employment Response



Note: Figure shows the cumulative response of employment when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Figure 2.2 presents the counterfactual employment response. Employment fell by more than 4 percent between 2008-2009 (relative to trend) and remained at this low level thereafter. The counterfactual exercise shows that the productivity/markup, discount rate, and leisure shocks contributed about the same amount to the initial decline during the 2008-2009 period (each explaining roughly one-third of the aggregate employment decline). However, the discount rate and leisure shocks do not explain any of the persistence in the employment decline post 2009. Instead, it is the productivity/markup shock that explains most of the sluggish response of employment post 2009.

Figures B.7 and B.8 help understand the “missing price and wage deflation puzzle”. Figure B.7 shows the counterfactual price response to each of the shocks. Aggregate prices fell relative to trend between 2008 and 2009 and quickly stabilized thereafter despite the weak employment situation post-2009. This is the sense in which there was “missing price deflation”. Both the discount rate and the leisure shock put downward pressure on aggregate prices. However, the productivity/markup shock put upward pressure on aggregate prices. The counterfactual analysis shows that if the economy had only been hit by the productiv-

Negro et al. (2015).

ity/markup shock, prices would have risen (by upwards of 1 percent) relative to trend during the Great Recession. Instead, we find that it is this countervailing productivity/markup shock that arises as the explanation for the missing deflation puzzle—particularly post 2009. This finding is consistent with the results of Christiano et al. (2015a).

Figure B.8 shows the cumulative nominal wage response to each of the shocks. Again, the figure shows the “missing wage deflation puzzle” during the Great Recession. Throughout the recession, nominal wage growth was close to zero (relative to trend). However, if the economy had only experienced the discount rate shock, nominal wages would have fallen by roughly 1.5 percent relative to trend by 2009 and would have remained below trend in 2011. It is the leisure and productivity/markup shocks that explain why nominal wages did not fall during the Great Recession.

2.7.2 SENSITIVITY TO ALTERNATIVE PARAMETERS λ AND ϕ

How do our estimated parameters affect our employment, price and wage counterfactuals? In Table B.5, we report the contribution of each shock to aggregate employment declines implied by different combinations of $\{\phi, \lambda\}$. We do this for both the initial years of the recession (2008 to 2009) and over the longer period encompassing the recovery (2008 to 2012). Each cell in Table B.5 shows how much of the employment change during the time period can be attributed to the discount rate shock (γ) and how much can be explained by the productivity/markup shock (z). The sum of all three shocks sums to 100 percent. So, the difference between the sum of the γ and z contributions and 100 percent is attributed to the leisure shock (ϵ). The qualitative conclusions of the previous section still hold for the range of $\{\phi, \lambda\}$ estimates in Table 2.2. These go from roughly 0.7 to 0.8 for λ and from roughly 1.0 to 2.5 for ϕ .

Table B.5 offers several further results worth discussing. First, we observe that the relative importance of the leisure shock vis-a-vis the discount rate and productivity/markup shocks combined is governed by the Frisch labor supply elasticity (ϕ). We estimate a relatively large elasticity, in the range of that used to calibrate standard macro models.⁴⁰ However,

⁴⁰This result may be of independent interest to the reader familiar with the macro v. micro labor supply elasticities (see Chetty, Guren, Manoli, and Weber (2011)). Using cross-sectional data (same as in most of the micro labor-supply elasticity literature) we arrive at an estimate similar to the macro elasticity (estimated from aggregate time-series data). We believe this is because the regional variation in employment rates that we use to estimate this elasticity only incorporates the extensive margin adjustment in the labor supply, which is the same margin that is most important in accounting for aggregate fluctuations in total hours over the business cycle.

suppose we used a much lower elasticity instead, e.g., $\phi = 0.5$, which is in line with some microeconomic estimates in the literature. In this case, the leisure shock would account for a much larger fraction of the employment decline in the Great Recession. In other words, if labor supply is fairly elastic, large movements in employment are consistent with relatively small movements in real wages (as we observe in US data), without the need of large leisure shocks. While this sensitivity analysis results from re-estimating the shocks under different parameterizations for ϕ using our procedure, the intuition is in line with the benchmark monetary union model from Section 4.

The intuition for the decomposition between discount rate and productivity/markup shocks is more subtle but also in line with the our benchmark monetary union model. We find that the degree of wage flexibility (λ) affects their relative contribution to the remaining, unexplained part by the leisure shock alone. For example, if we increased the degree of wage flexibility,⁴¹ the productivity/markup shock would account for a much larger fraction of the employment decline in the Great Recession. Theoretically, it is clear that when λ is large, the discount rate shock should not matter much for the determination of employment. To see this, consider the extreme case where wages are perfectly flexible and the discount rate shock is only composed of the monetary shock. Then the equilibrium in the simple theoretical model satisfies monetary neutrality. We formalized this point in Section 2.4.6 when we derived the model’s implied elasticity of aggregate employment to a discount rate shock. Conversely, when wages are very rigid ($\lambda = 0.1$), our procedure suggest that discount rate shocks explain essentially all of the decline in the early part of the recession and much of the persistence in employment decline during the 2008-2012 period.

2.7.3 SENSITIVITY TO ALTERNATIVE IDENTIFYING ASSUMPTIONS

The above results are based on the particular functional form of our wage setting equation:

$$W_{kt} = (P_{kt} e^{\epsilon_{kt}} (N_{kt})^{\frac{1}{\phi}})^{\lambda} (W_{kt-1})^{1-\lambda}$$

This wage setting equation reflects our assumption of GHH preferences as well as no forward looking behavior when wages are reset. Both of these assumptions were made for

⁴¹It is worth mentioning that for large values of λ and small values of ϕ , the results in Table B.5 become rather sensitive to small variations in parameters. This is because our shock identification procedure needs a certain degree of wage stickiness, as explained in the Section 2.5. For example, for values of λ around 0.95 it is not possible to identify the productivity/markup and discount rate shocks.

tractability. In this sub-section, we explore the sensitivity of our results to relaxing both of these assumptions.

In Appendix B1, we derive the aggregate and local wage setting equations under a broad set of utility functions where consumption and leisure are non-separable. This class of utility functions allows for arbitrarily large income and substitution effects. As we show in the appendix, the use of local consumption data allows us to estimate the extent of wage stickiness as well as to estimate the parameters that encompass both the income and substitution effects on labor supply. In particular, we can estimate the following equation using local data:

$$\pi_{kt}^w = \tilde{b}_t + \tilde{b}_1 \pi_{kt} + \tilde{b}_2 (n_{kt} - n_{kt-1}) + \tilde{b}_3 \pi_{kt-1}^w + \tilde{b}_4 (c_{kt} - c_{kt-1}) + \Psi D_t + \Gamma X_k + e_{kt}.$$

This equation is identical to our estimating equation above aside from the addition of local consumption growth (i.e., $c_{kt} - c_{kt-1}$). As outlined in Appendix B1, the coefficients \tilde{b}_1 and \tilde{b}_3 sum to 1 even under the broader preference specification. We impose this restriction when estimating the modified equation. We measure local real consumption growth using the log-change in real retail expenditures at the state level computed within the Nielsen sample. We obtain real expenditures by deflating nominal expenditures with our local price indices.⁴² For our base specification, our estimates of \tilde{b}_1 , \tilde{b}_2 , and \tilde{b}_4 are 0.72 (standard error = 0.12), 0.25 (standard error = 0.08), and 0.16 (standard error = 0.06), respectively. A positive and significant coefficient on real consumption growth (\tilde{b}_4) reflects the presence of income effects on labor supply. Controlling for this income effect, our estimate of wage flexibility (\tilde{b}_1) is slightly higher than our base specification where income effects are not allowed.

In the aggregate wage setting equation, we can substitute out consumption growth using the model definition ($c_t = w_t + n_t - p_t$). Appendix B1 shows that the aggregate wage setting equation still takes the following form:

$$\pi_t^w = \lambda \pi_t + \frac{\lambda}{\phi} (n_t - n_{t-1}) + (1 - \lambda) \pi_{t-1}^w + \frac{\lambda}{1 - \omega} \epsilon_t$$

where ω is a parameter that represents the strength of the income effect on labor supply

⁴²This measure of real expenditures is (1) highly correlated with measures of local employment and (2) highly correlated with the BEA's recent state level personal expenditures measure. Our results are similar if we use the BEA's local consumption measure. However, we prefer our measure given that much of the BEA's local consumption measure is imputed (where the imputation uses local employment measures).

(and maps directly to \tilde{b}_4 from the above local labor supply regression, see equation (B.1) in Appendix B1). Aside from the coefficient scaling the aggregate leisure shock, this equation is identical to the identification restriction we imposed when estimating the aggregate SVAR. The only difference is that there is no longer a direct mapping between λ and λ/ϕ —in the above aggregate wage setting equation that we impose to identify the SVAR—and the reduced-form parameters \tilde{b}_1 and \tilde{b}_2 —in the local wage setting equation. However, as shown in Appendix B1, there is still a one-to-one mapping between the parameters we estimate from the local regression (\tilde{b}_1 , \tilde{b}_2 , and \tilde{b}_4) and the aggregate parameters we need to identify the SVAR (λ and ϕ). With the correctly specified λ and ϕ , we can just use the matrix in Table B.5 to read off the decomposition of shocks during the Great Recession. While λ and ϕ are no longer structural parameters (instead being combinations of structural parameters), knowing them still helps identifying the aggregate SVAR. Using our estimates of \tilde{b}_1 , \tilde{b}_2 , and \tilde{b}_4 and the procedure developed in Appendix B1, we estimate λ and ϕ (allowing for income effects on labor supply) to be 0.68 and 2.0, respectively. These parameters are nearly identical to our base specification without income effects. The take-away from this sensitivity exercise is that abstracting from income effects on labor supply is not biasing our decomposition of the shocks driving aggregate employment declines during the Great Recession in any meaningful way.

In Appendix B2, we specify an alternative wage setting equation allowing for forward looking behavior when wages are reset. We show that ignoring forward looking wage setting behavior biases up our estimates of wage flexibility. That is, the amount of wage flexibility that we estimate using cross-state variation is too large relative to the true amount of wage flexibility in the aggregate. We show that the bias depends on two parameters: (1) the extent to which firms put weight on forward looking behavior when setting wages (which we call κ) and (2) the underlying persistence process of local wages (which we call $\bar{\rho}_w$). Assuming that at the aggregate level the monetary authority wants to stabilize expected nominal wage growth, we quantify the extent to which our estimates of λ from the local regressions are biased upwards as a function of these two parameters. Under a range of plausible parameter estimates for κ and $\bar{\rho}_w$, we show that the bias is quite small. For our baseline parameter estimates, we show that $\kappa(1 - \bar{\rho}_w)$ must exceed 2.5 for the true λ to be lower than 0.4. This is an order of magnitude larger than any plausible parametrization for either κ or $\bar{\rho}_w$. Moreover, as seen in Table B.5, our estimates of the role of demand shocks in explaining employment decline during the Great Recession are broadly similar for values of λ between

0.4 and 0.69. These results suggest that our abstraction from including expectations in our wage setting equation is not quantitatively altering the paper’s conclusions.

2.7.4 DISCUSSION: AGGREGATE V. REGIONAL SHOCK DECOMPOSITION

The above results suggest that aggregate wages appeared sticky during the Great Recession because a combination of aggregate shocks resulted in relatively offsetting effects on aggregate wages. By extending our procedure to allow for the estimation of regional shocks in a regional SVAR, we find that the discount rate shock explains roughly 80 percent of the change in non-tradable employment between 2007 and 2010 across states. This is consistent with results found in Mian and Sufi (2014) suggesting that housing price declines explained much of the cross-state variation in non-tradable employment. We conclude that even though discount rate shocks only explained a portion of aggregate employment decline early in the Great Recession and very little of its persistence, the discount rate shock was primarily responsible for much of the cross-state variation during this time period.

We relegate the discussion of the estimation procedure and identification assumptions to the Online Appendix—given that the estimation of regional shocks is not central to the paper. However, this exercise allow us to highlight the difficulty in extrapolating findings from cross-region variation to interpret aggregate time-series patterns. The fact that local discount rate shocks explain much of the cross-state variation in employment during the Great Recession does not imply that an aggregate discount rate shock explains much of the aggregate time-series variation in employment during the Great Recession. If the discount rate shock would have been the main driver of aggregate employment decline during the Great Recession then aggregate wages would have behaved similarly to state-level wages.

2.8 CONCLUSION

Regional business cycles during the Great Recession in the US were strikingly different than their aggregate counterpart. This is the cornerstone observation on which we built this paper. Yet, the aggregate US economy is just a collection of these regions connected by trade of goods and assets. We argued that their aggregation cannot be arbitrary and that regional business cycle patterns have interesting implications for aggregate business cycles.

Our paper offers four takeaways. The first is that the relationship between wages and

employment in the aggregate time-series during the 2006-2011 period is very different than the cross-state relationship between these variables during the same time period. For example, while aggregate wages appeared to be sticky despite aggregate employment falling sharply, both local nominal and real wages co-varied strongly with local employment growth in the cross-section of US states. Both documenting the regional facts and the creation of the underlying local price and wage indices are the first innovations of the paper.

The second take-away is that wages seem to be modestly sticky when using cross-state variation to estimate our wage setting equation. The amount of wage stickiness is often a key parameter in many macro models. Despite its importance, there are not many estimates of the frequency with which wages adjust (particularly relative to estimates of price adjustments). We develop a procedure to estimate the amount of wage stickiness using cross-region variation. The wage stickiness parameter is key to our empirical methodology to estimate the underlying shocks and elasticities. The fact that we estimate that wages are only modestly sticky limits the importance of “demand” shocks at the aggregate level in explaining the Great Recession. If wages are only modestly sticky, aggregate demand shocks should have resulted in falling aggregate wages—which was not observed in the aggregate time-series. Regardless of the use of this parameter in our empirical work, our estimate of wage stickiness could be of independent interest to researchers.

The third take-away from this paper is developing a methodology that allows us to estimate aggregate shocks by combining aggregate and regional data. This methodology is a hybrid method that merges restrictions imposed by a theoretical model with aggregate and cross-sectional data when estimating a SVAR and identifying the corresponding shocks. We view this as a contribution to the growing literature that uses model-based structure to estimate SVARs.

Finally, the fourth take-away is perhaps the most important for the goals of the paper. We show that a combination of both “demand” and “supply” shocks are necessary to account for the joint dynamics of aggregate prices, wages and employment during the 2007-2012 period in the US. In contrast with the aggregate results, we find that discount rate shocks explain most of the observed employment, price and wage dynamics across states. These results suggest that solely using cross-region variation to explain aggregate fluctuations is insufficient when some shocks do not have a substantive regional component. The fact that aggregate wages did not fall cannot be explained by large degrees of wages stickiness. The reason aggregate wages did not fall is that the series of shocks experienced by the aggregate economy were such that

some shocks put downward pressure on prices and wages (discount rate shocks) while other shocks put upward pressure on prices and wages (productivity/markup and leisure shocks). In the cross-section, however, the discount rate shocks caused prices, wages and employment to move in the same direction. Lastly, in our calibrated monetary union model, we show that the local employment elasticity to a local discount rate shock is substantially larger than the aggregate employment elasticity to an identically-sized aggregate discount rate shock. These results suggest that even when the aggregate and regional shocks are the same, it is hard to draw inferences about the aggregate economy using regional variation. Collectively, our results suggest that researchers should be cautious when extrapolating cross-sectional variation to make statements about aggregate business cycles.

Chapter 3

FROM HYPERINFLATION TO STABLE PRICES: ARGENTINA'S EVIDENCE ON MENU COST MODELS (WITH FERNANDO ALVAREZ, MARTIN GONZALEZ-ROZADA AND PABLO ANDRES NEUMEYER)

3.1 INTRODUCTION

Infrequent nominal price adjustments are at the center of much literature studying positive and normative implications of inflation. In this paper, we study how price setting behavior changes with inflation, both theoretically and empirically. We consider menu cost models with idiosyncratic firm-level shocks, i.e., models where monopolistic firms set prices subject to a fixed cost of adjustment and are hit by shocks to their real marginal cost. We derive sharp predictions concerning how changes in steady state inflation affect price setting behavior at near-zero and high inflation rates. Then, we confront these predictions with their empirical counterparts using firm-level data underlying Argentina's consumer price index from 1988 to 1997. Argentina's experience provides a unique opportunity to analyze price setting behavior through the lens of menu cost models because it encompasses several years of price stability (and even deflation) as well as years of sustained, very high inflation.

We find evidence supporting many predictions of menu cost models. Our empirical findings involve confronting a number of variables (summarizing price setting behavior) across different time periods with different inflation rates. However, our theoretical results involve comparative statics about how these same variables change across steady states with different inflation rates. We show that these comparative statics essentially depend on the ratio of inflation to the variance of firm-level idiosyncratic shocks. Therefore, they are comparable to our empirical findings under the assumption that the variance of idiosyncratic shocks remains approximately constant across time periods. The first set of results concerns low inflation

economies. Both theoretically and empirically, we show that in a neighborhood of zero inflation: (i) the frequency of price changes is unresponsive to inflation, (ii) the dispersion of relative prices is unresponsive to inflation, (iii) the frequency of price increases is equal to the frequency of price decreases, (iv) conditional on a price change, the size of price increases is equal to the size of price decreases, and (v) inflation rises (falls) because the frequency of price increases rises (falls) and that of price decreases falls (rises)—as opposed to the size of price increases and decreases changing. The second set of results concerns high inflation economies. We find that: (vi) the frequency of price adjustment becomes the same for all products/firms, (vii) the elasticity of the average frequency of price changes with respect to inflation converges to two-thirds, (viii) the elasticity of the average size of price changes with respect to inflation converges to one-third, and (ix) the frequency of price decreases converges to zero. We show these are consistent with menu cost models if idiosyncratic shocks are very persistent.

We believe these results are interesting both because they underlie the welfare costs of inflation in menu cost models (as well as other models of price stickiness) and serve to test this class of models. First, the menu cost paid when changing prices is a direct welfare cost of inflation since these resources are wasted. Second, the “extra” price dispersion created by nominal variation in prices is another avenue for inefficiency in menu cost models as well as other models with sticky prices. For example, in models with an exogenous frequency of price adjustment, as described in chapter 6 of Woodford (2003).¹ Therefore, our findings (i) and (ii) imply that the welfare cost of increasing the rate of inflation is negligible around zero inflation.

We begin the paper by presenting a model of monopolistic firms that are hit by idiosyncratic shocks to their real marginal cost and face a fixed cost of changing prices. Similar models have been introduced by Barro (1972) and Sheshinski and Weiss (1977), and augmented to include idiosyncratic firm level shocks by Bertola and Caballero (1990), Danziger (1999), Golosov and Lucas (2007), and Gertler and Leahy (2008), among others. Then, in Section 3.2.1 and Section 3.2.1, we derive our comparative static results about how inflation affects price setting behavior and illustrate them with a numerical example based on the model in Golosov and Lucas (2007). The results for low inflation are new. Mainly, the

¹See Benabou (1992) and Burstein and Hellwig (2008) for earlier and recent examples of analysis that takes both effects of inflation into account, the former using heterogeneous consumers that search for products and homogeneous firms, and the latter using differentiated products in the demand side and heterogeneity in the firm’s cost.

prediction that most of the changes in inflation (ninety percent to be precise) are accounted for by changes in the difference between the frequency of price increases and decreases.² Furthermore, a new insight of this paper is that, under some simplifying assumptions, price setting behavior depends only on the ratio between inflation and the variance of the idiosyncratic shocks. In this sense, high inflation economies are equivalent to economies in which firms do not face idiosyncratic shocks for nominal price adjustment decisions. Hence, we can apply benchmark results on the effect of inflation on price setting behavior in the deterministic case (Sheshinski and Weiss (1977) and especially Benabou and Konieczny (1994)) to the case in which firms face persistent idiosyncratic shocks and there is high inflation.

In Section 3.3 we describe our dataset. We use the micro-data underlying the construction of the Argentine Consumer Price Index index from 1988 to 1997 for 506 goods covering 84 percent of expenditure. The unique feature of this data is the range of inflation during this time period. The inflation rate was almost 5000 percent during 1989 and 1500 percent during 1990. After the stabilization plan of 1991 there is a quick disinflation episode and after 1992 there is virtual price stability with some deflationary periods. Then, in Section 3.4, we describe how we estimate the frequency of price changes, the size of price changes, and the dispersion of relative prices.

We follow with our main empirical findings in Section 3.5. Here we discuss the most notable ones as well as their relationship to previous literature. We find that the frequency of price changes as a function of the rate of inflation is flat at low inflation levels and has a constant elasticity for high inflation levels (a novel empirical finding). We estimate this elasticity to be between one half and two thirds, which is close to the theoretical prediction of two thirds.³ There is a large literature that estimates the frequency of price changes. Various papers do this for different countries, for different time periods, for different rates of inflation, and for different sets of goods.⁴ A feature of our dataset (and a success of the

²Alvarez et al. (2011) and Alvarez and Lippi (2014) show that the frequency of price changes and the dispersion of relative prices are insensitive to inflation at low inflation rates under a more restrictive set of assumptions.

³To accommodate the range of estimates in the data, our simple model could be extended allowing firms to freely change prices at random times. Nakamura and Steinsson (2010) consider a version of this model (also see Dotsey and King (2005), Caballero and Engel (2007) and Alvarez et al. (2014b)).

⁴For example, for low inflation economies Álvarez et al. (2006) for the Euro area, Bils and Klenow (2004), Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008) for the US. For moderate levels of inflation, Gagnon (2009) for México, Barros et al. (2009) for Brazil, and Wulfsberg (2010) for Norway. Baharad and Eden (2004), and Lach and Tsiddon (1992) for Israel as well as Konieczny and Skrzypacz (2005) for Poland look at high inflation economies, but not as high as our Argentine sample.

theory) is that our estimates of the frequency of price changes for each level of inflation in the Argentine data are similar to the estimates of the other studies with the same rate of inflation. This is illustrated in Figure 3.7 in Section 3.5.1.1.

Furthermore, we find that even though the frequency of price changes is unresponsive to inflation when inflation is low, the difference between the frequency of price increases and the frequency of price decreases is an increasing function of inflation. This is consistent with previous evidence (see, for example, Nakamura and Steinsson (2008) for the US, Gagnon (2009) for Mexico, Berardi et al. (2013) for France and Cavallo (2015) for cross-country evidence). Our contribution is to provide a new theoretical interpretation of this fact through the lens of menu cost models, which suggests a particular test of this hypothesis (see Figure 3.8). Then, we document that the cross-good dispersion of the frequency of price changes falls with inflation. As it is well known in the literature studying low inflation economies, there is large variation in the frequency of price changes across firms selling different goods. The new finding is that, as inflation rises, the frequency of price changes becomes similar across different goods—thus reflecting how idiosyncratic differences across firms disappear as a motive for changing prices when inflation is high.

Next, we are concerned with the dispersion of relative prices and size of price changes. Despite being of theoretical and practical importance because of its welfare implications, this is the first paper to look at the relationship between inflation and the dispersion of relative prices across stores selling the same good.⁵ We document that the dispersion of relative prices across stores is insensitive to inflation for low inflation rates and it becomes an increasing function of inflation for high inflation rates. However, the elasticity of the dispersion of relative prices does not reach the theoretical elasticity of one third for the inflation rates in our sample; a fact that is consistent with our numerical simulations when idiosyncratic shocks are very persistent. The average size of price changes, conditional on a price change taking place, exhibits a similar pattern. It is insensitive to inflation at low inflation rates and then starts to increase with the rate of inflation. Recently, Nakamura et al. (2016) found similar results with post 1977 BLS micro-data on US consumer prices. They find that for inflation rates under 14% per year the absolute value of price changes is insensitive to inflation.

Finally, several online appendices provide supplementary material. Online appendix C.1 contains the proofs of the propositions in Section 3.2.1. Some of these proofs are of inde-

⁵Many papers look at the dispersion of *inflation rates* across goods, which is a different concept.

pendent interest as they fully characterize the solution of the menu cost model with a closed form solution. Online appendix C.2 contains an analytical characterization of our version of the Golosov and Lucas (2007) model. Online appendix C.3 describes some details of our data. In online appendix C.4, we perform extensive robustness checks to evaluate the sensitivity of our estimates. Our findings are robust to different treatment of sales, product substitutions, and missing values in the estimation of the frequency of price changes and with respect to the level of aggregation of price changes. They are also robust to using contemporaneous inflation or an estimate of expected future inflation, for the relevant time frame, and to excluding observations corresponding to periods with inflation above some threshold for which we have reasons to believe that discrete sampling might bias the estimates. Online appendix C.5 contains a short description of the history of economic policy and inflation in Argentina for the years before and during our sample. Further details on the samples and inflation ranges of other studies in the literature are provided in online appendix C.6.

3.2 COMPARATIVE STATIC PROPERTIES OF MENU COST MODELS

In this section we study how inflation affects price setting behavior in menu cost models. We show several theoretical predictions of a class of models where a single competitive monopolist firm faces a fixed cost of changing its nominal price in the presence of idiosyncratic real marginal cost shocks and constant inflation.

This section has two parts. In Section 3.2.1 we write down a simple set up and obtain the main analytical results. In Section 3.2.1.1 and Section 3.2.1.2 we obtain the results for low and high inflation, respectively, explain the nature of the assumptions needed for the results, highlight which form of these results are already present in the literature, and discuss pros and cons of applying these comparative statics results to time series data. In Section 3.2.2 we illustrate the results with a numerical example that follows Golosov and Lucas (2007). We relax some assumptions and parameterize the model in an empirically reasonable fashion in order to show that theoretical results of Section 3.2.1 are applicable to the Argentinean experience.

3.2.1 THE BASIC MENU COST MODEL

This section presents a benchmark model to study the sensitivity of the firm's pricing decisions to changes in the rate of inflation. We study the problem of a monopolist adjusting the

nominal price of its product in an environment with inflation, idiosyncratic real marginal cost shocks, and a fixed cost (the menu cost) of changing nominal prices. We think of this problem as similar to the firm's problem in Golosov and Lucas (2007).

We assume that the instantaneous profit of the monopolist depends on its price relative to the economy (or industry) wide average price and on an idiosyncratic shock. We let $F(p, z)$ be the real value of the profit per period as a function of the log of the nominal price charged by the firm relative to the the nominal economy (or industry) wide price, p , and the idiosyncratic shock, z . We assume that the economy wide price grows at a constant inflation rate π so that when the firm does not change its nominal price its relative log-price evolves according to $dp = -\pi dt$. We also assume that F is strictly concave in its first argument. The variable $z \in Z$ is a shifter of the profit function. We also allow the menu cost to depend on z , in which case we write $C_t = c\zeta(z_t)$, where $c \geq 0$ is a constant, so $c = 0$ represents the frictionless problem. We assume that $\{z_t\}$ is a diffusion with coefficients $a(\cdot)$ and $b(\cdot)$:

$$dz = a(z) dt + b(z) \sigma dW \quad (3.1)$$

where $\{W(t)\}$ is a standard Brownian Motion so $W(t) - W(0) \sim N(0, t)$. We keep the parameter σ separately from $b(\cdot)$ so that when $\sigma = 0$ the problem is deterministic. We use $r \geq 0$ for the real discount rate of profits and adjustment costs. We let $\{\tau_i\}$ be the stopping times at which prices are adjusted and $\{\Delta p(\tau_i)\}$ the corresponding price changes, so that the problem of the firm can be written as

$$V(p, z) = \max_{\{\tau_i, \Delta p_i\}_{i=0}^{\infty}} \mathbb{E} \left[\int_0^{\infty} e^{-rt} F(p(t), z(t)) dt - \sum_{i=0}^{\infty} e^{-r\tau_i} c \zeta(z(\tau_i)) \mid z(0) = z \right] \quad (3.2)$$

with $p(t) = p(0) + \sum_{i=0}^{\tau_i < t} \Delta p(\tau_i) - \pi t$ for all $t \geq 0$ and the initial state is given by $p(0)$.

The optimal policy that solves this problem can be described by the vector $\Psi(z; \pi, \sigma^2) = [\underline{\psi}(z; \pi, \sigma^2), \bar{\psi}(z; \pi, \sigma^2); \hat{\psi}(z; \pi, \sigma^2)]$. Each element of the vector is a function of z, π, σ^2 . We include the parameters π and σ^2 as explicit arguments of the decisions rules in order to conduct some comparative statics. The functions $\underline{\psi}(z; \pi, \sigma^2)$ and $\bar{\psi}(z; \pi, \sigma^2)$ define the inaction set

$$\mathcal{I}(\pi, \sigma^2) = \{(p, z) \in \mathbb{R} \times Z : \underline{\psi}(z; \pi, \sigma^2) \leq p \leq \bar{\psi}(z; \pi, \sigma^2)\}.$$

If the firm's relative price is within the inaction set, $(p, z) \in \mathcal{I}$, then it is optimal not to change

prices. Outside the interior of the inaction set the firm will adjust prices so that its relative price just after adjustment is given by $p = \hat{\psi}(z; \pi, \sigma^2)$. Since $\{z(t)\}$ has continuous paths, all adjustments will occur at the boundary of the inaction set (given additional regularity conditions). For instance, a firm with a relative price p and an idiosyncratic shock z such that the relative price hits the lower boundary of the inaction set—i.e., $p = \underline{\psi}(z; \pi, \sigma^2)$ —will raise its price by $\Delta p = \hat{\psi}(z; \pi, \sigma^2) - \underline{\psi}(z; \pi, \sigma^2)$.

Using the optimal decision rules we can compute the density of the invariant distribution of the state, $g(p, z; \pi, \sigma^2)$, as well as the expected time between adjustments $\mathcal{T}(p, z; \pi, \sigma^2)$ starting from the state (p, z) . Note that using $g(\cdot)$ we can readily find the distribution of relative prices in the economy (or industry) and we can compute the expected time elapsed between consecutive adjustments under the invariant distribution, and its reciprocal, the expected number of adjustments per unit of time, which we denote by $\lambda_a(\pi, \sigma^2)$.

We denote by $\lambda_a^+(\pi, \sigma^2)$ and $\lambda_a^-(\pi, \sigma^2)$ the frequencies of price increases and decreases respectively. Furthermore, we let $\Delta_p^+(\pi, \sigma^2)$ be the expected size of price changes, conditional on of having an increase, and $\Delta_p^-(\pi, \sigma^2)$ the corresponding expected size of price changes, conditional on having a decrease. The expectation is $\Delta_p^+(\pi, \sigma^2) = \int_Z \left[\hat{\psi}(z) - \underline{\psi}(z) \right] \frac{g(\underline{\psi}(z), z)}{\int_Z g(\underline{\psi}(z'), z') dz'} dz$ where we omit (π, σ^2) as arguments of $g, \hat{\psi}$ and $\underline{\psi}$ to simplify notation.

3.2.1.1 COMPARATIVE STATICS WITH LOW INFLATION.

In this section we show that when inflation is zero and firms face idiosyncratic profit shocks, changes in the rate of inflation do not have a first order effect neither on the frequency of price changes nor on the distribution of relative prices. The intuition for this result is that at zero inflation, price changes are triggered by idiosyncratic shocks and small variations in inflation have only a second order effect. Moreover, we show that there is a strong symmetry in this case: the frequency of price increases and decreases as well as the size of price increases and decreases are the same.

In order to prove this result we assume that the idiosyncratic shock z has strictly positive volatility and we make some mild symmetry assumptions on the firm's problem, i.e., that the process for the shocks is symmetric around zero, and that the profit function is symmetric in the log of the static profit maximizing relative price as well as in its shifter.

More precisely we define symmetry as follows. Assume that $z \in Z = [-\bar{z}, \bar{z}]$ for some strictly positive \bar{z} . Let the profit maximizing relative price given z be $p^*(z) \equiv$

$\arg \max_x F(x, z)$. We say that $a(\cdot), b(\cdot), \zeta(z)$ and $F(\cdot)$ are *symmetric* if

$$a(z) = -a(-z) \leq 0 \text{ and } b(z) = b(-z) > 0 \text{ for all } z \in [0, \bar{z}] , \quad (3.3)$$

$$p^*(z) = -p^*(-z) \geq 0 \text{ for all } z \in [0, \bar{z}] \quad (3.4)$$

$$F(\hat{p} + p^*(z), z) = F(-\hat{p} + p^*(-z), -z) + f(z) \text{ for all } z \in [0, \bar{z}] \text{ and } \hat{p} \geq 0 , \quad (3.5)$$

$$\zeta(z) = \zeta(-z) > 0 \text{ for all } z \in [0, \bar{z}] , \quad (3.6)$$

for some function $f(z)$ and with the normalization $p^*(0) = 0$.

Let $\mu(z)$ be the density of the invariant distribution of z , when it exists. Equation (3.3) implies that the invariant distribution μ as well as the transition densities of the exogenous process $\{z_t\}$ are symmetric around $z = 0$. Equations (3.4)-(3.5) state that the profit function is symmetric around the (log) maximizing price and its cost shifter. Thus, if the price is \hat{p} *higher* than the optimal for a firm with z , profits deviate from its optimal value by the same amount as with prices \hat{p} *lower* than the optimal when the shifter is $-z$. The function f allows to have an effect of the shifter z on the level of the profits that is independent of the price.

An example of a symmetric case is

$$a(z) = -a_0 z , \quad b(z) = b_0 , \quad \zeta(z) = \zeta_0 , \quad F(p, z) = d_0 - e_0 (p - z)^2 - f_0 z \text{ so } p^*(z) = z \quad (3.7)$$

for non-negative constants $a_0, b_0, \zeta_0, e_0, d_0$ and f_0 . One way to think about the symmetry assumption is to consider a second order approximation of the profit function around the profit maximizing price, so that

$$F(p, z) = F(p^*(z), z) + \frac{1}{2} F_{pp}(p^*(z), z) (p - p^*(z))^2 + o((p - p^*(z))^2) . \quad (3.8)$$

We note that firm adjusts prices more frequently as the fixed adjustment cost c becomes smaller, thus making the term $(p - p^*(z))^2$ smaller and the quadratic approximation increasingly accurate.

We let $h(\hat{p}; \pi, \sigma^2) = \int g(\hat{p}, z; \pi, \sigma^2) dz$ be the invariant distribution of the relative prices \hat{p} , for an economy, or industry, with (π, σ) . Using h we can compute several statistic of interest, such as $\bar{\sigma}(\pi, \sigma^2)$ the standard deviation of the relative prices $\hat{p} = p - \bar{p}$. As in the case of the frequency of price changes, we include (π, σ^2) explicitly as arguments of this statistic.

Proposition 3 Assume that $z \in Z = [-\bar{z}, \bar{z}]$ for some strictly positive \bar{z} and that $F(\cdot)$, $a(\cdot)$, $b(\cdot)$ and $\zeta(\cdot)$ satisfy the symmetry conditions (3.3)-(3.6). Then, the following are true:

- i. if the frequency of price changes $\lambda_a(\pi, \sigma^2)$ is differentiable at $\pi = 0$, then the frequency of price changes is insensitive to inflation,

$$\frac{\partial}{\partial \pi} \lambda_a(0, \sigma^2) = 0,$$

- ii. if the density of the invariant $h(\hat{p}; \cdot, \sigma^2)$ is differentiable at $\pi = 0$, then the dispersion of relative prices under the invariant distribution is insensitive to inflation,

$$\frac{\partial}{\partial \pi} \bar{\sigma}(0, \sigma^2) = 0,$$

- iii. the frequencies of price changes and the size of price adjustment are symmetric at $\pi = 0$ in the sense that

$$\lambda_a^+(0, \sigma^2) = \lambda_a^-(0, \sigma^2), \quad \frac{\partial \lambda_a^+(0, \sigma^2)}{\partial \pi} = -\frac{\partial \lambda_a^-(0, \sigma^2)}{\partial \pi} \text{ and}$$

$$\Delta_p^+(0, \sigma^2) = \Delta_p^-(0, \sigma^2), \quad \frac{\partial \Delta_p^+(0, \sigma^2)}{\partial \pi} = -\frac{\partial \Delta_p^-(0, \sigma^2)}{\partial \pi}$$

where λ_a^+ is the frequency of price increases, Δ_p^+ is the average size of price increases and λ_a^- , Δ_p^- are the analogous concepts for price decreases.

The proof is in online appendix C.1.1. The main idea is to use the symmetry of F to show the results. The expected number of adjustments is symmetric around zero inflation, i.e., $\lambda_a(\pi, \sigma^2) = \lambda_a(-\pi, \sigma^2)$ for all π . Given the symmetry of the profit function we view this property as quite intuitive: a 1% inflation should give rise to as many price changes as a 1% deflation. Symmetry implies that if λ_a is differentiable then $\frac{\partial}{\partial \pi} \lambda_a(\pi, \sigma^2) = -\frac{\partial}{\partial \pi} \lambda_a(-\pi, \sigma^2)$, which establishes the first result.

Analogously, for the distribution of relative prices, the main idea is to show that the marginal distribution of relative prices is symmetric in the sense that $h(\hat{p}; \pi, \sigma^2) = h(-\hat{p}, -\pi, \sigma^2)$ for all \hat{p}, π , i.e., the probability of high relative prices with positive inflation is the same as that of low relative prices with deflation. As these symmetric functions are locally unchanged with respect to π when $\pi = 0$, inflation has no first order effect on the second moment of

inflation at $\pi = 0$. Similarly, the symmetry of the frequency and of the average size of price increases and of price decreases also follow from the symmetry assumptions.

The assumption of differentiability of λ_a and $\bar{\sigma}$ with respect to π is not merely a technical condition. The function $\lambda_a(\cdot, \sigma^2)$ could have a local minimum at $\pi = 0$ without being smooth, as it is in the case of $\sigma^2 = 0$ to which we will turn in the next subsection.⁶ Likewise, the differentiability of $h(\hat{p}; \pi, \sigma^2)$ at $\pi = 0$ requires $\sigma > 0$. In Sheshinski and Weiss's (1977) model, i.e., when σ and $\pi = 0$ the distribution h is degenerate, uniform at $\pi \neq 0$ but non-differentiable at $\pi = 0$.

Remarks and relation to the literature

We find Proposition 3's theoretical predictions interesting because they extend an important result on the welfare cost of inflation from sticky price models with exogenous price changes (e.g., the Calvo model) to menu cost models with endogenous frequency of price changes. The result is that in cashless economies with low inflation there is no first order welfare effect of inflation, i.e., the welfare cost of inflation can be approximated by a “purely quadratic” function of inflation.

Inflation imposes welfare costs through two channels in cashless economies. First, the “extra” price dispersion created by inflation is an avenue for inefficiency in models with sticky prices, since it creates “wedges” between the marginal rates of substitution in consumption and the marginal rates of transformation in production. See, for example, chapter 6 of Woodford (2003) and references therein for the analysis of this effect. Part (ii) of Proposition 3 extends this result to the menu cost model. Second, the higher endogenous frequency of price adjustments due to inflation is an obvious source of welfare losses when these adjustments are costly. Part (i) of Proposition 3 establishes that this second channel is also negligible for low inflation rates.⁷

The results of Proposition 3 are likely to apply to a wider class of models. The alert reader

⁶ We conjecture, but have not proved at this level of generality, that as long as $\sigma^2 > 0$, the problem is regular enough to become smooth, i.e., the idiosyncratic shocks will dominate the effect of inflation. For several examples one can either compute all the required functions or show that they are smooth, given the elliptical nature of the different ODE's involved. Based on this logic, as well as on computations for different models, we believe that the length of the interval for inflations around zero for which $\lambda_a(\cdot, \sigma^2)$ is approximately flat is increasing in the value of σ^2 .

⁷There are other papers that take more than one effect into account. Burstein and Hellwig (2008) compute numerical examples in a model closer to ours, which also includes the traditional money demand cost. Benabou (1992) uses a different framework, with heterogeneous consumers that search for products and homogeneous firms subject to menu costs but constant production costs.

will realize that the essential assumption is the symmetry of the profit function around the profit maximizing price, and hence if this is maintained, the result should hold. For example, a version of Proposition 3 applies to models with both observations and menu cost, Alvarez et al. (2011), to models that have multi-product firms, Alvarez and Lippi (2014), and to models that combine menu costs and Calvo type adjustments (Nakamura and Steinsson (2010) and Alvarez et al. (2014b)). In those papers the second order approximation of the profit function stated in equation (3.8) is used.⁸ Note that one can either assume symmetry directly, or obtain it approximately if the menu cost is small, since in this case a second order approximation around the static profit maximizing price is symmetric. In Section 3.2.2 we solve numerically a version of the model that is not symmetric, but with menu costs that are empirically reasonable, and confirm the results of this section. In Section 3.2.2 we also show that for reasonable parameter values the functions λ_a and $\bar{\sigma}$ are approximately flat for a wide value of inflation rates around zero.

We don't know of other theoretical results analyzing the sensitivity of $\lambda_a(\pi)$ and $\bar{\sigma}(\pi)$ to inflation around $\pi = 0$ in this set-up. However, there is a closely related model that contains a complete analytical characterization by Danziger (1999). In fact, we can show that for a small cost of changing prices, Proposition 3 holds in Danziger's characterization.

3.2.1.2 COMPARATIVE STATICS WITH HIGH INFLATION

Now we turn to the analysis of price setting behavior for large values of inflation. In highly inflationary environments, the main reason for firms to change nominal prices is to keep their relative price in a target zone as the aggregate price level grows. Idiosyncratic shocks in the high inflation case become less important and, therefore, the analysis of the deterministic case is instructive. This leads us to proceed in two steps. First we derive comparative static results in the deterministic case—i.e., when $\sigma^2 = 0$. This is a version of the problem studied by Sheshinski and Weiss (1977). Then, we study the conditions under which these comparative statics are the same as for the case of $\sigma^2 > 0$ and very large π .

Sheshinski and Weiss (1977) study a menu cost model similar to the deterministic case in our basic setup. The firm's problem is to decide when to change prices and by how much when aggregate prices grow at the rate π . In Sheshinski and Weiss's (1977) model the time elapsed between adjustments is simply a constant, which we denote by $\mathcal{T}(\pi)$. Sheshinski

⁸Furthermore $p^*(z) = z$ is assumed to follow a random walk with no drift, so that $\bar{z} = \infty$, $a(z) = 0$ and $b(z) = 1$.

and Weiss (1977) find sufficient conditions so that the time between adjustments decreases with the inflation rate (see their Proposition 2), and several authors have further refined the characterization by concentrating on the case where the fixed cost c is small. Let $p^* = \arg \max_p F(p, 0)$ be the log price of the static monopolist maximization profit, where $z = 0$ is a normalization of the shifter parameter which stays constant. In the deterministic set-up the optimal policy for $\pi > 0$ is to let the log price reach a value s at which time it adjusts to S , where $s < p^* < S$. The time between adjustments is then $\mathcal{T}(\pi) = (S - s)/\pi$. Furthermore, we highlight another implication obtained in the Sheshinski and Weiss's (1977) model, i.e., the set-up with $\sigma^2 = 0$. The distribution of the log relative price is uniform in the interval $[s, S]$. Thus the standard deviation of the log relative prices in this economy, denoted by $\bar{\sigma}$, is given by $\bar{\sigma} = \sqrt{1/12}(S - s)$. As established in Proposition 1 in Sheshinski and Weiss (1977), the range of prices $S - s$ is increasing in the inflation rate π . Obviously the elasticities of λ_a and of $\bar{\sigma}$ with respect to π are related since $S - s = \pi\mathcal{T}$ and $\lambda_a = 1/\mathcal{T}$.

Lemma 6 *Assume that $\sigma^2 = 0$ and $\pi > 0$. Then it follows immediately that $\lambda_a^-(\pi, \sigma^2) = 0$ and that $\Delta_p^+(\pi, \sigma^2) = S - s$. Furthermore assume that $F(\cdot, 0)$ is three times differentiable, then*

$$\lim_{c \rightarrow 0} \frac{\partial \lambda_a}{\partial \pi} \frac{\pi}{\lambda_a} = \frac{2}{3} \quad \text{and} \quad (3.9a)$$

$$\lim_{c \rightarrow 0} \frac{\partial \bar{\sigma}}{\partial \pi} \frac{\pi}{\bar{\sigma}} = \frac{1}{3} \quad (3.9b)$$

Proof. See online appendix C.1.2.

The lemma establishes that in the deterministic case when menu costs c are small, there are no price decreases, and the magnitude of price increases, $S - s$, increases with inflation at a rate of $1/3$. Also, as inflation increases the time between consecutive price changes shrinks and the frequency of price adjustment increases with an elasticity of $2/3$.

Next, Lemma 7 analyzes the conditions under which the limiting values of the elasticities in Lemma 6 for the Sheshinski and Weiss (1977) model are the same as for the case with idiosyncratic costs, $\sigma > 0$, and very large π . Lemma 7 establishes that when the idiosyncratic shocks z are very persistent and interest rates and menu costs are very small, the frequency of price adjustment is homogeneous of degree one in (π, σ^2) so that it can be written as a function of the ratio σ^2/π .

For the next results we write the frequency of price adjustment as a function of the rate of inflation, π , the variance of the idiosyncratic shock, σ^2 , the discount factor, r , and the inverse of the menu cost, $1/c$, that is, $\lambda_a(\pi, \sigma^2, r, \frac{1}{c})$. We also write the policy rules as functions of the parameters for each z ; $\Psi(\pi, \sigma^2, r, \frac{1}{c}; z) = \left[\underline{\psi}(z; \pi, \sigma^2, r, \frac{1}{c}), \bar{\psi}(z; \pi, \sigma^2, r, \frac{1}{c}), \hat{\psi}(z; \pi, \sigma^2, r, \frac{1}{c}) \right]$ and the expected price change functions as $\Delta_p^+(\pi, \sigma^2, r, \frac{1}{c})$ and $\Delta_p^-(\pi, \sigma^2, r, \frac{1}{c})$.

Lemma 7 *Let $a(z) = 0$ for all z . Then $\lambda_a(\pi, \sigma^2, r, \frac{1}{c})$ is homogenous of degree one and the policy functions $\Psi(\pi, \sigma^2, r, \frac{1}{c}; z)$ are homogeneous of degree zero in all the parameters for a fixed z . Therefore,*

$$\begin{aligned} \lim_{\pi \rightarrow \infty} \left[\lim_{r \downarrow 0, c \downarrow 0} \frac{\partial \lambda_a(\pi, \sigma^2, r, \frac{1}{c})}{\partial \pi} \frac{\pi}{\lambda_a(\pi, \sigma^2, r, \frac{1}{c})} \right]_{\sigma > 0} &= & (3.10a) \\ \lim_{\sigma \rightarrow 0} \left[\lim_{r \downarrow 0, c \downarrow 0} \frac{\partial \lambda_a(\pi, \sigma^2, r, \frac{1}{c})}{\partial \pi} \frac{\pi}{\lambda_a(\pi, \sigma^2, r, \frac{1}{c})} \right]_{\pi > 0}, & \end{aligned}$$

for each z ,

$$\lim_{\pi \rightarrow \infty} \left[\lim_{r \downarrow 0, c \downarrow 0} \Psi(\pi, \sigma^2, r, \frac{1}{c}; z) \right]_{\sigma > 0} = \lim_{\sigma \rightarrow 0} \left[\lim_{r \downarrow 0, c \downarrow 0} \Psi(\pi, \sigma^2, r, \frac{1}{c}; z) \right]_{\pi > 0} \quad (3.10b)$$

and

$$\lim_{\pi \rightarrow \infty} \left[\lim_{r \downarrow 0, c \downarrow 0} \Delta_p^+(\pi, \sigma^2, r, \frac{1}{c}) \right]_{\sigma > 0} = \lim_{\sigma \rightarrow 0} \left[\lim_{r \downarrow 0, c \downarrow 0} \Delta_p^+(\pi, \sigma^2, r, \frac{1}{c}) \right]_{\pi > 0}. \quad (3.10c)$$

Proof. See online appendix C.1.3.

The intuition underlying Lemma 7's proof is that multiplying r, π, σ^2 and the functions $a(\cdot)$ and $F(\cdot)$ in the firm's problem— equation (3.2)—by a constant $k > 0$ is akin to changing the units in which we measure time. Moreover, the objective function in the right hand side of equation (3.2) is homogeneous of degree one in $F(\cdot)$ and c and, hence, the policy function is the same whether we multiply $F(\cdot)$ by k or divide c by it. Thus, if the function $a(\cdot)$ equals zero, then λ_a is homogeneous of degree one in $(\pi, \sigma^2, r, \frac{1}{c})$. Likewise, $\lambda_a(\pi, \sigma^2)$ is homogeneous of degree one in (π, σ^2) when menu costs are small, $c \downarrow 0$, the interest rate is very low, $r \downarrow 0$, and shocks are very persistent, $a(\cdot) = 0$. The interpretation of r going to zero is that instead of maximizing the expected discounted profit, the firm is maximizing the expected average profit, a case frequently analyzed in stopping time problems. If the idiosyncratic shock process has $a(\cdot) = 0$ and $b(\cdot)$ is bounded, then z is a Martingale, i.e., the

shifter has permanent shocks.

Lemma 7 extends the result on the elasticity of the frequency of price adjustment with respect to inflation of equation (3.9a) in Lemma 6 to the case with $\sigma > 0$ and with an arbitrarily large π . This lemma requires that the shifter z has only permanent shocks, i.e., that $a(z) = 0$. A similar argument does not apply to the cross sectional dispersion of relative prices $\bar{\sigma}$ since setting $a(z) = 0$ implies that there is no invariant distribution of z , and hence no invariant distribution of relative prices.

Using Lemma 6 and Lemma 7 we obtain the following result for the high inflation case:

Proposition 4 *Assume that $a(z) = 0$ for all z and F is three times differentiable. Consider two firms with $\sigma_1, \sigma_2 > 0$. Then,*

$$\lim_{\pi \rightarrow \infty} \left[\frac{\lim_{r \downarrow 0, c \downarrow 0} \lambda_a(\pi, \sigma_1^2)}{\lim_{r \downarrow 0, c \downarrow 0} \lambda_a(\pi, \sigma_2^2)} \right] = 1 \quad (3.11a)$$

$$\lim_{\pi \rightarrow \infty} \left[\lim_{r \downarrow 0, c \downarrow 0} \frac{\partial \lambda_a(\pi, \sigma_i^2)}{\partial \pi} \frac{\pi}{\lambda_a(\pi, \sigma_i^2)} \right] = \frac{2}{3} \quad \text{for } i = 1, 2 \quad \text{and} \quad (3.11b)$$

$$\lim_{\pi \rightarrow \infty} \left[\lim_{r \downarrow 0, c \downarrow 0} \frac{\partial \Delta_p^+(\pi, \sigma_i^2)}{\partial \pi} \frac{\pi}{\Delta_p^+(\pi, \sigma_i^2)} \right] = \frac{1}{3} \quad \text{for } i = 1, 2. \quad (3.11c)$$

Proposition 4 contains strong predictions about the limiting behavior of the frequency of price adjustment as inflation becomes large. The first part is a direct consequence of Lemma 7. It implies that if we think that different industries have systematically different idiosyncratic shocks, we would expect the variance of these shocks to differ across industries and, hence, the frequency of price adjustment to be different across industries when inflation is low. Equation (3.11a) implies that differences in the frequency of price adjustment observed with low inflation should wash away as inflation becomes large. This is illustrated in the numerical example in the next section (see Figure 3.2) and verified in the data (see Section 3.5.2). The intuition is that, when inflation is low, the main driver of idiosyncratic nominal price changes are idiosyncratic shocks and, when inflation is high, the main driver of price changes is the growth of aggregate prices. The second part of Proposition 4 is a sharp prediction about the rate at which firms change the frequency of price adjustment when inflation grows. It states that this elasticity should be $2/3$ in the limit when $\pi \rightarrow \infty$. Finally, equation (3.11c) states that the elasticity of price increases with respect to inflation converges to $1/3$ as $\pi \rightarrow \infty$. It follows from the fact that Δ_p^+ is $S - s = \pi/\lambda_a$ when $\sigma = 0$.

The results of Proposition 4 apply to a wider set of models such as those mentioned in the comments to Proposition 3.

3.2.1.3 DECOMPOSITION OF CHANGES IN THE RATE OF INFLATION

This section shows how steady state changes in the rate of inflation can be decomposed into changes in the extensive and in the intensive margins of price adjustment—i.e., changes in inflation accounted for by the frequency of price changes and changes in inflation accounted for by the size of price changes conditional on a price change taking place. Our main result is that for low inflation the extensive margin accounts for ninety percent of changes in inflation while for large inflation it accounts for two thirds of inflation changes.

We decompose the inflation rate as the difference between the product of the frequency of price increases times the average size of price increases, and the product of the frequency of price decreases times the average size of price decreases. Formally:

$$\pi = \lambda_a^+ \Delta_p^+ - \lambda_a^- \Delta_p^-$$

Totally differentiating the previous expression with respect to the inflation rate, we decompose changes in the inflation rate into those accounted for by changes in the frequency, denoted by δ , and those accounted for by changes in the average size, denoted by $1 - \delta$:

$$1 = \underbrace{\frac{\partial \lambda_a^+}{\partial \pi} \Delta_p^+ - \frac{\partial \lambda_a^-}{\partial \pi} \Delta_p^-}_{\delta(\pi)} + \underbrace{\frac{\partial \Delta_p^+}{\partial \pi} \lambda_a^+ - \frac{\partial \Delta_p^-}{\partial \pi} \lambda_a^-}_{1-\delta(\pi)}$$

We derive the decomposition of inflation for $\pi = 0$ and for $\pi \rightarrow \infty$, in the special case with a quadratic profit function,⁹ no discounting, and where z represents the (log of the) product cost and follows a drift-less continuous time random walk.

Proposition 5 *Assume that $a(z) = 0$ and $b(z) = 1$, $\sigma > 0$ and $F(p - \bar{p}, z) = B (p - z)^2$. Consider the limit of the policy functions $\lim_{r \rightarrow 0} \Psi(\pi, \sigma^2, r, \frac{1}{c}, B; z) = \Psi(\pi, \sigma^2, \frac{1}{c}, B; z)$. Then,*

$$\delta(0) = \frac{9}{10} \text{ and } \lim_{\pi \rightarrow \infty} \delta(\pi) = \frac{2}{3}$$

⁹Equivalently, we can write the result for small a fixed cost c , so that prices are close to the profit maximizing value, and thus a second order expansion of the profit function is accurate.

Proof. See online appendix C.1.4 ■

Notably, while Proposition 5 states that ninety percent of changes in inflation around zero inflation are accounted for by the extensive margin of price adjustment, Proposition 3 states that the frequency of price changes is insensitive to inflation near zero inflation. To gain insight into the interplay between the two propositions observe that for zero inflation there is no trend in relative prices and symmetry implies that the frequency of price increases and decreases are the same, $\lambda_a^+ = \lambda_a^-$. Also, that $\Delta_p^+ = \Delta_p^-$. Proposition 3 implies that $\frac{\partial \lambda_a^+}{\partial \pi} = -\frac{\partial \lambda_a^-}{\partial \pi}$ so the extensive margin at zero is $\delta(0) = 2\Delta_p^+ \frac{\partial \lambda_a^+}{\partial \pi}$. Since inflation introduces a negative trend in relative prices, it induces them to hit more often the lower limit of the inaction set, prompting more price increases and less price decreases. The characterization of optimal policies in the proof of Proposition 5 shows that these changes in the frequency of price increases and decreases account for *ninety percent* of changes in the rate of inflation at $\pi = 0$. A similar argument holds for the decomposition of the change of inflation in a mild deflation.

The second part of Proposition 5 restates the results in Proposition 4 under slightly different assumptions. The results in Proposition 4 hold in the limit when $c \rightarrow 0$ while Proposition 5 holds for a specific functional form of the profit function F . Not surprisingly, for small adjustment costs the quadratic profit function in Proposition 5 is a good approximation of F .

The key technical insight in the proof of Proposition 5 is that $\lim_{r \rightarrow 0} rV(x, r)$ is finite and independent of x . This allows us to obtain an analytical solution of the value function and to characterize optimal policies as simple functions of the ratio σ^2/π .

Finally, the following corollary to Proposition 5 presents a sharp prediction about how changes in inflation affect the frequency of price increases and decreases when inflation is low.

Corollary 1 *Around $\pi = 0$, the difference between the frequency of price increases and decreases rises with inflation. Formally,*

$$\frac{\partial(\lambda_a^+ - \lambda_a^-)}{\partial \pi} \Big|_{\pi=0} = \frac{\delta(0)}{\Delta_p^+ \Big|_{\pi=0}} = \frac{\frac{9}{10}}{\Delta_p^+ \Big|_{\pi=0}} > 0$$

Taken together, the results at low inflation in Proposition 5 and its corollary imply that

inflation rises when inflation is low mostly because the frequency of price increases rises and that of price decreases falls (the extensive margin), as opposed to the size of price increases rising and that of price decreases falling (the intensive margin).

General Remarks

We conclude this section with a few remarks on the applicability of these comparative static results to the time series variation in our dataset. The propositions in this section were obtained under the assumption that inflation is to remain *constant* at the rate π , and that the frequency of price changes is computed under the invariant distribution. Thus, strictly speaking, our propositions are *not* predictions for time series variation but comparative static results.

We give three comments on this respect. First, this should be less of a concern for very high inflation, since the model becomes close to static, i.e., firms change prices very often and thus the adjustment to the invariant distribution happens very fast. Second, when we analyze the Argentinean data we correlate the current frequency of price changes with an average of the current and future inflation rates. We experiment with different definition of these averages and find that the estimates of the elasticities in the first two propositions of this section are not sensitive to this. Moreover, with Argentina's experience in mind, Beraja (2013) studies the transition dynamics in a menu cost model where agents anticipate a disinflation in the future and performs the same comparative statics with artificial data generated from such model. He finds that the theoretical results in this section are robust to conducting the analysis in a non-stationary economy during a disinflation process calibrated to the Argentine economy. Third, in Section 3.2.2 we numerically solve a standard version of the menu cost model for reasonable parameter values for menu cost c and discount rate r , which are positive but small, and for finite but large inflation rates π , of the order that are observed in Argentina. We find that the propositions in this section, which use limit values for c, r and π , accurately predict the behavior of the statistics of interest computed in the calibrated model.

3.2.2 ILLUSTRATING THE THEORY WITH AN EXAMPLE

In this section we specify a version of the firm's problem studied in Section 3.2.1 to illustrate the theory. We characterize the solution of the model analytically and numerically and

show how changes in the rate of inflation affect optimal pricing rules, the frequency of price changes, and the size of price adjustments.

We compute this example to verify the robustness of the analytical predictions obtained so far. In Section 3.2.1 we obtained sharp analytical results under a variety of simplifying assumption such as limit values of parameters (e.g., vanishing menu cost c and or discount rate r), or the shape of profit functions (e.g., symmetry of F). Also our analytical results were obtained at two extreme values of inflation. In this section we check the robustness of the simplifying assumptions by computing a version of the model away from the limit cases, and also consider values of inflation within Argentina's experience.

The example assumes a constant elasticity of demand, a constant returns to scale production technology, idiosyncratic shocks to marginal cost that are permanent, an exponentially distributed product life and a cost of changing prices that is proportional to current profits (but independent of the size of the price change). This version of the Golosov and Lucas (2007) model is identical to the one in Kehoe and Midrigan (2010).¹⁰

We assume that period profits are given by a demand with constant elasticity η and with a constant return to scale technology with marginal cost given by e^z , so that

$$F(p, z) = e^{-\eta p} (e^p - e^z) ,$$

where $p - z$ is the log of the gross markup (the net markup).

The shocks on the log of the cost are permanent in the sense that

$$dz = \mu_z dt + \sigma dW - z dN.$$

where N is the counter of a Poisson process with constant arrival rate per unit of time ρ . We interpret this as products dying with Poisson arrival rate $\rho > 0$ per unit of time, at which time they are replaced by a new one which starts with $z = 0$ and must set its initial price. A positive value of μ_z can be interpreted as a vintage effect, i.e., the technology for new products grows at rate μ_z . The disappearance of products assures the existence of an ergodic distribution of relative prices and we view this device as realistic given the rate at which product are substituted in most datasets.

The menu cost, $\zeta(z)$, is assumed to be $\zeta(z) = c F(p^*(z), z)$ for some constant $c > 0$.

¹⁰We zero out the transitory shock that gives rise to sales in Kehoe and Midrigan (2010) and set it in continuous time.

Here, $p^*(z) = z + m$ where $m = \log(\frac{\eta}{\eta-1})$ is the log of the gross optimal static markup. Thus $F(p^*(z), z) = e^{(1-\eta)z} \left(\frac{\eta-1}{\eta}\right)^\eta \frac{1}{\eta-1}$. Note that $F(p^*(z), z)$ is decreasing and strictly convex on z for $\eta > 1$.

We will assume that $\eta > 1$ so that the static monopolist problem has a solution, and that

$$r + \rho \geq (1 - \eta) \left[\mu_z + (1 - \eta) \frac{\sigma^2}{2} \right] \quad (3.12)$$

so that the profits of the problem with zero fixed cost, $\hat{c} = 0$ are finite. Since for $\eta > 1$, period profits are decreasing and convex on z , and hence discounted expected profits are finite if the discount rate $r + \rho$ is high enough, or if the cost increases at a high enough rate μ_z , or if $\sigma^2/2$ is low enough.¹¹

The firm's optimal pricing policy for the $\sigma > 0$ case can be characterized in terms of three constants $X \equiv (\underline{x}, \bar{x}, \hat{x})$. This is due to the combination of assumptions of constant elasticity of demand, constant returns to scale and permanent shocks to cost while the product last. In this example, the policy function takes the simple form $\Psi(z) = X + z$. Letting $x = p - z$ be the log of the real gross mark-up, which we refer to as the net markup, we write the inaction set as $\mathcal{I} = \{x : \underline{x} < x < \bar{x}\}$. It is optimal to keep the price unchanged when the net markup x is in the interval (\underline{x}, \bar{x}) . When prices are not changed, the real markup evolves according to $dx = -(\mu_z + \pi)dt + \sigma dW$. When the real markup hits either of the two thresholds, prices are adjusted so that the real markup is \hat{x} and thus the optimal return relative price is $\hat{\psi}(z) = \hat{x} + z$. For the case where $\sigma = 0$ and $\pi + \mu_z > 0$, we obtain a version of Sheshinski and Weiss's (1977) model, and the optimal policy can be characterized simply by two thresholds $s \equiv \underline{x} < \hat{x} \equiv S$.¹²

In online appendix C.2 we present several propositions with an analytical characterization of the solution of this model. A novel contribution of this paper is to derive a system of three equations in three unknowns for X , as well as the explicit solution to the value function, which depends on the parameters $\Theta \equiv (\pi, \mu_z, \sigma^2, \rho, r, \eta, c)$. We also derive an explicit solution for the expected number of adjustment per unit of time λ_a given a policy X and parameters $(\pi, \mu_z, \sigma^2, \rho)$ and we characterize the density g for the invariant distribution of (p, z) implied

¹¹In this case the expected discounted value of profits, starting with z_0 is given by $\mathbb{E}_0 \int_0^\infty e^{-(r+\rho)t} F(p^*(z_t), z_t) dt = \frac{(\eta-1)^{\eta-1} e^{(1-\eta)z_0}}{\eta^\eta} \int_0^\infty e^{[-(r+\rho)+(1-\eta)\mu_z+(1-\eta)^2\frac{\sigma^2}{2}]t} dt$.

¹²In this case $\lambda_a = \rho / \left[1 - \exp\left(-\frac{\rho}{\pi+\mu_z}(\hat{x}-\underline{x})\right) \right] = (\pi + \mu_z) / (\hat{x} - \underline{x}) + o(\rho / (\pi + \mu_z))$, so that it coincides with the expression used for $\sigma = 0$ if ρ is small relative to π .

by the policy X and the parameters $(\pi, \mu_z, \sigma^2, \rho, \eta)$.

The remainder of this section contains several figures that describe numerically how changes in the rate of inflation affect the optimal pricing rules (the triplet X), the frequency of price changes, and the size of price adjustments. For the numerical examples we follow Kehoe and Midrigan (2010) and use $\eta = 3$, which implies a very large markup, but it is roughly inline with marketing/IO estimates of demand elasticities. We set $\rho = 0.1$ so products have a lifetime of 10 years, and $r = 0.06$ so yearly interest rates are 6%. We let $c = 0.002$ so that adjustment cost is 20 basis point of yearly frictionless profits. We let $\mu_z = 0.02$, i.e., a 2 percent per year increase in cost (or a 2% increase in vintage productivity). We consider three values for $\sigma \in \{0, 0.15, 0.20\}$, the first corresponds to Sheshinski and Weiss's (1977) model, and the others are 15% and 20% standard deviation in the change in marginal cost, at annual rates. The values of $c/(\eta(\eta - 1))$ and $\sigma = 0.15$ were jointly chosen so that at zero inflation the model matches both the average number of price changes $\lambda_a = 2.7$ and the average size of price changes of $\Delta_p = 0.10$, roughly the values corresponding to zero inflation in our data set.¹³


Figure 3.1 illustrates how the optimal threshold policies vary with inflation for two cases, $\sigma = 0.15$ and $\sigma = 0$. The blue lines depict Golosov and Lucas's (2007) case with $\sigma > 0$. The dashed center line is the optimal return mark-up and the outer lines are the boundaries of the inaction set. With no inflation, the mark-up will drift away from the starting optimal markup, driven by the idiosyncratic shock. The firm will keep its nominal price fixed as long as it does not hit the boundaries. Once the mark-up hits either boundary, \underline{x} or \bar{x} , the firm resets the price and the mark-up returns to \hat{x} . The red lines depict the optimal thresholds for Sheshinski and Weiss's (1977) case with $\sigma = 0$. Mark-ups always fall when there are no idiosyncratic shocks after the firm resets its nominal price. Hence, the upper limit of the inaction set becomes irrelevant. The firm resets its nominal price to $\hat{x} + z$ when the mark-up hits the lower bound \underline{x} and waits for it to fall again.

Figure 3.1 shows several properties of the model.¹⁴ At very low inflation rates, and when $\sigma > 0$, the thresholds are symmetric, i.e., the distance between \underline{x} and \hat{x} is the same as

¹³Using results in Alvarez et al. (2011), when z is a random walk and the fixed cost is proportional to profits, there are mappings from observations to parameters $\lambda_a = \left[\frac{\sigma^2 (1/2) \eta (\eta - 1)}{c} \right]^{\frac{1}{2}}$ and $\Delta_p = \left[\sigma^2 \frac{c}{(1/2) \eta (\eta - 1)} \right]^{\frac{1}{4}} \sqrt{\frac{2}{3.14159 \dots}}$

¹⁴This example does not exactly satisfies all the assumptions of the model in Section 3.2.1 since the profit function F derived from a constant elasticity demand is not symmetric. Yet, for small cost c the terms in the quadratic expansion, which are symmetric by construction, should provide an accurate approximation.


Figure 3.1: Optimal thresholds for different inflation rates



the distance between \bar{x} and \hat{x} . This symmetry implies that the size of price increases is equal to the size of price decreases, $\Delta_p^+(0, \sigma) = \Delta_p^-(0, \sigma)$, and that the frequency of price increases is equal to the one of price decreases, $\lambda_a^+(0, \sigma) = \lambda_a^-(0, \sigma)$. These illustrate the results obtained in part (iii) of Proposition 3. At very high inflation rates, the models with $\sigma > 0$ and with $\sigma = 0$ are equivalent in the sense that the critical values \underline{x} and \hat{x} in Golosov and Lucas's (2007) model converge to the Ss bands in Sheshinski and Weiss's (1977) model as established in equation (3.10b) in Lemma 7. As a result, the magnitude of price changes in the two models is the same as in equation (3.10c) in Lemma 7, $\Delta_p^+(\pi, 0) = S - s = \hat{x} - \underline{x}$. For rates of inflation above 250 percent per year, Figure 3.1 also shows that the elasticity of $\Delta_p^+(\pi, 0)$ with respect to inflation is close to 1/3—equation (3.11c) in Proposition 4.

Figure 3.2 displays the frequency of price increases λ_a^+ , together with the frequency of all adjustments λ_a , for two values of the cost volatility σ . There are several interesting

Figure 3.2: Frequency of Price Adjustments λ_a and of Price Increases λ_a^+



observations about this figure. First, the frequency λ_a is insensitive to inflation in the neighborhood of zero inflation as established in part (i) of Proposition 3. Second, the length of the inflation interval around $\pi = 0$ where λ_a is approximately constant increases with σ —see the discussion in footnote 6. Third, the last part of Proposition 3 predicts that the frequency of price increases and price decreases is the same when $\pi = 0$. The figure shows that for low inflation the frequency of price increases is about half of the frequency of price changes, indicating that half the price changes are increases and half are decreases. Fourth, for values of inflation above 250% per year, the frequency of price changes λ_a for different values of σ are approximately the same, consistent with the limiting results in equation (3.11a) of Proposition 4. Fifth, since the graph is in log scale, it is clear that the common slope is approximately constant for large inflation, and close to $2/3$ as established in equation (3.11b) in Proposition 4. Finally, as inflation becomes large all price adjustments

are price increases—as it can be seen by the fact that λ_a^+ converges to λ_a for each value of σ .

We now turn to the distribution of relative prices (which is derived in online appendix C.2). Let $g(p, z; \pi, \sigma)$ be the invariant joint distribution of p and z and let the marginal distribution of p be $h(p) = \int_{p-\bar{x}}^{p-\underline{x}} g(p, z; \pi, \sigma) dz$. We are mostly interested in the *marginal*, or unconditional distribution of relative prices $h(p)$, as opposed to the conditional distribution $g(p, z)$ since the marginal (or unconditional) distribution is the object we can hope to measure in actual data because z is not readily observable.


Our focus is the analysis of the standard deviation of relative prices, especially its elasticity with respect to inflation. There are two sources of dispersion for relative prices: the idiosyncratic cost shocks and the asynchronous price adjustments to inflation. It is helpful to decompose the unconditional variance of relative prices $\bar{\sigma}(p; \pi, \cdot)$ for a given inflation rate π as follows:

$$\sigma(p; \pi, \cdot) = \mathbb{E}[\text{Var}(p|z; \pi, \cdot)] + \text{Var}[\mathbb{E}(p|z; \pi, \cdot)] \quad (3.13)$$

We first explain the behavior of the first term in equation (3.13). The results in online appendix C.2 show that the distribution $g(\cdot, z)$ has the same shape for any value of z . This implies that the first term mostly captures dispersion in relative prices coming from asynchronous price adjustments to inflation. In Figure 3.3, we fix $z = 0$ and plot the standard deviation of p corresponding to the distribution $g(\cdot, 0)$ and $\sigma > 0$ for different values of π . The figure shows that the elasticity of the standard deviation of relative prices conditional on z with respect to inflation is approximately zero for $\pi = 0$ (as in Proposition 3) and it is approximately one-third for large π (as in Lemma 6). Moreover, the variance of relative prices is independent of σ for large π (as in Lemma 7). Next, we discuss the second term in equation (3.13). This source of dispersion is mostly determined by price changes due to firm idiosyncratic shocks z . To see this, note that the variance of the average price is equal to the cross sectional dispersion of z for all values of π when menu costs are zero.

In sum, we observe that relative price dispersion mostly results from idiosyncratic firm shocks when inflation is low—neither of the two terms of the unconditional variance changes with inflation—whereas, for large enough inflation rates, relative price dispersion mostly results from asynchronous price adjustments to inflation. In our numerical examples, however, it takes much higher inflation rates than the ones observed in the peak months in Argentina for this to happen (see Figure C.2 in online appendix C.2). While the first term becomes

Figure 3.3: Standard deviation of log relative prices, conditional on $z = 0$.



insensitive to changes in the variance of idiosyncratic firm shocks, as seen in Figure 3.3, the second term does not for this range of inflation.

We briefly comment on the difference between the behaviour of λ_a and of $\bar{\sigma}$ as functions of π . In this model the value for the frequency of price changes $\lambda_a(\pi, \sigma^2)$ converges to $\lambda_a(\pi, 0)$ as inflation increases much faster than $\bar{\sigma}(\pi, \sigma^2)$ does. The reason is that given the permanent nature of the shocks to a product cost, z , the expected time until the next adjustment \mathcal{T} is only a function of $x = p - z$. Recall that $\lambda_a = 1/\mathcal{T}$ so that the cross sectional distribution of z is essentially irrelevant for the frequency of adjustment. Instead, the standard deviation of relative prices $\bar{\sigma}$ depends on the cross sectional distribution of z on a crucial way. Indeed, if the idiosyncratic shocks were completely permanent, there will be no invariant distribution of relative prices. In our example, the reason why there is an invariant distribution is that $\rho > 0$, so products are returned to $z = 0$ at exponentially distributed times.

3.3 DESCRIPTION OF THE DATASET

Our dataset contains 8,618,345 price quotes underlying the consumer price index for the Buenos Metropolitan Area in the period December 1988 to September 1997. Each price quote represents an *item*, i.e., a good or service of a determined brand sold in a specific outlet in a specific period of time. Goods¹⁵ and outlets are chosen to be representative of consumer expenditure in the 1986 consumer expenditure survey¹⁶. Price quotes are for 506 goods that account for about 84 percent of household expenditures.

Goods are divided into two groups: homogeneous and differentiated goods. Differentiated goods represent 50.5 percent of the expenditure in our sample while homogeneous goods account for the remaining 49.5 percent.¹⁷ Prices are collected every two weeks for all homogeneous goods and for those differentiated goods sold in super-market chains; and are gathered every month for the rest of the differentiated goods. The dataset contains 233 prices collected every two weeks and 302 prices collected every month. 29 of each of these goods are gathered both monthly and biweekly.¹⁸

An important feature of the dataset is the rich cross section of outlets where prices are recorded at each point in time. Over the whole sample there are 11,659 outlets. Roughly around 3200 outlets per month for homogeneous goods and about the same number for differentiated goods. On average, across the 9 years, there are 166 outlets per good (81 outlets per product collected monthly and 265 per good collected bimonthly). Online appendix C.3 contains further information on data collection and on the classification of goods.

We exclude from the sample price quotes for baskets of goods, rents and fuel prices. Baskets correspond to around 9.91% of total expenditure and are excluded because their prices are gathered for any good in a basket, i.e., if one good is not available, it is substituted by any another in the basket. Examples are medicines and cigarettes. Rents are sampled monthly for a fixed set of representative properties. Reported prices are for the average of the sampled properties and include what is paid on that month, as opposed to what is paid for a new contract. Rents represent 2.33 percent of household expenditure. Fuel prices

¹⁵To simplify the exposition, when it is clear, we use goods to refer to either goods or services.

¹⁶Encuesta Nacional de Gasto de los Hogares


¹⁷Examples of homogeneous goods are: barley bread, chicken, lettuce, etc. Examples of differentiated goods are: moccasin shoes, utilities, tourism, and professional services.

¹⁸The outlets are divided into 20 waves, corresponding to the 20 working days of the month. Each outlet is visited roughly in the same working day every 10 working days, in the case of homogeneous goods and differentiated goods gathered at super-markets. The dataset includes the particular day when each price is gathered.

account for 4 percent of total expenditure and we exclude them because they were gathered in a separate database that we do not have access to.

The dataset has some missing observations and flags for stock-outs, price substitutions and sales. We treat stock-outs (10.5 percent of observations) and price quotes with no recorded information (2.25 percent of observations) as missing observations. The statistical agency substitutes the price quote of an item for a similar item, typically, when the good is either discontinued by the producer or not sold any longer by an outlet. Using this definition, across the 9 years of our dataset, we have an average of 2.39 percent of price quotes that have been substituted. The data set contains an indicator of whether an item was on sale or not. Around 5 percent of items have a sale flag. This is small compared with the 11 percent frequency of sales reported by Klenow and Kryvtsov (2008) for the US. 70 percent of the sales correspond to homogeneous items (this is similar to Klenow and Kryvtsov (2008). They report that sales are more frequent for food items). The time series data for the number of outlets per good and for the frequencies of missing observations, substitutions and sales are depicted in Figure 3.4.

Figure 3.4: Number Outlets per Good, and Frequencies of Substitution and Sales



Note: For the homogeneous goods during a month we count a sale or substitution if there was one such event in any of the two fifteen days subperiods. Missing includes stock-outs.

3.4 ESTIMATING THE FREQUENCY OF PRICE CHANGES

We extend the methodology of Klenow and Kryvtsov (2008) to the case of time varying frequencies of price changes. We assume a constant probability of a price change per unit of time (a month for differentiated goods and two weeks for homogeneous goods) so that the arrival rate of a price change follows a Poisson process. In this case, the maximum likelihood estimator of the frequency of price changes is

$$\lambda_t = -\ln(1 - \text{fraction of outlets that changed price between } t \text{ and } t-1). \quad (3.14)$$


The fraction of outlets that changed price between periods can be calculated for individual goods or for the aggregate by pooling the data for all outlets and all goods together. In this computation, we drop observations with missing price quotes. This *simple estimator* just counts the fraction of price changes in a period of time, and transform it into a per unit of time rate, λ . We refer to λ as the “instantaneous” frequency of price changes.

Later, we perform robustness checks by using different methods of aggregation across goods, by considering different treatments for sales, substitutions and missing observations, and by dropping the assumption that price changes follow a Poisson process.

Figure 3.5 plots the monthly time series of the simple pooled estimator of λ and the expected inflation rate. It assumes that all homogeneous and all differentiated goods have the same frequency of price changes and estimates this aggregate frequency by using the simple pooled estimator for the homogeneous and for the differentiated goods. The bi-weekly estimates of the homogeneous goods are aggregated to a monthly frequency¹⁹, and the plot shows the weighted average of these two estimators, using the share of household expenditures as weights. Finally, the expected inflation is computed as the average inflation rate $1/\hat{\lambda}_t$ periods ahead. We observe that the two variables are correlated. For instance, during the mid-1989 hyperinflation, the implied expected duration of a price spell is close to one week; while after 1993 the implied expected duration is close to half a year.

¹⁹The monthly frequency is the sum of the bi-weekly frequencies of each month.

Figure 3.5: Estimated Frequency of Price Changes λ and Expected Inflation



Note: Simple estimator of λ , $\hat{\lambda} = -\log(1 - f_t)$, where f_t is the fraction of outlets that changed price in period t . λ is estimated separately for homogeneous goods (bi-weekly sample) and for differentiated goods (monthly sample). Homogeneous goods frequencies are converted to monthly by adding the bi-weekly ones for each month pair. The aggregate number is obtained by averaging with the respective expenditure shares in the Argentine CPI. Inflation is the average of the log-difference of monthly prices multiplied by 1200 and weighted by expenditure shares. Expected inflation is the average inflation rate $1/\hat{\lambda}_t$ periods ahead.

3.5 ARGENTINA'S EVIDENCE ON MENU COST MODELS OF PRICE DYNAMICS

In Section 3.2 we uncovered several properties of menu cost models that can be contrasted with data. The presentation of the empirical results in this section is organized around those predictions. As a reminder, these are:

1. The elasticity of the frequency of price changes λ with respect to changes in the rate of inflation is zero at low inflation rates and it approximates two thirds as inflation becomes very large.²⁰

²⁰In models in which the menu cost is zero at some random times, so that they combine menu costs

2. The dispersion of the frequency of price changes across goods decreases with inflation. It is zero when inflation goes to infinity and the model converges to the Sheshinski and Weiss (1977) model with no idiosyncratic shocks.
3. Intensive and extensive margins of price increases and decreases.
 - (a) The frequency of price increases and of price decreases are similar at low inflation rates.
 - (b) The size of price increases and of price decreases are similar at low inflation rates.
 - (c) At low inflation rates, as inflation grows, the frequency of price changes remains constant while the frequency of price increases rises and the frequency of price decreases falls.
 - (d) For high inflation rates, the frequency of price increases converges to λ and the frequency of price decreases converges to zero.
 - (e) The size of price changes is an increasing function of the inflation rate.
4. The elasticity of the dispersion of prices across stores with respect to inflation is zero for low inflation rates and approaches one-third when inflation goes to infinity. This elasticity, however, may be smaller than one third for high inflation rates if idiosyncratic shocks are very persistent.

Next, we look at each of these predictions in the Argentinean data.

3.5.1 THE FREQUENCY OF PRICE CHANGES AND INFLATION

In this section, we report how the estimated frequency of price changes changes with inflation. We find that, as predicted by the menu cost model, the frequency of price changes is insensitive to inflation when inflation is low. Moreover, for high inflation rates, we find that the elasticity of the frequency of price changes with respect to inflation is between one-half and the theoretical two-thirds.

Figure 3.6 plots the frequency of price changes against the rate of inflation using log scale for both variables.²¹ On the right axis we indicate the implied instantaneous duration, and Calvo type price adjustments, this elasticity can be lower than 2/3 at high inflation. See, for example Nakamura and Steinsson (2010) and Alvarez et al. (2014b)

²¹See Section 3.2 for the caveats on these results and on the interpretation of contemporaneous correlations.

i.e., $1/\lambda$. In interpreting this figure, as well as other estimates presented below, it is worth noting that $1/\lambda_t$ is the expected duration of prices at time t if λ_t would remain constant in the future and provided that the probability of a price change is the same within the smallest period of observation (1 month for differentiated goods and 2 weeks for homogeneous goods).

Motivated by the theoretical considerations in Section 3.2, as well as the patterns evidenced in Figure 3.6, we fit (estimated via non-linear least squares) the following statistical model to the data:


$$\log \lambda = a + \epsilon \min \{\pi - \pi^c, 0\} + \nu(\min \{\pi - \pi^c, 0\})^2 + \gamma \max \{\log \pi - \log \pi^c, 0\}. \quad (3.15)$$

This model assumes that $\log \lambda$ is a quadratic function of inflation for inflation rates below the critical value, π^c and that $\log \lambda$ is a linear function of $\log \pi$ for inflation rates above π^c . In Figure 3.6 we observe that λ is insensitive to inflation at low inflation rates. Increasing inflation from 0 to 1 percent per year increases the frequency of price changes by only 0.04 percent. Moreover, the behavior of λ is symmetric around zero. For high inflation rates, the elasticity of λ with respect to inflation is captured by the parameter γ . We estimate γ to be at least $1/2$ but smaller than the theoretical limit $2/3$. Also, as predicted by the menu cost model, as inflation rises this elasticity becomes constant—i.e., the linear fit for $\log \lambda$ as a function of $\log \pi$ works well for high inflation rates. In this estimation, the critical value π^c in the statistical model, which has no theoretical interpretation, is of 14 percent per year.²² The expected duration of price spells for zero inflation is 4.5 months, which is consistent with international evidence as the next section shows.

The strong results in Figure 3.6 are surprising as the model applies to steady states in which the inflation rate has been at the same level for a long time. As a robustness check, we replicated Figure 3.6 using expected inflation instead of inflation (computed as the simple average of inflation rates $1/\lambda$ months ahead). The results are similar to those in Figure 3.6 but somewhat weaker.

²²The comparative static of the models discussed in Section 3.2 does not imply a kink as the one in equation (3.15), we merely use this specification because it is a low dimensional representation of interesting patterns in the data that provides a good fit and has properties at the extreme values that are consistent with our interpretation of the theory.

Figure 3.6: The Frequency of Price Changes (λ) and Expected Inflation.



Note: Simple estimator of λ , $\hat{\lambda} = -\log(1 - f_t)$, where f_t is the fraction of outlets that changed price in period t . λ is estimated separately for homogeneous goods (bi-weekly sample) and for differentiated goods (monthly sample). Homogeneous goods frequencies are converted to monthly by adding the bi-weekly ones for each month pair. The aggregate number is obtained by averaging with the respective expenditure shares in the Argentine CPI. Inflation is the average of the log-difference of monthly prices weighted by expenditure shares. The fitted line is $\log \lambda = a + \epsilon \min \{\pi - \pi^c, 0\} + \nu (\min \{\pi - \pi^c, 0\})^2 + \gamma \max \{\log \pi - \log \pi^c, 0\}$. The red squares represent negative expected inflation rates and the blue circles positive ones.

3.5.1.1 INTERNATIONAL EVIDENCE ON THE FREQUENCY OF PRICE CHANGES AND INFLATION

The previous section shows that certain aspects of price setting behavior in Argentina are consistent with the predictions of menu cost models. In particular, the elasticity of the frequency of price changes is close to zero at low inflation rates and close to two-thirds for high inflation rates. Here we show that Argentina's inflationary experience is of special

interest because it both spans and extends previous findings in the literature.

There are several studies that estimate the frequency of price changes for countries experiencing different inflation rates. Figure 3.7 provides a visual summary of these studies and compares them to ours by adding the international evidence to Figure 3.6. First, observe how the wide range of inflation rates covered by our sample makes this paper unique: none of the other papers covers inflation rates ranging from a mild deflation to 7.2 million percent per year (annualized rate of inflation in July 1989). This is what allows us to estimate the elasticity of the frequency of price changes with respect to inflation both at low and high inflation rates. In the other samples it is hard to test these hypothesis because of their limited inflation range. Second, we note how the patterns of the data for each country are consistent with the two predictions of the menu cost model. Third, we note that, in most cases, the level of the estimated frequency of price changes is similar to Argentina's. The similarity between our results and the existing literature is remarkable since the other studies involve different economies, different goods and different time periods. It is a strong indicator that our results are of general interest, as the theory suggests, and are not a special feature of Argentina.²³


The studies included in the figure are all the ones we could find covering a wide inflation range. For the low inflation range we included studies for the United States by Bils and Klenow (2004), Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008) and for the Euro Area by Álvarez et al. (2006). Our estimates of the frequency of price changes are consistent with all of them.²⁴ We have three data points for Israel corresponding to an inflation rate of 16 percent per year between 1991 and 1992 (Baharad and Eden (2004)), 64 percent per year between 1978 and 1979, and 120 percent per year between 1981 and 1982 (Lach and Tsiddon (1992)). The frequency of price changes for these three points is well aligned with the Argentine data. The same is true for the Norwegian data (Wulfsberg (2010)) that ranges from 0.5 percent to 14 percent per year. For Poland, Mexico and Brazil we were able to obtain monthly data for a wide range of inflation rates. The Polish sample ranges from 18 percent to 249 percent per year (Konieczny and Skrzypacz (2005)) and the

²³Table C.12 in online appendix C.6 provides a succinct comparison of the data sets used in these studies and of the inflationary environment in place in each case. The table shows that in addition to covering a wider range of inflation rates our data set is special due to its broad coverage that includes more than 500 goods representing 85 percent of Argentina's consumption expenditures.

²⁴There are other studies for low inflation countries, especially for the Euro area, but since they mostly yield estimates similar to those of Álvarez et al. (2006) we do not report them (see Álvarez et al. (2006) and Klenow and Malin (2011) for references to these studies).

Mexican one ranges from 3.5 percent to 45 percent per year (Gagnon (2009)). In both cases, the observations are aligned with the Argentina sample. The Brazilian data (Barros et al. (2009)) yields an elasticity of the frequency of price changes at high inflation that is consistent with ours. However, it yields a higher level of the frequency of price changes than ours and other studies.

Figure 3.7: The Frequency of Price Changes (λ) and Expected Inflation: International Evidence



Note: price changes per month for Argentina are the simple pooled estimator of λ . For the other cases we plot $-\log(1 - f)$, where f is the reported frequency of price changes in each study. The (λ, π) pairs for Argentina, Mexico and Brazil are estimated once a month and for the other countries once a year. Expected inflation is the average inflation $1/\lambda$ months ahead. Data for the Euro area is from Álvarez et al. (2006), for the US from Bils and Klenow (2004), Klenow and Kryvtsov (2008) and Nakamura and Steinsson (2008), for México from Gagnon (2009), for Israel from Baharad and Eden (2004), and Lach and Tsiddon (1992), for Poland from Konieczny and Skrzypacz (2005), for Brazil Barros et al. (2009), and for Norway from Wulfsberg (2010). Logarithmic scale for the both axis.

3.5.1.2 ROBUSTNESS

In online appendix C.4, we conduct many robustness exercises to evaluate the sensitivity of the main results in this section. The first set of exercises deals with recurrent issues when analyzing micro-price datasets, such as missing observations and price changes due to substitutions or sales. Secondly, we discuss issues of aggregation across products. Third, we address biases resulting from discrete sampling. Fourth, we present results using different measures of expected inflation. Finally, we address the possibility that the theoretical propositions which hold in the steady state are a poor description of the Argentine experience in the high inflation period leading to the stabilization plan in 1991 where agents are likely to have anticipated the strong disinflation that followed.

The conclusions of online appendix C.4 are twofold. First, at low inflation rates the empirical findings of this section go through intact. Second, at high inflation, we observe some quantitative but not qualitative differences. Most notably, depending on the estimator used to aggregate the data, the elasticity of the frequency of price changes can range from approximately one-half to the theoretical two-thirds.

3.5.2 INFLATION AND THE DISPERSION OF THE FREQUENCY OF PRICE CHANGES

This section reports how the dispersion of the frequency of price changes varies as inflation grows. Proposition 4 states that, under certain conditions, the firm's pricing behavior when inflation is high is independent of the variance of the idiosyncratic shocks. This implies that as inflation becomes higher it swamps the effect of idiosyncratic differences across firms that result in differences in the frequency with which they change prices. Figure 3.2 illustrates this point in the numerical example of our version of Golosov and Lucas (2007) model in Section 3.2.2.

We will interpret this result as applying not only to individual firms but to particular individual industries as well. In Table 3.1 we estimate λ for each narrowly defined industry (at a 5-digit level of aggregation),²⁵ calculate the implied average duration $\frac{1}{\lambda}$ and present two measures of the dispersion of λ s across such industries. We observe a significant decline in dispersion as inflation rises both across homogeneous and differentiated good industries.

²⁵Examples of 5 digit aggregation are citric fruits, soaps and detergents. See Table C.1 in online appendix C.3.

For homogeneous goods, the 90-10 percentile difference in λ s when inflation is above 500 percent per year is only 12 percent of the percentile difference at single-digit inflation. For differentiated goods, the 90-10 percentile difference of λ s at high inflation is 1.6 percent of the percentile difference at low inflation.

Table 3.1: Cross Industry Dispersion of Duration $\frac{1}{\lambda}$

Annual Inflation Range (%)	Median Duration	75-25 pct Difference	90-10 pct Difference
Homogeneous Goods			
< 10	6.37	5.47	10.30
[10, 100)	3.36	2.70	5.25
[100, 500)	2.29	1.67	3.26
≥ 500	1.03	0.77	1.28
Differentiated Goods			
< 10	11.55	11.81	24.67
[10, 100)	3.82	2.84	5.39
[100, 500)	1.05	0.65	1.23
≥ 500	0.39	0.20	0.41

Note: duration is in months and calculated as $\frac{1}{\lambda}$ for each 5-digit industry. The cross-industry statistic, e.g., 75-25 pct, is the average over all observations corresponding to inflation rates in the interval.

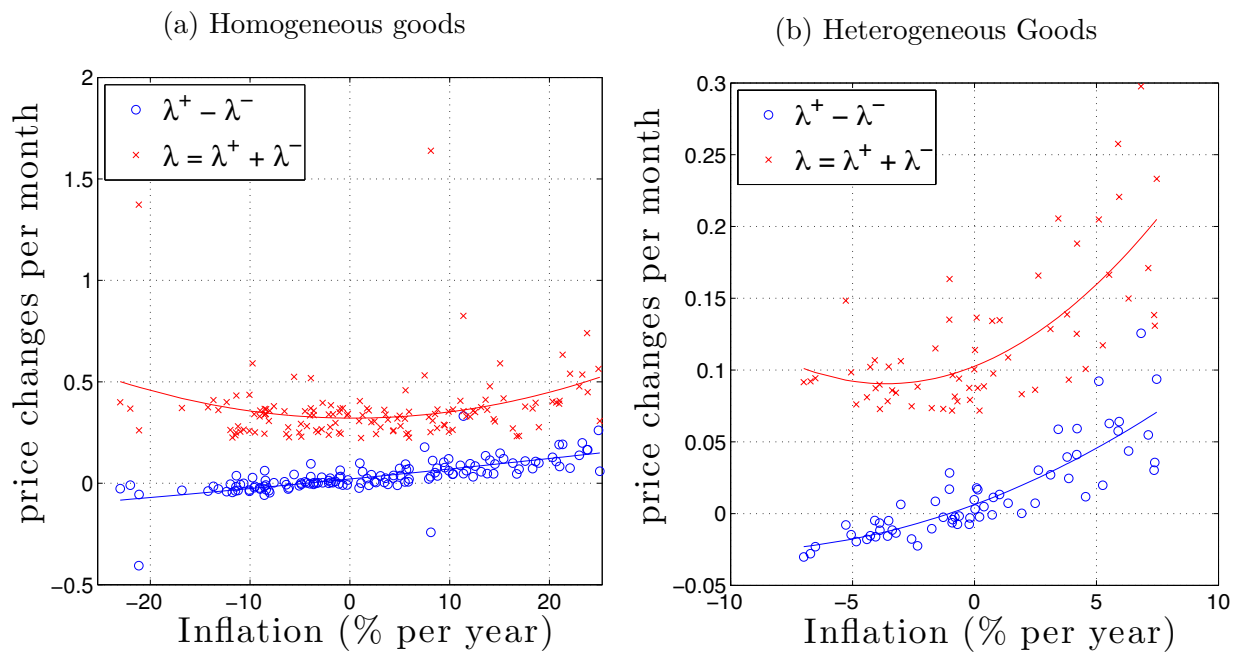
3.5.3 INFLATION AND THE INTENSIVE AND EXTENSIVE MARGINS OF PRICE ADJUSTMENTS

In this section, we confront theoretical predictions about the behavior of the intensive and extensive margins²⁶ of price changes with the data.

We first look at the predictions of the theory (propositions 3 and 5 and corollary 1) with respect to the frequency of price changes (the extensive margin of price adjustments) for near-zero inflation rates. According to theory, for near-zero inflation, the frequency of price increases and of price decreases is the same, the frequency of price changes is insensitive to inflation and the difference between the frequency of price increases and decreases rises with inflation. Figure 3.2 illustrates some of these properties of the menu cost model in the numerical example in Section 3.2.2.

²⁶See Section 3.2.1.3 for a definition of these margins

Figure 3.8: Decomposition of inflation for low inflation rates



Note: λ is the frequency of price changes per month. λ^+ (λ^-) is the frequency of price increases (decreases) per month. Inflation is the annualized log difference of the average price between two consecutive periods. The inflation range is chosen by picking the 1-percentile inflation (minimum inflation rate removing outliers) and its positive opposite. Lines are least squares second degree polynomials.

Figure 3.8 takes these predictions to the data for our two groups of goods. The red crosses plot the frequency of price changes λ against inflation while the blue circles represent the difference between the frequency of price increases and that of price decreases, $\lambda^+ - \lambda^-$. The range of inflation in the figure was chosen by picking the lowest, negative rate of inflation (excluding outliers) and its positive opposite. The quadratic function fitting the red crosses reflects both the insensitivity of the frequency of price changes to inflation as well as the symmetry between the frequency of price increases and of price decreases—i.e., quadratic functions have zero derivative and are symmetric around zero. Figure 3.9 shows the same fact by plotting the frequency of price increases (green stars) and the frequency of price decreases (red squares) against the absolute value of inflation.²⁷ Furthermore, the two panels in Figure 3.8 show that the prediction that the derivative of $\lambda^+ - \lambda^-$ with respect to inflation is positive for low rates of inflation seems to be consistent with the data.

We conclude the analysis of the extensive margin of price changes at low inflation with a variance decomposition of inflation. Proposition 5 states that, near-zero inflation, most of the changes in inflation (90 percent to be precise) result from the extensive margin. For both homogenous and heterogenous goods and ranges of annual inflation as high as 20 to -20 percent and as low as 5 to -5 percent, we find that the contribution of the extensive margin to total inflation variance is between 80 and 90 percent.²⁸

Next, we look at extensive margin predictions for high inflation rates. We showed that for high inflation rates the frequency of price increases, λ^+ , should converge to λ and that the frequency of prices decreases, λ^- converges to zero (Lemma 6 and Lemma 7). Figure 3.9 shows that this is indeed the case in the Argentine data.


Finally, the empirical behavior of the intensive margin is described in Figure 3.10, which shows the absolute value of the magnitude of price increases and of price decreases as a function of the absolute value of inflation (in semilog scale).²⁹ Figure 3.10 shows that, for low inflation rates, the size of price changes insensitive to inflation and that the size of price increases and decreases is the same (approximately 10 percent). This is consistent with

²⁷The axes in Figure 3.8 are not in log-scale, unlike Figure 3.6 and Figure 3.9, thus assuring that the insensitivity of the frequency of price adjustment to inflation is not an artifact of the scale.

²⁸The extensive margin contribution is calculated as follows. Remember that inflation can be written as $\pi = (\lambda^+ - \lambda^-)\Delta_p^+ + (\Delta_p^+ - \Delta_p^-)\lambda^-$. Then, $Var(\pi) = Cov((\lambda^+ - \lambda^-)\Delta_p^+, \pi) + Cov((\Delta_p^+ - \Delta_p^-)\lambda^-, \pi)$. Therefore, the extensive margin contribution is $\frac{Cov((\lambda^+ - \lambda^-)\Delta_p^+, \pi)}{Var(\pi)}$. Since at $\pi = 0$ the frequency of price increases and decreases are identical, this calculation approximates the theoretical extensive margin contribution $\delta(0)$.

²⁹Figure C.1 in online appendix C.2.1 is the analog of Figure 3.10, which we computed with the numerical example described in 3.2.2. The theoretical and the empirical figures are qualitatively very similar


Figure 3.9: Inflation and the Extensive Margin of Price Adjustments



Note: The frequency of price increases and decreases is calculated as $-\log(1 - f)$, where f is the fraction of outlets increasing or decreasing price in a given date.

the last part of Proposition 3. As inflation rises, the size of price increases and decreases rises, with the magnitude of price increases becoming larger than that of price decreases. This is consistent with the properties of our numerical example, shown in Figure 3.1 and in Figure C.1.

Figure 3.10: Inflation and the Intensive Margin of Price Adjustments



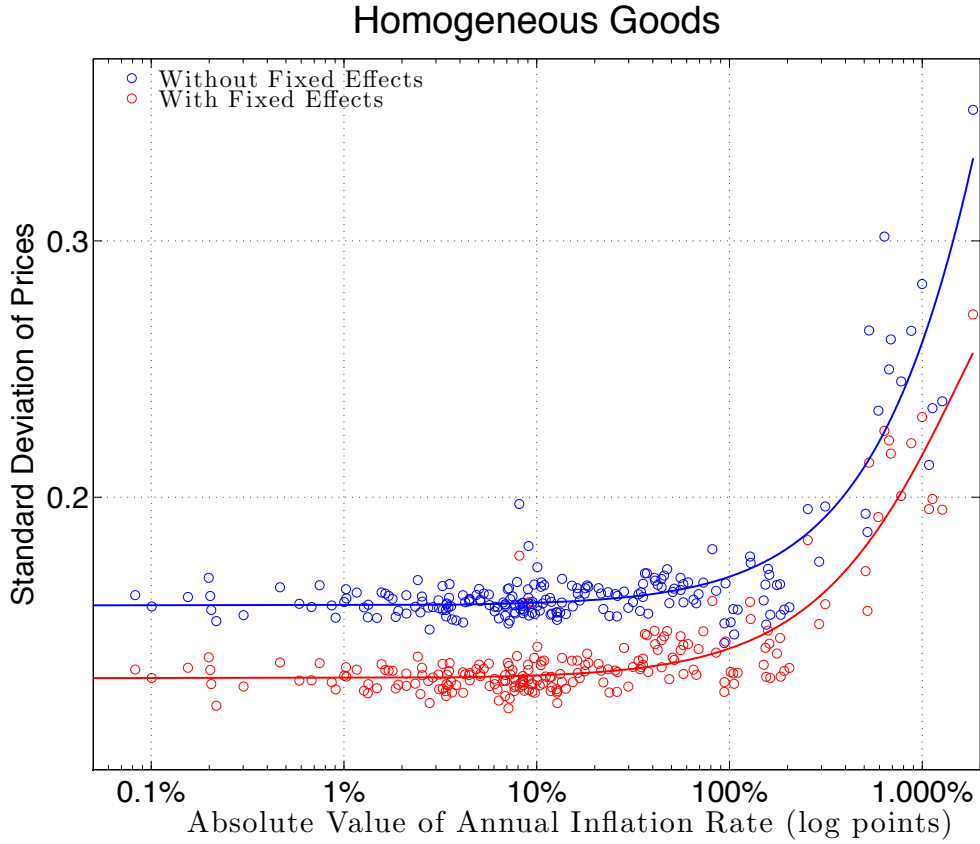
Note: The average price change is the log difference in prices, conditional on a price change taking place, averaged with expenditure weights over all homogeneous and differentiated goods in a given date.

3.5.4 INFLATION AND THE DISPERSION OF RELATIVE PRICES

In Section 3.2, we analyzed how inflation affects the dispersion of relative prices in menu cost models. As a summary, we showed that the dispersion of relative prices is insensitive to changes in inflation when inflation is low whereas it increases with inflation when inflation is high.

In this section, we explore the relationship between average price dispersion across goods

Figure 3.11: Average Dispersion of Relative Prices and Inflation



Standard deviation of prices at t is $\bar{\sigma}_t = \sum_{j=1}^N \omega_j \left[\sum_{i \in O_j} \left[(\log p_{i,j,t})^2 - \left(\frac{1}{\#O_j} \sum_{i \in O_j} \log p_{i,j,t} \right)^2 \right] \right]^{\frac{1}{2}}$, where O_j are the set of outlets that sell the good j , ω_j the expenditure share of the good j and $p_{i,j,t}$ is the price of the good j sold at outlet i at time t .

and inflation. We begin by measuring price dispersion across outlets selling the same good or service at a given month. Then, we report a weighted average of the dispersion of prices, where the weights are given by each product's expenditure share in the consumer survey. Formally, the average dispersion of relative prices at time t is given by

$$\bar{\sigma}_t = \sum_{j=1}^N \omega_j \left[\sum_{i \in O_j} \left[(\log p_{i,j,t})^2 - \left(\frac{1}{\#O_j} \sum_{i \in O_j} \log p_{i,j,t} \right)^2 \right] \right]^{\frac{1}{2}},$$

where O_j are the set of outlets that sell the good j , ω_j the expenditure share of the good

j , and $p_{i,j,t}$ is the price of the good j sold at outlet i at time t .³⁰ We compute the time series for $\bar{\sigma}_t$ among differentiated goods and homogeneous goods separately. In Figure 3.11, we plot (in a log-scale) these measures of average dispersion of relative prices against the corresponding inflation for homogeneous and differentiated goods. For robustness, we also introduce a similar dispersion measure where we control for store fixed effects.³¹

Figure 3.11 presents the empirical counterpart to the comparative statics illustrated in Figure 3.3. It plots the standard deviation of prices for homogeneous goods against the absolute value of inflation. Deflation points are represented by squares and inflation points by circles. We find that the dispersion of relative prices is unresponsive to inflation at low inflation rates and it eventually begins to increase as inflation becomes higher, as is the case in Figure 3.3. This is true whether we control for store fixed effects or not. Finally, the elasticity of the standard deviation of relative prices to inflation is close to the theoretical one-third when inflation is high.³² We conclude that the welfare costs of inflation coming from dispersion in relative prices—as emphasized in chapter 6 of Woodford (2003)—is likely to be relevant only for high rates of inflation, as evidenced by the insensitivity to inflation for inflation rates below 50 percent.

3.6 CONCLUSIONS

Guided by predictions of menu cost models of nominal price setting, we empirically analyzed how inflation affects price setting behavior by using a novel micro-dataset underlying Argentina’s consumer price index. Argentina’s experience is unique because it encompasses periods of very high and near-zero inflation, thus allowing to test sharp predictions of menu cost models in this extreme scenarios.

We found that, when inflation is low, the frequency of price changes, the dispersion of relative prices, and the absolute size of price changes are insensitive to inflation. Furthermore, we showed, both theoretically and empirically, that the difference between the frequency of price increases and decreases rises with inflation when inflation is low. These findings are consistent with predictions of menu cost models where idiosyncratic firm shocks swamp inflation as a motive for changing prices, when inflation is low.

³⁰In the computation of $\bar{\sigma}$, the set of outlets that sell a given good vary across time. Also, we exclude the goods whose prices are missing

³¹We tried other measures of dispersion as well, such as the 90-10 percentile difference, with very similar qualitative results.

³²This is true for homogenous goods but it is less apparent for differentiated goods.

At high inflation, we found that inflation swamps idiosyncratic shocks as a driver of price changes. The frequency of price changes across different products becomes similar, and the frequency of price changes, the dispersion of relative prices, and the size of price changes all rise with inflation. These findings are consistent with predictions in Sheshinski and Weiss (1977)'s menu cost model with no idiosyncratic shocks.

Finally, we confirmed and extended available evidence for the relationship between the frequency of price changes and inflation for countries that experienced either very high or low inflation. Despite large structural differences between these countries, we view these findings as reflecting common, robust economic mechanisms—captured by menu cost models—driving price changes and inflation.

BIBLIOGRAPHY

- ADAO, R., A. COSTINOT, AND D. DONALDSON (2015): “Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade,” Tech. rep., National Bureau of Economic Research.
- ALVAREZ, F., H. L. BIHAN, AND F. LIPPI (2014a): “Small and large price changes and the propagation of monetary shocks,” Tech. rep., National Bureau of Economic Research.
- (2014b): “Small and Large Price Changes and the Propagation of Monetary Shocks,” NBER Working Papers 20155, National Bureau of Economic Research, Inc.
- ALVAREZ, F. AND F. LIPPI (2014): “Price setting with menu cost for multi-product firms,” *Econometrica*, 82, 89–135.
- ALVAREZ, F. E., F. LIPPI, AND L. PACIELLO (2011): “Optimal price setting with observation and menu costs,” *The Quarterly Journal of Economics*, 126, 1909–1960.
- ALVAREZ, F. E. AND S. ZELDES (2005): “Reducing Inflation in Argentina: Mission Impossible?” *Case Study, Columbia GSB*.
- ÁLVAREZ, L. J., E. DHYNE, M. HOEBERICHTS, C. KWAPIL, H. L. BIHAN, P. LNNE-MANN, F. MARTINS, R. SABBATINI, H. STAHL, P. VERMEULEN, AND J. VILMUNEN (2006): “Sticky Prices in the Euro Area: A Summary of New Micro-Evidence,” *Journal of the European Economic Association*, 4, 575–584.
- AN, S. AND F. SCHORFHEIDE (2007): “Bayesian analysis of DSGE models,” *Econometric reviews*, 26, 113–172.
- ARKOLAKIS, C., A. COSTINOT, AND A. RODRÍGUEZ-CLARE (2012): “New Trade Models, Same Old Gains?” *American Economic Review*, 102, 94–130.

- ASDRUBALI, P., B. E. SORENSEN, AND O. YOSHA (1996): “Channels of interstate risk sharing: United States 1963-1990,” *The Quarterly Journal of Economics*, 1081–1110.
- ATKESON, A. AND T. BAYOUMI (1993): “Do private capital markets insure regional risk? Evidence from the United States and Europe,” *Open Economies Review*, 4, 303–324.
- ATKESON, A. AND A. BURSTEIN (2015): “Aggregate Implications of Innovation Policy,” .
- AUCLERT, A. (2014): “Monetary policy and the redistribution channel,” *Unpublished Manuscript, Massachusetts Institute of Technology*.
- AUTOR, D. H., D. DORN, AND G. H. HANSON (2013): “The China Syndrome: Local Labor Market Effects of Import Competition in the United States,” *The American Economic Review*, 103, 2121–2168.
- BAHARAD, E. AND B. EDEN (2004): “Price Rigidity and Price Dispersion: Evidence from Micro Data,” *Review of Economic Dynamics*, 7, 613–41.
- BALL, L. AND S. MAZUMDER (2011): “Inflation dynamics and the great recession,” *Brookings Papers on Economic Activity*, 337–406.
- BARRO, R. (1972): “A theory of monopolistic price adjustment,” *The Review of Economic Studies*, 39, 17–26.
- BARROS, R., M. BONOMO, C. CARVALHO, AND S. MATOS (2009): “Price Setting in a Variable Macroeconomic Environment: Evidence from Brazilian CPI,” Mimeo.
- BAUMEISTER, C. AND J. D. HAMILTON (2015): “Sign restrictions, structural vector autoregressions, and useful prior information,” *Econometrica*, 83, 1963–1999.
- BAYOUMI, T. AND P. R. MASSON (1995): “Fiscal flows in the United States and Canada: lessons for monetary union in Europe,” *European Economic Review*, 39, 253–274.
- BENABOU, R. (1992): “Inflation and efficiency in search markets,” *The Review of Economic Studies*, 59, 299–329.
- BENABOU, R. AND J. KONIECZNY (1994): “On inflation and output with costly price changes: A simple unifying result,” *The American Economic Review*, 290–297.

- BERAJA, M. (2013): “A Brief Note on Disinflation in a Menu Cost Model,” University of Chicago.
- (2015): “A Semi-Structural Methodology for Policy Counterfactuals and an Application to Fiscal Unions,” *University of Chicago Working Paper*.
- BERAJA, M., A. FUSTER, E. HURST, AND J. VAVRA (2015a): “Regional heterogeneity and monetary policy,” Tech. rep., Staff Report Federal Reserve Bank of New York.
- BERAJA, M., E. HURST, AND J. OSPINA (2015b): “The Aggregate Implications of Regional Business Cycles,” *Unpublished Manuscript, University of Chicago*.
- BERARDI, N., E. GAUTIER, AND H. L. BIHAN (2013): “More Facts about Prices: France Before and During the Great Recession,” Tech. rep., Banque de France.
- BERGER, D. AND J. VAVRA (2015): “Dynamics of the US price distribution,” Tech. rep., National Bureau of Economic Research.
- BERTOLA, G. AND R. CABALLERO (1990): “Kinked adjustment costs and aggregate dynamics,” *NBER Macroeconomics Annual*, 5, 237–288.
- BILS, M. AND P. J. KLENOW (2004): “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, 112, 947–985.
- BLANCHARD, O. J. AND C. M. KAHN (1980): “The solution of linear difference models under rational expectations,” *Econometrica: Journal of the Econometric Society*, 1305–1311.
- BLANCHARD, O. J., L. F. KATZ, R. E. HALL, AND B. EICHENGREEN (1992): “Regional evolutions,” *Brookings papers on economic activity*, 1–75.
- BLANCHARD, O. J. AND D. QUAH (1989): “The Dynamic Effects of Aggregate Demand and Supply Disturbances,” *The American Economic Review*, 79, 655–673.
- BLAUM, J., C. LELARGE, AND M. PETERS (2015): “The Gains From Trade in Firm-Based Models of Importing,” .
- BUCOVETSKY, S. (1998): “Federalism, equalization and risk aversion,” *Journal of Public Economics*, 67, 301–328.

- BUERA, F. AND J. P. NICOLINI (2010): “The Fiscal and Monetary History of Argentina 1960-2007,” Manuscript, UCLA.
- BUERA, F. J. AND B. MOLL (2015): “Aggregate implications of a credit crunch: The importance of heterogeneity,” *American Economic Journal: Macroeconomics*, 7, 1–42.
- BURSTEIN, A. AND C. HELLWIG (2008): “Welfare costs of inflation in a menu cost model,” *American Economic Review*, 98, 438–443.
- BURSTEIN, A. T., J. C. NEVES, AND S. REBELO (2003): “Distribution costs and real exchange rate dynamics during exchange-rate-based stabilizations,” *Journal of Monetary Economics*, 50, 1189–1214.
- CABALLERO, R. J. AND E. M. ENGEL (2007): “Price stickiness in Ss models: New interpretations of old results,” *Journal of Monetary Economics*, 54, 100–121.
- CAGE, R., J. GREENLEES, AND P. JACKMAN (2003): “Introducing the chained consumer price index,” in *International Working Group on Price Indices (Ottawa Group): Proceedings of the Seventh Meeting*, Paris: INSEE, 213–246.
- CANAY, I. AND A. SHAIKH (2015): “Practical and Theoretical Advances for Inference in Partially Identified Models,” .
- CARRIERO, A. AND R. GIACOMINI (2011): “How useful are no-arbitrage restrictions for forecasting the term structure of interest rates?” *Journal of Econometrics*, 164, 21–34.
- CAVALLO, A. (2015): “Scraped Data and Sticky Prices,” Working Paper 21490, National Bureau of Economic Research.
- CAVALLO, D. F. AND J. A. COTTANI (1997): “Argentina’s Convertibility Plan and the IMF,” *The American Economic Review*, 87, pp. 17–22.
- CHARI, V. V. AND P. J. KEHOE (2007): “On the need for fiscal constraints in a monetary union,” *Journal of Monetary Economics*, 54, 2399–2408.
- CHARLES, K., E. HURST, AND M. NOTOWIDIGDO (2015): “Masking and Employment Changes During the 2000s: Housing Booms and Manufacturing Decline,” *University of Chicago Working Paper*.

- CHARLES, K. K., E. HURST, AND M. J. NOTOWIDIGDO (2013): “Manufacturing decline, housing booms, and non-employment,” Tech. rep., National Bureau of Economic Research.
- (2014): “Housing Booms, Labor Market Outcomes, and Educational Attainment,” *Manuscript, University of Chicago*.
- CHETTY, R. (2008): “Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods,” Tech. rep., National Bureau of Economic Research.
- CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2011): “Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins,” *The American Economic Review*, 101, 471–475.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): “Monetary policy shocks: What have we learned and to what end?” *Handbook of macroeconomics*, 1, 65–148.
- CHRISTIANO, L. J., M. S. EICHENBAUM, AND M. TRABANDT (2015a): “Understanding the Great Recession,” *American Economic Journal: Macroeconomics*, 7, 110–67.
- (2015b): “Unemployment and Business Cycles,” *Econometrica* (*forthcoming*).
- CLARIDA, R., J. GALI, AND M. GERTLER (2000): “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory,” *Quarterly Journal of Economics*, 147–180.
- CLARIDA, R. H. AND M. GERTLER (1997): “How the Bundesbank conducts monetary policy,” in *Reducing inflation: Motivation and strategy*, University of Chicago Press, 363–412.
- COIBION, O., Y. GORODNICHENKO, AND G. H. HONG (2015): “The Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications,” *American Economic Review*, 105, 993–1029.
- DANZIGER, L. (1999): “A Dynamic Economy with Costly Price Adjustments,” *American Economic Review*, 89, 878–901.
- DEJONG, D. N. AND C. DAVE (2011): *Structural macroeconometrics*, Princeton University Press.

- DEL NEGRO, M., M. P. GIANNONI, AND F. SCHORFHEIDE (2015): “Inflation in the Great Recession and New Keynesian Models,” *American Economic Journal: Macroeconomics*, 7, 168–196.
- DEL NEGRO, M. AND F. SCHORFHEIDE (2004): “Priors from general equilibrium models for vars*,” *International Economic Review*, 45, 643–673.
- (2009): “Monetary Policy Analysis with Potentially Misspecified Models,” *The American Economic Review*, 99, 1415.
- DIEWERT, W. E. (1976): “Exact and superlative index numbers,” *Journal of econometrics*, 4, 115–145.
- DOTSEY, M. AND R. G. KING (2005): “Implications of state-dependent pricing for dynamic macroeconomic models,” *Journal of Monetary Economics*, 52, 213–242.
- EGGERTSSON, G. B. AND P. KRUGMAN (2012): “Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach*,” *The Quarterly Journal of Economics*, qjs023.
- EVERS, M. P. (2015): “Fiscal Federalism and Monetary Unions: A Quantitative Assessment,” *Journal of International Economics*.
- FARHI, E. AND I. WERNING (2014): “Fiscal unions,” *Unpublished Manuscript, Massachusetts Institute of Technology*.
- FEYRER, J. AND B. SACERDOTE (2013): “How Much Would US Style Fiscal Integration Buffer European Unemployment and Income Shocks?(A Comparative Empirical Analysis),” *The American Economic Review*, 103, 125–128.
- FITZGERALD, T. AND J. P. NICOLINI (2014): “Is there a stable relationship between unemployment and future inflation? evidence from us cities,” *Federal Reserve Bank of Minneapolis, Working Paper*, 713.
- GAGNON, E. (2009): “Price Setting during Low and High Inflation: Evidence from Mexico,” *Quarterly Journal of Economics*, 124, 1221–1263.
- GALÍ, J. (2010): “Monetary Policy and Unemployment,” *NBER Working Paper Series*, 15871.

- GERTLER, M. AND J. LEAHY (2008): “A Phillips curve with an Ss foundation,” *Journal of Political Economy*, 116.
- GILCHRIST, S., J. W. SIM, AND E. ZAKRAJŠEK (2014): “Uncertainty, financial frictions, and investment dynamics,” Tech. rep., National Bureau of Economic Research.
- GIROUD, X. AND H. M. MUELLER (2015): “Firm Leverage and Unemployment during the Great Recession,” *NBER Working Paper Series*, 21076.
- GOLOSOV, M. AND R. E. J. LUCAS (2007): “Menu Costs and Phillips Curves,” *Journal of Political Economy*, 115, 171–199.
- GUERRIERI, V. AND G. LORENZONI (2011): “Credit crises, precautionary savings, and the liquidity trap,” Tech. rep., National Bureau of Economic Research.
- HALL, R. E. (2011): “The Long Slump,” *THE AMERICAN ECONOMIC REVIEW*, 101, 431–469.
- HALL, R. E. AND P. R. MILGROM (2008): “The Limited Influence of Unemployment on the Wage Bargain,” *The American Economic Review*, 98, 1653–1674.
- HAMILTON, J. D. (1994): *Time series analysis*, vol. 2, Princeton university press Princeton.
- HANSEN, G. D. (1985): “Indivisible labor and the business cycle,” *Journal of monetary Economics*, 16, 309–327.
- HANSEN, L. P. AND T. J. SARGENT (2008): *Robustness*, Princeton university press.
- HARBERGER, A. C. (1964): “The measurement of waste,” *The American Economic Review*, 58–76.
- HURST, E., B. J. KEYS, A. SERU, AND J. S. VAVRA (2015): “Regional Redistribution Through the US Mortgage Market,” Tech. rep., National Bureau of Economic Research.
- IVANCIC, L., W. E. DIEWERT, AND K. J. FOX (2011): “Scanner data, time aggregation and the construction of price indexes,” *Journal of Econometrics*, 161, 24–35.
- JAIMOVICH, N. AND H. E. SIU (2014): “The Trend is the Cycle: Job Polarization and Jobless Recoveries,” *Duke University Working Paper*.

- KAPLAN, G. AND G. MENZIO (2015): “The morphology of price dispersion,” *International Economic Review*, 56, 1165–1206.
- KEHOE, P. J. AND V. MIDRIGAN (2010): “Prices are Sticky After All,” NBER Working Papers 16364, National Bureau of Economic Research, Inc.
- KIGUEL, M. A. (1991): “Inflation in Argentina: Stop and go since the Austral plan,” *World Development*, 19, 969 – 986.
- KLENOW, P. J. AND O. KRYVTSOV (2008): “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?” *The Quarterly Journal of Economics*, 123, 863–904.
- KLENOW, P. J. AND B. A. MALIN (2011): “Microeconomic Evidence on Price-Setting,” in *Handbook of Monetary Economics 3A*, ed. by B. Friedman and M. Woodford, Elsevier, 231–284.
- KOLLMANN, R. (2001): “Explaining international comovements of output and asset returns: the role of money and nominal rigidities,” *Journal of Economic Dynamics and Control*, 25, 1547–1583.
- KONIECZNY, J. D. AND A. SKRZYPACZ (2005): “Inflation and price setting in a natural experiment,” *Journal of Monetary Economics*, 52, 621–632.
- KYDLAND, F. E. AND E. C. PRESCOTT (1982): “Time to build and aggregate fluctuations,” *Econometrica: Journal of the Econometric Society*, 1345–1370.
- LACH, S. AND D. TSIDDON (1992): “The Behavior of Prices and Inflation: An Empirical Analysis of Dissaggregated Price Data,” *Journal of Political Economy*, 100, 349–89.
- LOCKWOOD, B. (1999): “Inter-regional insurance,” *Journal of Public Economics*, 72, 1–37.
- LUCAS, R. E. (1976): “Econometric policy evaluation: A critique,” in *Carnegie-Rochester conference series on public policy*, Elsevier, vol. 1, 19–46.
- MCKAY, A. AND R. REIS (2013): “The role of automatic stabilizers in the US business cycle,” Tech. rep., CEPR Discussion Papers.

- MELSER, D. (2011): “Constructing High Frequency Price Indexes Using Scanner Data,” in *twelfth meeting of the Ottawa Group*, 4–6.
- MIAN, A., K. RAO, A. SUFI, ET AL. (2013): “Household Balance Sheets, Consumption, and the Economic Slump,” *The Quarterly Journal of Economics*, 128, 1687–1726.
- MIAN, A. AND A. SUFI (2011): “House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis,” *American Economic Review*, 101, 2132–56.
- (2014): “What Explains the 2007–2009 Drop in Employment?” *Econometrica*, 82, 2197–2223.
- MIDRIGAN, V. AND T. PHILIPPON (2011): “Household leverage and the recession,” Tech. rep., National Bureau of Economic Research.
- MONDRAGON, J. (2015): “Household credit and employment in the Great Recession,” *University of California ta Berkeley Working Paper*.
- MULLIGAN, C. B. (2012): *The redistribution recession: How labor market distortions contracted the economy*, Oxford University Press.
- NAKAMURA, E. AND J. STEINSSON (2008): “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *Quarterly Journal of Economics*, 123, 1415–1464.
- (2010): “Monetary Non-Neutrality in a Multisector Menu Cost Model,” *The Quarterly Journal of Economics*, 125, 961–1013.
- (2014): “Fiscal Stimulus in a Monetary Union: Evidence from US Regions,” *The American Economic Review*, 104, 753–792.
- NAKAMURA, E., J. STEINSSON, AND D. VILLAR (2016): “The Elusive Costs of Inflation: Price Dispersion during the US Great Inflation,” Columbia University.
- OH, H. AND R. REIS (2012): “Targeted transfers and the fiscal response to the great recession,” *Journal of Monetary Economics*, 59, S50–S64.
- PERSSON, T. AND G. TABELLINI (1996a): “Federal fiscal constitutions: Risk sharing and moral hazard,” *Econometrica: Journal of the Econometric Society*, 623–646.

- (1996b): “Federal fiscal constitutions: risk sharing and redistribution,” *Journal of Political Economy*, 979–1009.
- PESARAN, M. H., R. G. PIERSE, AND K. C. LEE (1993): “Persistence, cointegration, and aggregation: A disaggregated analysis of output fluctuations in the US economy,” *Journal of Econometrics*, 56, 57–88.
- RUGGLES, S., M. SOBEK, C. A. FITCH, P. K. HALL, AND C. RONNANDER (1997): *Integrated public use microdata series: Version 2.0*, Historical Census Projects, Department of History, University of Minnesota.
- SALA-I MARTIN, X. AND J. D. SACHS (1991): “Fiscal Federalism and Optimum Currency Areas: Evidence for Europe from the United States,” *NBER Working Paper*.
- SCHMITT-GROHÉ, S. AND M. URIBE (2003): “Closing small open economy models,” *Journal of international Economics*, 61, 163–185.
- SCHORFHEIDE, F. (2000): “Loss function-based evaluation of DSGE models,” *Journal of Applied Econometrics*, 15, 645–670.
- SERGEYEV, D. AND N. MEHROTRA (2015): “Financial Shocks and Job Flows,” in *2015 Meeting Papers*, Society for Economic Dynamics, 520.
- SHESHINSKI, E. AND Y. WEISS (1977): “Inflation and Costs of Price Adjustment,” *Review of Economic Studies*, 44, 287–303.
- SHIMER, R. (2010): *Labor Markets and Business Cycles*, Princeton University Press.
- SIMS, C. A. (1980): “Macroeconomics and reality,” *Econometrica: Journal of the Econometric Society*, 1–48.
- STOCK, J. H. AND M. W. WATSON (2001): “Vector autoregressions,” *Journal of Economic perspectives*, 101–115.
- STOCK, J. H. AND M. W. WATSON (2012): “Disentangling the Channels of the 2007-09 Recession,” *Brookings Papers on Economic Activity: Spring 2012*, 81.
- STROEBEL, J. AND J. VAVRA (2014): “House Prices, Local Demand, and Retail Prices,” *NBER Working Paper Series*, 20710.

- UHLIG, H. (1995): “A toolkit for analyzing nonlinear dynamic stochastic models easily,” .
- (2005): “What are the effects of monetary policy on output? Results from an agnostic identification procedure,” *Journal of Monetary Economics*, 52, 381–419.
- VAVRA, J. (2014): “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation,” *The Quarterly journal of economics*, 129, 215–258.
- WOODFORD, M. (2003): *Interest and prices*, Princeton University Press.
- WULFSBERG, F. (2010): “Price adjustments and inflation - evidence from Norwegian consumer price data 1975-2004,” .
- YAGAN, D. (2014): “Moving to opportunity? Migratory insurance over the great recession,” *Working Paper*.

Appendix A

A SEMI-STRUCTURAL METHODOLOGY FOR POLICY COUNTERFACTUALS WITH AN APPLICATION TO FISCAL UNIONS

A.1 PROOF OF LEMMA 1

The following equations characterize the log-linearized equilibrium

$$\begin{aligned}
 \tilde{w}_{kt} - \tilde{p}_{kt} &= \frac{1}{\phi} \tilde{n}_{kt} \\
 \tilde{w}_{kt} - \tilde{p}_{kt} &= (\alpha - 1) \tilde{n}_{kt} + \beta \tilde{x}_{kt} + \tilde{z}_{kt} \\
 \tilde{q}_t - \tilde{p}_{kt} &= \alpha \tilde{n}_{kt} + (\beta - 1) \tilde{x}_{kt} + \tilde{z}_{kt} \\
 0 &= \mathbb{E}_t \left(\tilde{m}u_{kt+1} - \tilde{m}u_{kt+1} - (\tilde{p}_{kt+1} - \tilde{p}_{kt}) - \gamma_{kt+1} - \Phi_0(\tilde{c}_{kt} - \tilde{c}_t) + \tilde{i}_{t+1} \right) \\
 \tilde{m}u_{kt+1} &= -\frac{\sigma}{C - \frac{\phi}{1+\phi} N^{\frac{1+\phi}{\phi}}} \left(C \tilde{c}_{kt+1} - N^{\frac{1+\phi}{\phi}} \tilde{n}_{kt+1} \right) \\
 \tilde{c}_{kt} &= \tilde{w}_{kt} - \tilde{p}_{kt} + \tilde{n}_{kt} \\
 B \tilde{b}_{kt} &= B(1+r)(\tilde{b}_{kt-1} + \tilde{i}_t) + \eta_{kt} - X(\tilde{q}_t + \tilde{x}_{kt}) + S \tilde{s}_{kt} \\
 \sum_k \tilde{x}_{kt} &= \sum_k \tilde{\eta}_{kt} \\
 B^g \tilde{b}_t^g + S \sum_k \tilde{s}_{kt} + G \tilde{q}_t &= B^g(1+r)(\tilde{b}_{t-1}^g + \tilde{i}_t) \\
 \tilde{s}_{kt} &= \vartheta_w \tilde{w}_{kt} + \vartheta_n \tilde{n}_{kt} + \vartheta_b \tilde{b}_{kt-1} \\
 \tilde{i}_{t+1} &= \phi_p \mathbb{E}_t[\tilde{p}_{t+1} - \tilde{p}_t]
 \end{aligned}$$

After adding up, the aggregate log-linearized equilibrium evolution of $\{\tilde{w}_t - \tilde{p}_t, \tilde{n}_t\}$ is characterized by

$$\begin{aligned}
0 &= \mathbb{E}_t(\tilde{m}u_{t+1} - \tilde{m}u_t + (\phi_p - 1)(\tilde{p}_{t+1} - \tilde{p}_t) + \tilde{\gamma}_{t+1}) \\
0 &= \frac{1}{\phi}\tilde{n}_t - (\tilde{w}_t - \tilde{p}_t) \\
\tilde{w}_t - \tilde{p}_t &= (\alpha - 1)\tilde{n}_t + \tilde{z}_t + \beta\tilde{\eta}_t \\
\tilde{m}u_{t+1} &\equiv -\frac{\sigma}{C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}} \left(C(\tilde{w}_{t+1} - \tilde{p}_{t+1} + \tilde{n}_{t+1}) - N^{\frac{1+\phi}{\phi}}(\tilde{n}_{t+1}) \right)
\end{aligned}$$

which is equivalent to the system of equations characterizing the log-linearized equilibrium in a representative agent economy with a production technology that utilizes labor alone with an elasticity of α , no endogenous discounting and only 2 exogenous processes $\{\tilde{z}_t + \beta\tilde{\eta}_t, \tilde{\gamma}_t\}$.

Next, take log-deviations from the aggregate in the original system. This results in the system characterizing the evolution of $\{p_{kt}, w_{kt}, n_{kt}, b_{kt}, c_{kt}\}$ for given $\{z_{kt}, \eta_{kt}, \gamma_{kt}\}$

$$\begin{aligned}
0 &= \mathbb{E}_t \left[\frac{\sigma\beta}{\frac{\alpha\phi}{1+\phi} - 1} (w_{kt+1} + n_{t+1} - w_{kt} - n_{kt}) + (\alpha + \beta - 1)(n_{kt+1} - n_{kt}) \right. \\
&\quad \left. + (\beta - 1)(w_{kt+1} - w_{kt}) + \left(1 + \frac{\sigma}{\frac{\alpha\phi}{1+\phi} - 1}\right)(z_{kt+1} - z_{kt}) + \Phi_0 c_{kt} - \gamma_{kt+1} \right] \\
\beta w_{kt} &= \left(\frac{1+\phi}{\phi} - (\alpha + \beta) \right) n_{kt} - z_{kt} \\
\frac{B}{S} b_{kt} &= \frac{B}{S} (1+r) b_{kt-1} - \frac{X}{S} (w_{kt} + n_{kt}) + \vartheta_n n_{kt} + \vartheta_w w_{kt} + \vartheta_b b_{kt-1} + \frac{1}{S} \eta_{kt} \\
c_{kt} &= w_{kt} - p_{kt} + n_{kt} \\
p_{kt} &= (1 - \beta)w_{kt} - (z_{kt} + (\alpha + \beta - 1)n_{kt})
\end{aligned}$$

This system is independent of all aggregate variables and is analogous to the system characterizing the equilibrium in a small open economy without movements in the terms of trade and nominal interest rate, proving the Lemma.

A.2 PROOF OF CLAIM 1

. For any ν^1 we can find an alternative parameterization,

$$\begin{aligned} \frac{1 - \lambda_1}{\lambda_1} &= \frac{1 - \lambda^0}{\lambda^0} \frac{\nu^1}{\nu^0}, \beta^1 = \beta^0 \frac{\nu^1}{\nu^0}, \alpha^1 = \alpha^0, \frac{\frac{1+\phi^1}{\phi^1} - \alpha^1}{\nu^1} = \frac{\frac{1+\phi^0}{\phi^0} - \alpha^0}{\nu^0} \\ 1 + \frac{\sigma^1}{\left(\frac{\alpha^1 \phi^1}{1+\phi^1} - 1\right)} &= \left(1 + \frac{\sigma_0}{\frac{\alpha^0 \phi^0}{1+\phi^0} - 1}\right) \frac{\nu^0}{\nu^1}, \sigma_z^1 = \sigma_z^0 \frac{\nu^1}{\nu^0} \\ \Phi_0^1 &= \Phi_0^0 + \frac{(\sigma^1 \frac{\nu^0}{\nu^1} - \sigma^0) \alpha^0}{\alpha^0 - \frac{1+\phi^0}{\phi^0}}, \Phi_1^1 = \Phi_1^0 - \frac{(\sigma^1 \frac{\nu^0}{\nu^1} - \sigma^0) \alpha^0}{\alpha^0 - \frac{1+\phi^0}{\phi^0}} \end{aligned}$$

Replacing these in the system of equations characterizing the equilibrium we obtain the exact same system of equations corresponding to ξ^0 . Since ν^1 was arbitrarily chosen, this implies that there are many parameterizations consistent with the same ρ^0, Λ^0 . Moreover, in all of these parameterizations $\frac{1-\lambda_1}{\lambda_1} = \frac{1-\lambda^0}{\lambda^0} \frac{\nu^1}{\nu^0}$.

A.3 PROOF OF PROPOSITION 1

The proof is by construction. Applying the method of undetermined coefficients to (FiscalSME) implies that $\{P, \Lambda\}$ solve

$$\begin{aligned} (FP + G + \Theta_c)P + H + \Theta_p &= 0 \\ (LN + M)\Sigma + F\Lambda\Sigma^{-1}N\Sigma + (FP + G + \Theta_c)\Lambda &= 0 \end{aligned}$$

Also from Claim 3 we have

$$\begin{aligned} \Lambda\Sigma^{-1}N\Sigma + (\rho_1 - P)\Lambda &= 0 \\ (P - \rho_1)P - \rho_2 &= 0 \end{aligned}$$

Replacing in the first system, we obtain a new system

$$\begin{aligned} (FP + G + \Theta_c)P + H + \Theta_p &= 0 \\ (LN + M)\Sigma + (F\rho_1 + G + \Theta_c)\Lambda &= 0 \end{aligned}$$

Given $\{\rho_1, \rho_2\}$, let P be a solution with all eigenvalues inside the unite circle to the matrix quadratic equation $0 = (P - \rho_1)P - \rho_2$. Without loss of generality, take equations involving the first line of structure ξ , i.e. the Euler equation. Because the first line of Θ is all zeros by Assumption 2, this system can be written as

$$\begin{bmatrix} (P')^2 & P' & I & 0 \\ \Lambda' \rho'_1 & \Lambda' & 0 & I \end{bmatrix}_{6 \times 12} \begin{bmatrix} F'_1 \\ G'_1 \\ H'_1 \\ ((LN + M)\Sigma)'_1 \end{bmatrix}_{12 \times 1} = \mathbf{0}_{6 \times 1}$$

There are 12 unknowns in the first line of structure ξ . There are 6 equations in this system. One of the unknown elements can be normalized to 1 because the system is equalized to zero. This implies that by specifying at least 5 zeros in the first line of structure ξ we obtain a linear system of equations that is exactly determined and can be solved for the remaining unspecified elements in the structure (as long as the matrix containing $\{P, \Lambda, \rho_1\}$ has full rank).

An identical proof holds for identifying the structure ξ for in the other two equations: the sequential budget constraint (SB) and (Labor market). Finally, we solve the following system for $\{P(\xi, \Theta'), \Lambda(\xi, \Theta')\}$ given $\{P(\xi, \Theta), \rho_1(\xi, \Theta), \Lambda(\xi, \Theta)\}$ and $\{F, G, \Theta\}$. The system results from subtracting (SME Fiscal) for policy Θ and $\Theta' = 0$.

$$\begin{aligned} 0 &= F(P(\xi, \Theta')^2 - P(\xi, \Theta)^2) + G(P(\xi, \Theta') - P(\xi, \Theta)) - \Theta_c P(\xi, \Theta) - \Theta_p \\ 0 &= (FP(\xi, \Theta') + G) (\Lambda(\xi, \Theta')\Lambda(\xi, \Theta)^{-1} - I) \\ &\quad + F (\Lambda(\xi, \Theta')\Lambda(\xi, \Theta)^{-1} - I) (\rho_1(\xi, \Theta) - P(\xi, \Theta)) + F(P(\xi, \Theta') - P(\xi, \Theta)) - \Theta_c \end{aligned}$$

And obtain $\{\rho_1(\xi, \Theta'), \rho_2(\xi, \Theta')\}$ as:

$$\begin{aligned} \rho_1(\xi, \Theta') &= P(\xi, \Theta') + \Lambda(\xi, \Theta')\Lambda(\xi, \Theta)^{-1}(\rho_1(\xi, \Theta) - P(\xi, \Theta)) (\Lambda(\xi, \Theta')\Lambda(\xi, \Theta)^{-1})^{-1} \\ \rho_2(\xi, \Theta') &= (P(\xi, \Theta') - \rho_1(\xi, \Theta'))P(\xi, \Theta') \end{aligned}$$

which concludes the proof. ■

A.4 PROOF OF PROPOSITION 2

The proof is by construction. We have that $P(\xi, \Theta), \Lambda(\xi, \Theta), \rho_1(\xi, \Theta), \rho_2(\xi, \Theta)$ satisfy the following system of matrix equations for all $\{\xi, \Theta\}$,

$$((F + \Theta_f)P + G + \Theta_c)P + H + \Theta_p = 0 \quad (1)$$

$$(LN + M)\Sigma + (F + \Theta_f)\Lambda\Sigma^{-1}N\Sigma + ((F + \Theta_f)P + G + \Theta_c)\Lambda = 0 \quad (2)$$

$$\Lambda\Sigma^{-1}N\Sigma + (\rho_1 - P)\Lambda = 0 \quad (3)$$

$$(P - \rho_1)P - \rho_2 = 0 \quad (4)$$

Under Assumption 8, equation (4) has a unique solution $P(\xi, \Theta)$ with all eigenvalues inside the unit circle for given $\rho_1(\xi, \Theta), \rho_2(\xi, \Theta)$. Subtracting equation (1) evaluated at Θ and Θ' from each other, and doing the same for equation 2 (where we replace N using 3) we obtain a system of matrix equations to solve for $\{P(\xi, \Theta'), \Lambda(\xi, \Theta')\}$ given $\{P(\xi, \Theta), \Lambda(\xi, \Theta), F, G\}$:

$$\begin{aligned} 0 &= (F + \Theta'_f)(P(\xi, \Theta')^2 - P(\xi, \Theta)^2) + (G + \Theta'_c)(P(\xi, \Theta') - P(\xi, \Theta)) \\ &\quad + (\Theta'_f - \Theta_f)P(\xi, \Theta)^2 + (\Theta'_c - \Theta_c)P(\xi, \Theta) + \Theta'_p - \Theta_p \\ 0 &= ((F + \Theta'_f)P(\xi, \Theta') + G + \Theta'_c) (\Lambda(\xi, \Theta')\Lambda(\xi, \Theta)^{-1} - I) \\ &\quad + (F + \Theta'_f) (\Lambda(\xi, \Theta')\Lambda(\xi, \Theta)^{-1} - I) (\rho_1(\xi, \Theta) - P(\xi, \Theta)) \\ &\quad + (F + \Theta'_f)(P(\xi, \Theta') - P(\xi, \Theta)) + (\Theta'_f - \Theta_f)\rho_1(\xi, \Theta) + \Theta'_c - \Theta_c \end{aligned} \quad (\text{SME2})$$

Finally, we can obtain $\{\rho_1(\xi, \Theta'), \rho_2(\xi, \Theta')\}$ as:

$$\begin{aligned} \rho_1(\xi, \Theta') &= P(\xi, \Theta') + \Lambda(\xi, \Theta')\Lambda(\xi, \Theta)^{-1}(\rho_1(\xi, \Theta) - P(\xi, \Theta)) (\Lambda(\xi, \Theta')\Lambda(\xi, \Theta)^{-1})^{-1} \\ \rho_2(\xi, \Theta') &= (P(\xi, \Theta') - \rho_1(\xi, \Theta'))P(\xi, \Theta') \end{aligned}$$

which concludes the proof. ■

A.5 PROOF OF LEMMA 2

From system (5)-(8), we can write the equations involving the first line of $\{F, G\}$ as:

$$\begin{aligned} & \begin{bmatrix} P'P' & P' \\ \Lambda'\rho'_1 & \Lambda' \\ \Lambda'(\rho'_1\rho'_1 + \rho'_2) & \Lambda'\rho'_1 \end{bmatrix} \begin{bmatrix} F' \\ G' \end{bmatrix} \begin{bmatrix} 1 \\ 0_{k.1} \end{bmatrix} + \begin{bmatrix} I_{k.k} \\ 0_{k.k} \\ \Lambda' \end{bmatrix} H' \begin{bmatrix} 1 \\ 0_{k.1} \end{bmatrix} \\ & + \begin{bmatrix} P' & I \\ \Lambda' & 0 \\ \Lambda'\rho'_1 & \Lambda' \end{bmatrix} \begin{bmatrix} \Theta'_c \\ \Theta'_p \end{bmatrix} \begin{bmatrix} 1 \\ 0_{k.1} \end{bmatrix} + \begin{bmatrix} 0_{k.k} \\ \Sigma \\ \Sigma N' \end{bmatrix} (LN + M)' \begin{bmatrix} 1 \\ 0_{k.1} \end{bmatrix} = 0_{3k.1} \end{aligned}$$

There are $3k$ elements in the first line of $\{F, G, H\}$. Let j be the number of specified elements in $[1 \ 0_{1.k}](LN + M)\Sigma[I \ N]$. Then there are $k + j$ zero elements in $[1 \ 0_{1.k}](LN + M)\Sigma[0_{k.k} \ I \ N]$. Also if $\text{rank}([1 \ 0_{1.k}][\Theta_c \ \Theta_p]) = 0$ we can normalize one of the elements in the first line of $\{F, G, H\}$ to 1. This implies that it is necessary to specify $n \equiv 3k - k - j - (1 - \text{rank}([1 \ 0_{1.k}][\Theta_c \ \Theta_p]))$ elements in the first line of $\{F, G, H\}$ alone.

A.6 PROOF OF LEMMA 3

There are k^2 elements in $\Lambda(\xi, \Theta)$ that we would like to identify. The proof of the lemma consist in deriving k^2 independent equations to generate a system that is exactly determined.

From the orthogonalization conditions in 4., we obtain $k(k+1)/2$ equations. Since $V(\xi, \Theta)$ is symmetric we do not obtain k^2 equations, but only a number of equations equal to the lower (or upper) triangle. Conditions 1. to 3. imply that the system derived from picking 's' lines in equations (6) and (7) and specifying $r(k) = k(k-1)/2$ elements in $A(s)$


$$A(s) + B(s)\Lambda(\xi, \Theta) = 0$$

generates $r(k)$ independent equations that includes all the elements in all the columns of $\Lambda(\xi, \Theta)$.

Finally, we have described a total of $k(k+1)/2 + r(k) = k^2$ equations to identify the k^2 parameters in $\Lambda(\xi, \Theta)$.

A.7 FIGURES AND TABLES

Figure A.1: Employment in the Great Recession: Channels Decomposition



Note: Figure shows kernel density estimates of the distribution of employment growth between 2008 and 2010 across states. This is done for the actual economy with rigidities and transfers, and for all other counterfactual economies constructed in Section 1.6.6.

Table A.1: Counterfactual Employment Dispersion without Fiscal Integration

	$(\vartheta_n, \vartheta_w, \vartheta_b)$			
	$(-1.6, -0.9, 0.03)$	$(-1.4, -1.2, 0)$	$(-1.2, -1.4, 0)$	$(-1.1, -1.1, 0)$
σ_n^{2010}	3.5	3.4	3.3	3.1
$\bar{\sigma}_n$	4.9	4.9	4.7	4.5

Note: Table shows selected results of the counterfactual exercise for alternative parameterizations of the prevalent transfer policy rule. σ_n^{2010} is the standard deviation of the counterfactual distribution of employment across states in the year 2010 in percents in an economy without transfers. $\bar{\sigma}_n$ is the counterfactual standard deviation in the stationary distribution.

Appendix B

THE AGGREGATE IMPLICATIONS OF REGIONAL BUSINESS CYCLES

B.1 PROOF OF LEMMA 4 AND 5

The following equations characterize the log-linearized equilibrium

$$\begin{aligned}
 w_{kt}^r &= \lambda(\epsilon_{kt} + \frac{1}{\phi}n_{kt}) + (1 - \lambda)(w_{kt-1}^r - \pi_{kt}) \\
 w_{kt}^r &= -(1 - (\alpha + \theta\beta))n_{kt}^y - \beta(1 - \theta)(n_{kt}^x - n_{kt}^y) + z_{kt}^y + \beta z_{kt}^x \\
 0 &= \mathbb{E}_t(mu_{kt+1} - mu_{kt} - \pi_{kt+1} - \gamma_{kt+1} - \Phi_0(c_{kt} - c_t) + \varphi_p\pi_t + \varphi_y(c_t - c_t^*)) \\
 mu_{kt+1} &= -\frac{\sigma}{C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}} \left(Cc_{kt+1} - N^{\frac{1+\phi}{\phi}} \left(\frac{1+\phi}{\phi} \epsilon_{kt+1} + n_{kt+1} \right) \right) \\
 Nn_{kt} &= N^x n_{kt}^x + N^y n_{kt}^y \\
 c_{kt} &= w_{kt}^r + n_{kt}^y \\
 b_{kt} &= (1 + r)(b_{kt-1} + i_t) + \frac{X}{B}(z_{kt}^x + \theta n_{kt}^x - x_{kt}) - r\tau_t \\
 0 &= z_{kt}^x - (w_{kt} - q_t) - (1 - \theta)n_{kt}^x \\
 x_{kt} &= n_{kt}^y + (w_{kt} - q_t) \\
 \sum_k x_{kt} &= \sum_k (z_{kt}^x + \theta n_{kt}^x)
 \end{aligned}$$

From the last 3 equations, after adding up, it holds that $n_t^x = n_t^y$. Then the aggregate log-linearized equilibrium evolution of $\{\pi_t^w, w_t^r, n_t\}$ is characterized by

$$\begin{aligned}
0 &= \mathbb{E}_t(mu_{t+1} - mu_t - \pi_{t+1} - \gamma_{t+1}) + \varphi_p \mathbb{E}_t[\pi_{t+1}] + \varphi_y(w_t^r + n_t) \\
\pi_t^w &= \frac{\lambda}{1-\lambda}(\epsilon_t + \frac{1}{\phi}n_t - w_t^r) \\
w_t^r &= -(1 - (\alpha + \theta\beta))n_t + z_t \\
mu_{t+1} &\equiv -\frac{\sigma}{C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}} \left(C(w_{t+1}^r + n_{t+1}) - N^{\frac{1+\phi}{\phi}} \left(\frac{1+\phi}{\phi} \epsilon_{t+1} + n_{t+1} \right) \right) \\
\pi_{t+1} &\equiv \pi_{t+1}^w - (w_{t+1}^r - w_t^r)
\end{aligned}$$

which is equivalent to the system of equations characterizing the log-linearized equilibrium in a representative agent economy with a production technology that utilizes labor alone with an elasticity of $\alpha + \theta\beta$, no endogenous discounting and only 3 exogenous processes $\{z_t, \epsilon_t, \gamma_t\}$. The top equation is the aggregate Euler equation. The second equation is the aggregate wage setting equation. The third equation is effectively the aggregate labor demand curve.

To prove Lemma 5, just take log-deviations from the aggregate in the original system. This results in the system characterizing the evolution of $\{\tilde{p}_t, \tilde{w}_t, \tilde{n}_t^y, \tilde{n}_t^x\}$ for given $\{\tilde{z}_t^y, \tilde{z}_t^x, \tilde{\gamma}_t, \tilde{\epsilon}_t\}$,

$$\begin{aligned}
\tilde{w}_t &= \lambda \left(\tilde{p}_t + \tilde{\epsilon}_t + \frac{1}{\phi} \left(\frac{N^x}{N} \tilde{n}_t^x + \frac{N^y}{N} \tilde{n}_t^y \right) \right) + (1-\lambda)\tilde{w}_{t-1} \\
\tilde{w}_t &= \tilde{p}_t - (1 - (\alpha + \theta\beta))\tilde{n}_t^y - \beta(1-\theta)(\tilde{n}_t^x - \tilde{n}_t^y) + \tilde{z}_t^y + \beta\tilde{z}_t^x \\
\tilde{w}_t &= \tilde{z}_t^x - (1-\theta)\tilde{n}_t^x \\
0 &= \mathbb{E}_t(\tilde{m}u_{t+1} - \tilde{m}u_t - (\tilde{p}_{t+1} - \tilde{p}_t) - \Phi_0(\tilde{w}_t - \tilde{p}_t + \tilde{n}_t^y) - \tilde{\gamma}_{t+1}) \\
\tilde{m}u_{t+1} &\equiv -\frac{\sigma}{C - \frac{\phi}{1+\phi}N^{\frac{1+\phi}{\phi}}} \left(C(\tilde{w}_{t+1} - \tilde{p}_{t+1} + \tilde{n}_{t+1}^y) \right. \\
&\quad \left. - N^{\frac{1+\phi}{\phi}} \left(\frac{1+\phi}{\phi} \tilde{\epsilon}_{t+1} + \left(\frac{N^x}{N} \tilde{n}_{t+1}^x + \frac{N^y}{N} \tilde{n}_{t+1}^y \right) \right) \right) \\
\tilde{b}_t &= (1+r)\tilde{b}_{t-1} + \frac{X}{B}(\tilde{n}_t^x - \tilde{n}_t^y)
\end{aligned}$$

This system is identical to the original where we have set $i_t = q_t = 0$ and dropped the market clearing condition in the intermediate goods market.

B.2 ALTERNATIVE WAGE SETTING SPECIFICATIONS

B.2.1 PREFERENCES WITH WEALTH EFFECTS IN LABOR SUPPLY

In our benchmark specification for the wage setting equation we assumed that the marginal rate of substitution between consumption and hours worked is independent of consumption (as is the case with GHH preferences). In this section we explore the consequences of moving away from this assumption for our econometric procedure in Section 2.5. For a general set of preferences represented by $u(c, n)$, we can write the marginal rate of substitution in log-deviations from steady state as,

$$\begin{aligned} mrs_{kt} &= \left(\frac{u_{cn}c}{u_n} - \frac{u_{cc}c}{u_c} \right) c_{kt} + \left(\frac{u_{nn}n}{u_n} - \frac{u_{nc}c}{u_n} \right) n_{kt} \\ &\equiv \omega c_{kt} + \left(\omega + \frac{1}{\hat{\phi}} \right) n_{kt} \end{aligned}$$

which nest the special case with no wealth effects ($\omega = 0$) and so we obtain the marginal rate of substitution from our benchmark specification. The aggregate and state level wage setting equations become,

$$\begin{aligned} \pi_t^w &= \hat{\lambda}(\pi_t + \epsilon_t - \epsilon_{t-1} + \left(\frac{1}{\hat{\phi}} + \omega\right)(n_t - n_{t-1}) + \omega(c_t - c_{t-1})) + (1 - \hat{\lambda})\pi_{t-1}^w \\ \tilde{w}_{kt} &= \hat{\lambda}(\tilde{p}_{kt} + \epsilon_{kt} + \left(\frac{1}{\hat{\phi}} + \omega\right)\tilde{n}_{kt} + \omega\tilde{c}_{kt}) + (1 - \hat{\lambda})\tilde{w}_{kt-1} \end{aligned}$$

Replacing aggregate consumption with the model implied $w_t + n_t - p_t$ we obtain

$$\pi_t^w = \lambda\pi_t + \frac{\lambda}{\hat{\phi}}(n_t - n_{t-1}) + (1 - \lambda)\pi_{t-1}^w + \frac{\lambda}{1 - \omega}(\epsilon_t - \epsilon_{t-1})$$

where $\lambda \equiv \frac{\hat{\lambda}(1-\omega)}{1-\lambda\omega}$ and $\frac{1}{\hat{\phi}} \equiv \frac{\frac{1}{\hat{\phi}} + 2\omega}{1-\omega}$. Also, we can re-write the state level equation as

$$\tilde{w}_{kt} = \lambda\tilde{p}_{kt} + \frac{\lambda}{\hat{\phi}}\tilde{n}_{kt} + (1 - \lambda)\tilde{w}_{kt-1} + \frac{\lambda\omega}{1 - \omega}(\tilde{p}_{kt} + \tilde{c}_{kt} - (\tilde{w}_{kt} + \tilde{n}_{kt})) + \frac{\lambda}{1 - \omega}\tilde{\epsilon}_{kt} \quad (\text{B.1})$$

Since state level economies are open economies, in general, the term $\tilde{p}_{kt} + \tilde{c}_{kt} - (\tilde{w}_{kt} + \tilde{n}_{kt})$ will be different from zero. By omitting it in our cross-sectional regressions we could be obtaining biased estimates of λ, ϕ .

B.2.2 FORWARD LOOKING WAGES

In our benchmark specification for the wage setting equation we assumed that there was no forward looking term in the target wage. In this section we explore the consequences of having a forward looking component in the wage setting equation for our econometric procedure in Section 2.5. In particular, consider the aggregate and state level wage setting equations

$$\begin{aligned}\pi^w_t &= \lambda(\pi_t + \epsilon_t - \epsilon_{t-1} + \frac{1}{\phi}(n_t - n_{t-1})) + \lambda\kappa\mathbb{E}_t[\pi^w_{t+1}] + (1 - \lambda)\pi^w_{t-1} \\ \tilde{w}_{kt} &= \lambda(\tilde{p}_{kt} + \tilde{\epsilon}_{kt} + \frac{1}{\phi}\tilde{n}_{kt}) + \lambda\kappa\mathbb{E}_t[\tilde{w}_{kt+1} - \tilde{w}_{kt}] + (1 - \lambda)\tilde{w}_{kt-1}\end{aligned}$$

where κ parameterizes the importance of the forward looking term. Also, let's consider the case where local wages follow an AR(1) process in equilibrium with coefficient $\tilde{\rho}_w$ and aggregate expected wage inflation is zero. Our model from Section 2.4, would imply this, for instance, when $\theta \rightarrow 1$ so that $\tilde{w}_{kt} = \tilde{z}_{kt}^x$ in equilibrium and $\tilde{\rho}_w = \rho_x$; and the monetary authority fully stabilizes expected aggregate nominal wage growth. We obtain,

$$\begin{aligned}\pi^w_t &= \lambda(\pi_t + \epsilon_t - \epsilon_{t-1} + \frac{1}{\phi}(n_t - n_{t-1})) + (1 - \lambda)\pi^w_{t-1} \\ \tilde{w}_{kt} &= \frac{\lambda}{1 + \lambda\kappa(1 - \tilde{\rho}_w)}(\tilde{p}_{kt} + \tilde{\epsilon}_{kt} + \frac{1}{\phi}\tilde{n}_{kt}) + \frac{1 - \lambda}{1 + \lambda\kappa(1 - \tilde{\rho}_w)}\tilde{w}_{kt-1}\end{aligned}$$

Then, we can write,

$$\lambda = \frac{1 - \beta_w}{1 + \beta_w\kappa(1 - \tilde{\rho}_w)}$$


where $\beta_w \equiv \frac{1 - \lambda}{1 + \lambda\kappa(1 - \tilde{\rho}_w)}$. From this expression we see that our estimates for λ using cross-state variation are upward biased. However, we can get a notion on the magnitude of the bias by asking what would $\kappa(1 - \tilde{\rho}_w)$ have to be in order for λ to be less than some λ_0 . We obtain,

$$\kappa(1 - \tilde{\rho}_w) > \frac{1 - \beta_w - \lambda_0}{\beta_w\lambda_0}$$

For example, given our lower estimate for $\beta_w = 0.5$, in order for λ to be below 0.1 we would need a $\kappa(1 - \tilde{\rho}_w)$ larger than 8.


B.3 FIGURES AND TABLES

Figure B.1: Nielsen Retail Price Index vs. CPI Food Price Index




Note: In this figure, we compare our monthly retail scanner price index for the U.S. as a whole (dashed line) to the CPI's aggregate monthly food price index (solid line). We normalize both indices to 1 in January of 2006.

Figure B.2: The Evolution of Aggregate Real and Nominal Composition Adjusted Wages




Note: Figure shows the evolution of aggregate real and nominal log wages within the U.S. between 2000 and 2012 using data from the Current Population Survey. The sample is restricted to only males between the ages of 21 and 55, who are currently employed, who report usually working 30 hours per week, and who worked at least 48 weeks during the prior 12 months. As discussed in the text, we adjust wages for the changing labor market condition over time by controlling for age, race, education, and usual hours worked. We compute real wages by deflating our nominal wage index by the CPI-U of the corresponding year.

Figure B.3: Impulse Response to a Discount rate Shock




Note: Figure shows the impulse response to a one standard deviation discount rate shock. The horizontal axis are years after the shock.

Figure B.4: Impulse Response to a Productivity / Markup Shock



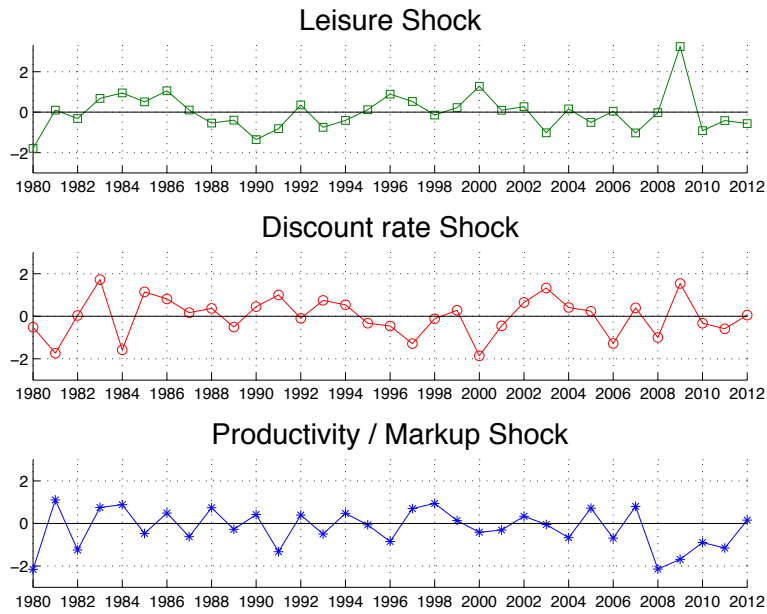
Note: Figure shows the impulse response to a one standard deviation productivity/markup shock. The horizontal axis are years after the shock.

Figure B.5: Impulse Response: Leisure Shock




Note: Figure shows the impulse response to a one standard deviation leisure shock. The horizontal axis are years after the shock.

Figure B.6: Shock time-series




Note: Figure shows the estimated aggregate shock realizations from 1980 to 2012.

Figure B.7: Counterfactual Price Response




Note: Figure shows the cumulative response of Prices when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Figure B.8: Counterfactual Wage Response



Note: Figure shows the cumulative response of Wages when we feed the VAR with the sequence of shocks between 2008 and 2012; one at a time.

Figure B.9: State Net Migration Rate 2009-2010 vs. State Employment Growth 2007-2010



Note: Figure shows state net migration rate between 2009 and 2010 against employment growth in the state during 2007-2010. Employment growth comes from the BLS and is defined in the text. State net migration rates come from American Community Survey.

Table B.1: Calibration

	Value	Target
α	0.62	Aggregate labor share in the non-tradable sector in 2006
θ	0.55	Aggregate labor share in the tradable sector in 2006
β	0.16	Aggregate labor share in 2006 (0.61). See on-line appendix
σ	2	Intertemporal elasticity of substitution
ϕ	2	Frisch elasticity. See Section 2.6.2 for estimation
λ	0.7	Wage stickiness. See Section 2.6.2 for estimation
φ_p	1.5	Taylor Rule from Galí (2010)
φ_y	0.5	Taylor Rule from Galí (2010)
Φ_0	0.01	Trade balance-output ratio volatility = 2 from Mendoza (1991).
R	0.03	Real interest rate
X	0.32	Intermediate inputs over output ratio US 2006
B	2.1	Median net worth to output ratio in the US 2006
ρ_γ	0.9	Persistence of discount rate shock
ρ_z	0.76	Persistence of productivity shock from the model. See on-line appendix
ρ_ϵ	0.66	Persistence of leisure shock from the model. See on-line appendix

Table B.2: Aggregate v. Regional Employment Impact Elasticities

	ρ_γ	0.9	0.9	0.9	0.9	0.6	0.1		
	β	0.16	0	0.16	0.16	0.16	0.16		
	φ_p	1.5	1.5	1.5	1.31	1.5	1.5		
	φ_y	0.5	0.5	0	0	0.5	0.5		
	(λ, ϕ)	(0.7, 2)	(0.7, 1)	(0.5, 2)	(0.7, 2)	(0.7, 2)	(0.7, 2)		
Aggregate	γ	0.77	0.53	1.29	0.74	1.11	1.76	0.71	0.31
	z	0.29	0.20	-0.35	0.28	-0.24	-1.07	0.29	0.29
	ϵ	-0.65	-0.51	-0.28	-0.69	-0.26	0.41	-0.65	-0.65
Regional	γ	1.76	1.26	2.95	2.85	1.76	1.76	0.36	0.03
	z^y	0.44	0.38	-0.18	0.10	0.44	0.44	0.44	0.44
	z^x	0.08	-0.06	0.23	0.05	0.08	0.08	0.08	0.08
	ϵ	-0.64	-0.52	-0.16	-0.43	-0.64	-0.64	-0.64	-0.64

Note: The table summarizes the response of employment on impact to each of the model shocks under our base calibration (column 1) and alternate calibrations (columns 2-8). The rows represent the employment response to different aggregate shocks (top three rows) and different local shocks (bottom three rows). Columns 2 and 3, explore the elasticities under different calibrations of wage stickiness and the Frisch elasticity of labor supply. Column 4 examines the robustness to changes in the tradable share of the intermediate good. Columns 5 and 6 examine the results under alternate Taylor Rule parameters. The final two columns change the persistence of the demand shock. The units are percentage deviations from the steady state, in the case of aggregate employment, and percentage deviations from the aggregate in the case of regional employment.

Table B.3: Aggregate Elasticities ($\lambda=0.7$, $\phi=2$)

		n	π	π^w	w^r
Short run	γ	0.77	1.64	1.42	-0.22
	z	0.29	-2.73	-1.81	0.92
	ϵ	-0.65	0.94	1.13	0.19
Long run	γ	0.56	1.01	1.03	-0.16
	z	0.12	-0.66	-0.79	0.40
	ϵ	-0.18	0.35	0.32	0.05

Note: The table summarizes the response of each aggregate variable to the shocks in percentage deviations from the steady state. The “short” elasticity is the response at date $t = 0$. The “long” elasticity is the response after 5 years.

Table B.4: Regional Elasticities ($\lambda=0.7$, $\phi=2$)

		n	π	π^w	w^r
Short run	γ	1.76	2.54	2.40	-0.14
	z^y	0.44	-2.00	-1.25	0.75
	z^x	0.08	0.29	0.31	0.02
	ϵ	-0.64	0.71	0.97	0.26
Long run	γ	0.12	-0.28	-0.30	0.19
	z^y	0.63	0.31	0.23	0.21
	z^x	-0.06	-0.03	-0.04	0.04
	ϵ	-0.44	-0.12	-0.20	0.14

Note: The table summarizes the response of each island level variable to the shocks in percentage deviations from the steady state. The “short” elasticity is the response at date $t = 0$. The “long” elasticity is the response after 5 years.

Table B.5: Discount/Interest rate (γ) and Productivity/Markup (z) shocks' contribution to aggregate employment change

λ	2008 to 2009						2008 to 2012				
	ϕ						ϕ				
		0.5	1	2	3	4	0.5	1	2	3	4
0.1	γ	103	83	108	107	108	40	-34	46	57	52
	z	-3	22	-3	-1	-2	47	126	48	38	43
0.3	γ	47	66	51	29	101	-13	-25	-33	-20	9
	z	2	16	45	71	1	98	123	134	121	92
0.5	γ	8	36	31	13	92	-1	-13	-19	2	47
	z	6	16	48	74	0	94	123	133	111	65
0.7	γ	0	0	30	21	11	47	48	-12	-12	0
	z	3	33	33	53	69	53	72	136	135	122
0.9	γ	0	0	0	2	3	45	47	41	15	9
	z	-1	24	54	64	69	58	79	91	114	118

Note: Table shows the percent contribution of the demand and supply shocks to the aggregate employment change implied by our procedure for different combinations of the parameters. For a given pair $\{\phi, \lambda\}$, the ' γ ' entry corresponds to the demand shock. The ' z ' entry to the supply shock. The percent contribution of the leisure shock can be calculated by subtracting the sum of both entries from 100. Entries with * are such that no decomposition of the shocks satisfy the identification restrictions for those parameter values.

Appendix C

FROM HYPERINFLATION TO STABLE PRICES: ARGENTINA'S EVIDENCE ON MENU COST MODELS

C.1 PROOFS OF COMPARATIVE STATIC RESULTS

C.1.1 PROOF OF PROPOSITION 3

First we establish that the value function, the optimal adjustment function and the inaction sets are all symmetric. We show that:

$$\begin{aligned} V(\hat{p} + p^*(z), z; \pi, \sigma^2) &= V(-\hat{p} + p^*(-z), -z; -\pi, \sigma^2) + v(z) , \\ \hat{\psi}(z; -\pi, \sigma^2) &= -\hat{\psi}(-z; \pi, \sigma^2), \text{ and} \\ (\hat{p} + p^*(z), z) \in \mathcal{I}(\pi, \sigma^2) &\implies (-\hat{p} + p^*(-z), -z) \in \mathcal{I}(-\pi, \sigma^2) \end{aligned}$$

for all $z \in [0, \bar{z}]$, $\hat{p} \geq 0$ and $\pi \in (-\bar{\pi}, \bar{\pi})$. The symmetry of these three objects can be established using a guess and verify argument in the Bellman equation. This argument has two parts, one deals with the instantaneous return and the second with the probabilities of different paths of z 's. First we show that the instantaneous return satisfy the analogous property that the required symmetry property for the value function stated above. We note that for $t \leq \tau$, where we let 0 the time where prices were last set, we have:

$$\begin{aligned} F(p(t) - \bar{p}(t), z(t)) &= F(p(0) - \bar{p}(0) - \pi t - p^*(z(t)) + p^*(z(t)), z(t)) \\ &= F(-p(0) + \bar{p}(0) + \pi t - p^*(-z(t)) + p^*(-z(t)), -z(t)) + f(z(t)) \\ &= F(-p(0) + \bar{p}(0) + \pi t + p^*(z(t)) + p^*(-z(t)), -z(t)) + f(z(t)) , \end{aligned}$$

where the second equality holds by symmetry of $F(\cdot)$ setting $\hat{p}(t) = p(0) - \bar{p}(0) - p^*(z(t)) - \pi t$. Thus fixing the path of $\{z(t)\}$ for $0 \leq t \leq \tau$, starting with $p(0) - \bar{p}(0)$ and $z(0)$ and having

inflation π , gives the same profits, assuming symmetry of $F(\cdot)$, than starting with $-p(0)+\bar{p}(0)$ and having inflation $-\pi$ and $-z(0)$. Finally the probability of the path $\{z(t)\}$ for $t \in [0, \tau]$ conditional on $z(0)$, given the symmetry of $a(\cdot)$ and $b(\cdot)$ is the same as the one for the path $\{-z(t)\}$ conditional on $-z(0)$. From here one obtains that the inaction set is symmetric. Likewise, from this property it is easy to see that the optimal adjustment is also symmetric. If with inflation π a firm adjust with current shock z setting $p = \bar{p} + \hat{\psi}(z; \pi, \sigma^2)$, then with inflation $-\pi$ and current shock $-z$ it will adjust to $p = \bar{p} + \hat{\psi}(-z, -\pi) = \bar{p} - \hat{\psi}(z, -\pi)$. To see this, let $t = 0$ be a date where an adjustment take place, let $p(0)$ the price right after the adjustment, and let τ the stopping time until the next adjustment. The value of $p(0)$ maximizes

$$\begin{aligned}
p(0) &= \arg \max_{\tilde{p}} \mathbb{E} \left[\int_0^\tau e^{-rt} F(\tilde{p} - \bar{p}(0) - \pi t, z(t)) \mid z(0) \right] \\
&= \arg \max_{\tilde{p}} \mathbb{E} \left[\int_0^\tau e^{-rt} F(\tilde{p} - \bar{p}(0) - \pi t - p^*(z(t)) + p^*(z(t)), z(t)) \mid z(0) \right] \\
&= \arg \max_{\tilde{p}} \mathbb{E} \left[\int_0^{\tau'} e^{-rt} F(-\tilde{p} + \bar{p}(0) + \pi t - p^*(-z(t)) + p^*(-z(t)), -z(t)) \mid -z(0) \right] \\
&= \arg \max_{\tilde{p}} \mathbb{E} \left[\int_0^{\tau'} e^{-rt} F(-\tilde{p} + \bar{p}(0) + \pi t, -z(t)) \mid -z(0) \right].
\end{aligned}$$

where τ' is the stopping time obtained from τ but defined flipping the sign of the z 's.

Given the symmetry of the inaction set and optimal adjustment it is relatively straightforward to establish the symmetry of the expected time to adjustment \mathcal{T} and of the invariant density g . With the \mathcal{T} and g symmetric, it is immediate to establish that λ_a is symmetric. Finally, if λ_a is differentiable, then $\frac{\partial}{\partial \pi} \lambda_a(\pi, \sigma^2) = -\frac{\partial}{\partial \pi} \lambda_a(-\pi, \sigma^2)$, which establish part (i) of the proposition.

Now we show that inflation has no first order effect of $\bar{\sigma}$, i.e., we establish part (ii). For that we first use that the symmetry of the decision rules and of the invariant distribution of the shocks (which follows from the symmetry of a and b) implies that $h(\hat{p}, \pi) = h(-\hat{p}, -\pi)$ where for simplicity we suppress σ^2 as an argument. Differentiating this expression with respect to π and evaluating at $\pi = 0$ we obtain: $h_\pi(\hat{p}, 0) = -h_\pi(-\hat{p}, 0)$. Let $f(\hat{p}, \pi)$ be any symmetric differentiable function in the sense that $f(\hat{p}) = f(-\hat{p})$. Then writing the expected value of f as $\mathbb{E}[f|\pi] = \int_{-\infty}^0 f(\hat{p})h(\hat{p}, \pi)d\hat{p} + \int_0^\infty f(\hat{p})h(\hat{p}, \pi)d\hat{p}$ and differentiating both terms with respect to π and evaluating it at $\pi = 0$, using the implications for symmetry for the

derivatives of h and the symmetry of f we have $\frac{\partial}{\partial \pi} \mathbb{E}[f|0] = 0$. Applying this to $f(\hat{p}) = \hat{p}^2$ we obtain that inflation does not have a first order effect on the second non-centered moment of the relative prices. Finally, to examine the effect of inflation on the variance of the relative prices, we need to examine the effect of inflation on the square of the average relative price,

$$\frac{\partial}{\partial \pi} \left[\int_{-\infty}^{\infty} \hat{p} h(\hat{p}, \pi) d\hat{p} \right]^2 \Big|_{\pi=0} = 2 \left[\int_{-\infty}^{\infty} \hat{p} h(\hat{p}, 0) d\hat{p} \right] \left[\int_{-\infty}^{\infty} \hat{p} \frac{\partial}{\partial \pi} h(\hat{p}, 0) d\hat{p} \right] = 0,$$

since by symmetry of $h(\cdot, 0)$ around $\hat{p} = 0$ we have $\int_{-\infty}^{\infty} \hat{p} h(\hat{p}, 0) d\hat{p} = 0$. Then, we have shown that inflation has no first order effect on the variance of relative prices around $\pi = 0$.

Finally, observe that by symmetry we have that for any inflation rate $\lambda_a^+(\pi) = \lambda_a^-(-\pi)$ and $\Delta_p^+(\pi) = \Delta_p^-(-\pi)$. Differentiating with respect to inflation yields $\frac{\partial \lambda_a^+(\pi)}{\partial \pi} = -\frac{\partial \lambda_a^-(-\pi)}{\partial \pi}$ and $\frac{\partial \Delta_p^+(\pi)}{\partial \pi} = -\frac{\partial \Delta_p^-(-\pi)}{\partial \pi}$. Evaluating it at $\pi = 0$ gives the result. Q.E.D.

C.1.2 PROOF OF LEMMA 6

Benabou and Konieczny (1994) compute the value of following an sS policy assuming that the period return function $F(p, 0)$ is cubic in terms of deviations from the profit maximizing price, i.e. $p - p^*$. This allows for explicit computation of the value of the policy and to obtain the first order conditions at s and S for any value of π . Adding equations (8) and (14) in Benabou and Konieczny (1994) we get the expression $S - s = 2\delta + \frac{2}{3} \left(\frac{r}{\pi} - a \right) \delta^2$ for $\delta = \left(\frac{3}{2} \frac{c\pi}{-F''} \right)^{\frac{1}{3}}$ and $a = -\frac{F'''}{2F''}$. Thus,

$$S - s = 2 \left(-\frac{3}{2} \frac{c}{F''} \right)^{\frac{1}{3}} \pi^{\frac{1}{3}} + \frac{2}{3} \left(-\frac{3}{2} \frac{c}{F''} \right)^{\frac{2}{3}} r \pi^{-\frac{1}{3}} - \frac{2}{3} \left(-\frac{3}{2} \frac{c}{F''} \right)^{\frac{2}{3}} a \pi^{\frac{2}{3}}$$

and

$$\frac{d(S - s)}{d\pi} \frac{\pi}{S - s} = \omega_1 \frac{1}{3} + \omega_2 \frac{-1}{3} + \omega_3 \frac{2}{3},$$

where $\omega_1 = 2 \left(-\frac{3}{2} \frac{c}{F''} \right)^{\frac{1}{3}} \pi^{\frac{1}{3}} / (S - s)$, $\omega_2 = \frac{2}{3} \left(-\frac{3}{2} \frac{c}{F''} \right)^{\frac{2}{3}} r \pi^{-\frac{1}{3}} / (S - s)$,

and $\omega_3 = -\frac{2}{3} \left(-\frac{3}{2} \frac{c}{F''} \right)^{\frac{2}{3}} a \pi^{\frac{2}{3}} / (S - s)$. Using L'hospital's rule $\lim_{c \rightarrow 0} \omega_1 = 1$, and $\lim_{c \rightarrow 0} \omega_2 = \lim_{c \rightarrow 0} \omega_3 = 0$, so that $\frac{d(S-s)}{d\pi} \frac{\pi}{S-s} = 1/3$.

The same argument for $\mathcal{T} = (S - s)/\pi$ yields

$$\mathcal{T} = 2 \left(-\frac{3}{2} \frac{c}{F''} \right)^{\frac{1}{3}} \pi^{-\frac{2}{3}} + \frac{2}{3} \left(-\frac{3}{2} \frac{c}{F''} \right)^{\frac{2}{3}} r \pi^{-\frac{4}{3}} - \frac{2}{3} \left(-\frac{3}{2} \frac{c}{F''} \right)^{\frac{2}{3}} a \pi^{-\frac{1}{3}}$$

As $\lambda_a = \mathcal{T}^{-1}$, taking limits as $c \rightarrow 0$ yields $\lim_{c \rightarrow 0} \frac{d\lambda_a}{d\pi} \frac{\lambda_a}{\pi} = \frac{2}{3}$. QED.

C.1.3 PROOF OF LEMMA 7

First notice that if in the problem described in equation (3.2) we multiply r, π, σ^2 and the functions $a(\cdot), F(\cdot)$ by a constant $k > 0$, we are just changing the units at which we measure time. Moreover, the objective function in the right hand side of equation (3.2) is homogeneous of degree one in $F(\cdot)$ and c and, hence, the policy function is the same whether we multiply $F(\cdot)$ by k or divide c by it. Thus, if the function $a(\cdot)$ equals zero, λ_a is homogeneous of degree one in $(\pi, \sigma^2, r, \frac{1}{c})$. Using this homogeneity:

$$\frac{1}{\lambda_a(1, \frac{\sigma^2}{\pi}, \frac{r}{\pi}, \frac{1}{\pi c})} \frac{\partial \lambda_a(1, \frac{\sigma^2}{\pi}, \frac{r}{\pi}, \frac{1}{\pi c})}{\partial \pi} = \frac{\pi}{\lambda_a(\pi, \sigma^2, r, \frac{1}{c})} \frac{\partial \lambda_a(\pi, \sigma^2, r, \frac{1}{c})}{\partial \pi}$$

Taking limits for $\sigma > 0$:

$$\begin{aligned} \lim_{r \downarrow 0, c \downarrow 0} \frac{1}{\lambda_a(1, \frac{\sigma^2}{\pi}, \frac{r}{\pi}, \frac{1}{\pi c})} \frac{\partial \lambda_a(1, \frac{\sigma^2}{\pi}, \frac{r}{\pi}, \frac{1}{\pi c})}{\partial \pi} &= \frac{1}{\lambda_a(1, \frac{\sigma^2}{\pi}, 0, \infty)} \frac{\partial \lambda_a(1, \frac{\sigma^2}{\pi}, 0, \infty)}{\partial \pi} \\ &= \lim_{r \downarrow 0, c \downarrow 0} \frac{\pi}{\lambda_a(\pi, \sigma^2, r, \frac{1}{c})} \frac{\partial \lambda_a(\pi, \sigma^2, r, \frac{1}{c})}{\partial \pi} = \frac{\pi}{\lambda_a(\pi, \sigma^2, 0, \infty)} \frac{\partial \lambda_a(\pi, \sigma^2, 0, \infty)}{\partial \pi} \end{aligned}$$

Thus

$$\lim_{\pi \rightarrow \infty} \left[\lim_{r \downarrow 0, c \downarrow 0} \frac{\frac{\partial \lambda_a(1, \frac{\sigma^2}{\pi}, \frac{r}{\pi}, \frac{1}{\pi c})}{\partial \pi}}{\lambda_a(1, \frac{\sigma^2}{\pi}, \frac{r}{\pi}, \frac{1}{\pi c})} \right]_{\sigma^2 > 0} = \lim_{\sigma \rightarrow 0} \left[\lim_{r \downarrow 0, c \downarrow 0} \frac{\frac{\partial \lambda_a(1, \frac{\sigma^2}{\pi}, \frac{r}{\pi}, \frac{1}{\pi c})}{\partial \pi}}{\lambda_a(1, \frac{\sigma^2}{\pi}, \frac{r}{\pi}, \frac{1}{\pi c})} \right]_{\pi > 0}.$$

For the same reason all the elements of Ψ for a given z are homogeneous of degree 0 so, equation (3.10b) follows.

Finally we establish equation (3.10c). Using the same notation we index the invariant density as follows $g(p - \bar{p}, z; \pi, \sigma^2, r, 1/c)$. Note that since $a(z) = 0$, scaling the four parameters scales the units of time of z . Moreover, scaling the four parameters does not change Ψ . Thus for each $(p - \bar{p}, z)$ the invariant density is homogeneous of degree zero in $(\pi, \sigma^2, r, 1/c)$. The result follows from the definition of Δ_p^+ and Δ_p^- in terms of g and Ψ . QED.

C.1.4 PROOF OF PROPOSITION 5

Let $x = p - z$ be the log of the real gross mark-up, which we refer to as the net markup. The firm's optimal pricing policy, in this case, can be characterized in terms of three constants

$X \equiv (\underline{x}, \bar{x}, \hat{x})$. The policy function takes the simple form $\Psi(z) = X + z$. We can write the inaction set in terms of the net markups as $\mathcal{I} = \{x : \underline{x} < x < \bar{x}\}$. It is optimal to keep the price unchanged when the net markup x is in the interval (\underline{x}, \bar{x}) . When prices are not changed, the real markup evolves according to $dx = -\pi dt + \sigma dW$. When the real markup hits either of the two thresholds, prices are adjusted so that the real markup is \hat{x} and thus the optimal return relative price is $\hat{\psi}(z) = \hat{x} + z$. Price increases are equal to $\Delta_p^+ = \hat{x} - \underline{x}$ and price decreases are equal to $\Delta_p^- = \bar{x} - \hat{x}$. In the case where $\sigma = 0$ and $\pi > 0$, we obtain a version of Sheshinski and Weiss's (1977) model, and the optimal policy can be characterized simply by two thresholds $\underline{x} < \hat{x}$.

We present a series of lemmas that yield the proof of Proposition 5. First, Lemma 8 simplifies the Hamilton-Jacobi-Bellman for the undiscounted case. It shows that $\lim_{r \rightarrow 0} rV(x, r)$ in problem 3.2 is a constant. Second, Lemma 9 represents the value function as a power series and finds its coefficients as functions of π and σ . Lemma 10 is an analytical solution for the zero inflation case that characterizes $(\underline{x}, \bar{x}, \hat{x})$, the size of price changes and the frequency of price increases and of price decreases. Lemma 11 characterizes the derivatives of these elements with respect to inflation at $\pi = 0$. Finally, Lemma 12 is a complete analytical characterization of Sheshinski and Weiss's (1977) case.

Lemma 8 *The limit as $r \downarrow 0$, of the value function and of the thresholds $\underline{x}, \hat{x}, \bar{x}$ can be obtained by solving for a constant A and a function $v : [\underline{x}, \bar{x}] \rightarrow \mathbb{R}$ which satisfy:*

$$A = Bx^2 - \pi v'(x) + \frac{\sigma^2}{2} v''(x) \text{ for all } x \in [\underline{x}, \bar{x}] \quad (\text{C.1})$$

and the boundary conditions:

$$v(\bar{x}) = v(x^*) + c,$$

$$v(\underline{x}) = v(x^*) + c$$

$$v'(\bar{x}) = 0,$$

$$v'(\underline{x}) = 0,$$

$$v'(\hat{x}) = 0.$$

Remark: $v(\hat{x})$ can be normalized to zero, and A has the interpretation of the expected profits per unit of time net of the expected cost of changing prices.

Proof. The solution of the firm's problem (3.2) can be characterized by the equation

$$rV(x_0, r) = F(x) - \pi V'(x, r) + \frac{\sigma^2}{2} V''(x, r) \text{ for all } x \in [\underline{x}, \bar{x}] \quad (\text{C.2})$$

and the following boundary conditions: two value matching conditions $V(\hat{x}, r) - V(\bar{x}, r) = V(\hat{x}, r) - V(\underline{x}, r) = c$ and the smooth pasting and optimal return conditions $V'(\hat{x}, r) = V'(\bar{x}, r) = V'(\underline{x}, r) = 0$.

Let $v(x) = \lim_{r \rightarrow 0} V(x, r)$ for $x \in [\underline{x}, \bar{x}]$ and $v'(x) = \lim_{r \rightarrow 0} V'(x, r)$ for $x \in [\underline{x}, \bar{x}]$.

We show that when $r \rightarrow 0$ the left hand side of equation (C.2) is a constant —i.e. $\lim_{r \rightarrow 0} rV(x, r) = A$ —so it becomes equation (C.1). Write $V(x, r)$ as $V(x, r) = V_1(x, r) + V_2(x, r)$, where the first term is the present value of expected profits and the second is the present value of the expected adjustment costs. Multiplying the former by r we get

$$rV_1(x_0, r) = \lim_{T \rightarrow \infty} E \left[\int_0^T r e^{-rt} F(x(t)) dt \mid x = x_0 \right]$$

for some profit function $F(x(t))$. Observe that rV_1 is a weighted average of $F(x(t))$ with positive weights, $r e^{-rt}$ that satisfy $\int_0^T r e^{-rt} dt = 1$. As $r \rightarrow 0$, the terms $r e^{-rt}$ become a constant, which has to be $1/T$ to still integrate to 1. Then,

$$\lim_{r \rightarrow 0} E \left[\lim_{T \rightarrow \infty} \int_0^T r e^{-rt} F(x(t)) dt \mid x = x_0 \right] = \lim_{T \rightarrow \infty} \int_0^\infty \frac{1}{T} E [F(x(t)) \mid x = x_0] dt.$$

If the path of $x(t)$ that solves the firm's problem (3.2) is ergodic then the average $E [F(x(t)) \mid x = x_0]$ is independent of the state so

$$\lim_{T \rightarrow \infty} \int_0^\infty \frac{1}{T} E [F(x(t)) \mid x = x_0] dt \rightarrow E [F(x)] \text{ for all } x(0) \in [\underline{x}, \bar{x}].$$

Now write the second part of the value function in problem (3.2), the expected costs of price adjustments, as $V_2(x_0, r) = \lim_{T \rightarrow \infty} E \left[\sum_{i=1}^{N(T)} e^{-r\tau_i} c \mid x = x(0) \right]$, where τ_i is the time of each price adjustment and $N(T) = \max[i : \tau_i \leq T]$ is the number of price adjustments before T . Letting τ_N be the time of the N^{th} adjustment, we can write $V_2(x, r)$ as $V_2(x_0, r) = E \left[\sum_{i=1}^N e^{-r\tau_i} c \mid x = x_0 \right] + E \left[e^{-r\tau_N} \sum_{i=N+1}^\infty e^{-r\tau_i} c \mid x = x_0 \right]$. Multiplying both terms by r , noticing that immediately after an adjustment x reverts to the reset value \hat{x} , adding and

subtracting $V(\hat{x}, r)$ on the left side of the equality, and collecting terms yields

$$rV_2(x_0, r) = \frac{r}{1 - E[e^{-r\tau_N} | x = x_0]} E \left[\sum_{i=1}^N e^{-r\tau_i} c | x = x(0) \right] - \frac{r}{1 - E[e^{-r\tau_N} | x = x_0]} [V_2(x_0, r) - V_2(\hat{x}, r)].$$

Taking limits as $r \rightarrow 0$, we get $\lim_{r \rightarrow 0} \sum_{i=1}^N e^{-r\tau_i} = N$ and $\lim_{r \rightarrow 0} \frac{r}{1 - E[e^{-r\tau_N} | x = x_0]} = \frac{1}{E[\tau_N | x = x_0]}$ so that

$$\lim_{r \rightarrow 0} rV_2(x_0, r) = c \frac{N}{E[\tau_N | x = x_0]} - \frac{1}{E[\tau_N | x = x_0]} [V_2(x_0, r) - V_2(\hat{x}, r)]$$

For $N \rightarrow \infty$, the first term in the right hand side converges to λ_a by the strong law of large numbers of renewal theory, and the second term vanishes since $|V_2(x_0, r) - V_2(\hat{x}, r)| \leq 1$ for all $r > 0$, so that

$$\lim_{r \rightarrow 0} rV_2(x_0, r) = \lambda_a \text{ for all } x_0$$

Equation (C.1) together with its boundary conditions is a functional equation to find a constant A , the values $(\underline{x}, \hat{x}, \bar{x})$ and a twice differentiable function $v : [\underline{x}, \bar{x}] \rightarrow \mathbb{R}$. Moreover, A is the expected profit net of the expected cost of adjustment—i.e. $A = E[F(x)] - \lambda_a c$.

In the case with $\sigma = 0$, the integral $\lim_{T \rightarrow \infty} \int_0^{\infty} \frac{1}{T} F(x(t)) dt$ converges to a value that is independent of $x(0)$ for the path of $x(t)$ that solves the firm's problem and a similar argument applies. Analogously for the expected costs of price adjustment V_2 .

■

Lemma 9 *Given \underline{x} and \bar{x} , the function v described in Lemma 8 is the power series*

$$v(x) = \sum_{i=1}^{\infty} \alpha_i x^i. \tag{C.3}$$

In the case with $\pi = 0$ and $\sigma > 0$ the coefficients are

$$\alpha_2 = \frac{A}{\sigma^2}, \quad \alpha_4 = -\frac{1}{6} \frac{B}{\sigma^2} \text{ and } \alpha_i = 0 \text{ for all } i \neq 2, 4$$

In the case with $\pi > 0$ and $\sigma = 0$, the coefficients are

$$\alpha_1 = -\frac{A}{\pi}, \quad \alpha_3 = \frac{1}{3} \frac{B}{\pi} \quad \text{and } \alpha_i = 0 \text{ for } i = 2 \text{ and for } i \geq 4.$$

Proof. We look for the coefficients of equation (C.3) that solve equation (C.1) for all $x \in [\underline{x}, \bar{x}]$. Substituting for v' and for v'' in equation (C.1) and matching coefficients yields the results. When $\sigma > 0$, the conditions for matching coefficients are

$$\begin{aligned} A &= -\alpha_1 \pi + \alpha_2 \sigma^2 \\ \alpha_3 &= \frac{1}{3} \left(\frac{2\pi}{\sigma^2} \right) \alpha_2 \\ \alpha_4 &= -\frac{1}{6} \frac{B}{\sigma^2} + \frac{1}{12} \left(\frac{2\pi}{\sigma^2} \right)^2 \alpha_2 \\ \alpha_i &= \left[-2 \frac{B}{\sigma^2} + \left(\frac{2\pi}{\sigma^2} \right)^2 \alpha_2 \right] \frac{2}{i!} \left(\frac{2\pi}{\sigma^2} \right)^{i-4} \quad \text{for } i \geq 5. \end{aligned}$$

For $\pi = 0$ we have that $\alpha_i = 0$ for all i except for α_2 and α_4 . For $i = 3$ and for $i \geq 5$ this follows directly from the conditions above. $\alpha_1 = 0$ is a condition for the symmetry of v when $\pi = 0$. ■

Finding $v(x)$ then requires to solve a two dimensional problem in α_1 and α_2 when $\sigma > 0$ (or in α_1 and α_3 when $\sigma = 0$) using the boundary conditions in Lemma 8.

Lemma 10 *The solution for the thresholds when $\pi = 0$ is $\hat{x} = 0$, $\bar{x} = -\underline{x} = \left(6c \frac{\sigma^2}{B}\right)^{\frac{1}{2}}$ and the constant $A = \left(\frac{2}{3} B c \sigma^2\right)^{\frac{1}{2}}$.*

Proof. We will use the normalization $v(\hat{x}) = 0$, the smooth pasting condition $v'(\bar{x}) = 0$ and the value matching conditions $v(\bar{x}) = v(\underline{x}) = c$ together with Lemma 9. From Lemma 9 we know that for the case of $\pi = 0$ the only terms in equation (C.3) are the powers 2 and 4 of x with $A = \alpha_2 \sigma^2$ and $\alpha_4 = -\frac{1}{6} \frac{B}{\sigma^2}$. This implies that for $v'(\bar{x}) = 0$ we need $A = B/3\bar{x}^2$, which, in turn, implies that $v(\bar{x}) = 1/6 \frac{B}{\sigma^2} \bar{x}^4 = c$ from where we obtain the expressions for \bar{x} and A . $\hat{x} = 0$ then follows from $v(\hat{x}) = 0$. Later we use the fact the $\alpha_2 = \frac{1}{3} B \frac{\bar{x}^2}{\sigma^2}$ ■

Lemma 11 *The derivatives of the thresholds are for $\pi = 0$ are given by:*

$$\frac{\partial \bar{x}}{\partial \pi} = \frac{\partial \underline{x}}{\partial \pi} = \frac{2}{15} \frac{1}{\lambda_a(0)} \quad \text{and} \quad \frac{\partial \hat{x}}{\partial \pi} = \frac{21}{90} \frac{1}{\lambda_a(0)}$$

Corollary 2 *The derivative of the size of price changes with respect to inflation for $\pi = 0$ is*

$$\frac{\partial \Delta_p^+(0)}{\partial \pi} = -\frac{\partial \Delta_p^-(0)}{\partial \pi} = 0.1 \frac{1}{\lambda_a(0)}$$

Proof. We solve first for the derivatives of α_1 and α_2 with respect for π at $\pi = 0$.

Derivative of α_2 : From the representation of $v(x; \pi)$ in equation (C.3) we have that $\frac{\partial^2}{(\partial x)^2} v(0; \pi) = 2\alpha_2$. The symmetry of the value function implies that $v(x, \pi) = v(-x, -\pi)$, which implies $\frac{\partial^3}{(\partial x)^2 \partial \pi} v(x, \pi) = -\frac{\partial^3}{(\partial x)^2 \partial \pi} v(-x, -\pi)$ for all $x \in [\underline{x}, \bar{x}]$ and for all π . Hence, $\frac{\partial^3}{(\partial x)^2 \partial \pi} v(0, 0) = -\frac{\partial^3}{(\partial x)^2 \partial \pi} v(0, 0) = 0$. Finally, $\frac{\partial^3}{(\partial x)^2 \partial \pi} v(0, 0) = 2\frac{\partial \alpha_2}{\partial \pi} = 0$.

Derivative of α_1 : From the representation of $v(x, \pi)$ in equation (C.3), using $\frac{\partial \alpha_2}{\partial \pi} = 0$ at $\pi = 0$, we have that

$$\frac{\partial v(x, 0)}{\partial \pi} = \frac{\partial \alpha_1(0)}{\partial \pi} x + \frac{2\alpha_2}{3\sigma^2} x^3 - \frac{1}{15} \frac{B}{\sigma^4} x^5$$

Differentiating value matching, $v(\hat{x}(\pi), \pi) + c = v(\bar{x}(\pi), \pi)$ with respect to inflation and using smooth pasting we get $\frac{\partial v(\hat{x}, \pi)}{\partial \pi} = \frac{\partial v(\bar{x}, \pi)}{\partial \pi}$. Evaluating at $\pi = 0$, $\frac{\partial v(\bar{x}, 0)}{\partial \pi} = \frac{\partial v(0, 0)}{\partial \pi}$, which is equal to zero since by symmetry $\frac{\partial v(0, 0)}{\partial \pi} = 0$. Therefore, evaluating the first equation at \bar{x} and dividing by \bar{x} we get

$$0 = \frac{\partial \alpha_1(0)}{\partial \pi} + \frac{2}{3\sigma^2} \alpha_2 \bar{x}^2 - \frac{1}{15} \frac{B}{\sigma^4} \bar{x}^4$$

On the other hand, smooth pasting evaluated at $\pi = 0$, yields $\frac{\partial v(\bar{x}, 0)}{\partial x} = 2\alpha_2 \bar{x} - 4\frac{1}{6} \frac{B}{\sigma^2} \bar{x}^3 = 0$, implying $\alpha_2 \bar{x}^2 = \frac{1}{3} \frac{B}{\sigma^2} \bar{x}^4$. Substituting above implies

$$0 = \frac{\partial \alpha_1(0)}{\partial \pi} + \frac{2}{3\sigma^2} \frac{1}{3} \frac{B}{\sigma^2} \bar{x}^4 - \frac{1}{15} \frac{B}{\sigma^4} \bar{x}^4$$

and we get

$$\frac{\partial \alpha_1(0)}{\partial \pi} = -\frac{7}{45} B \frac{\bar{x}^4}{\sigma^4}.$$

Now we are ready to take the derivative of the thresholds X with respect to π at $\pi = 0$.

Derivative of \hat{x} . Using the implicit function theorem on the smooth pasting condition

$\frac{\partial v(\hat{x},0)}{\partial x} = 0$ and recalling that at $\pi = 0$ $\hat{x} = 0$ we get

$$\frac{\partial \hat{x}(0)}{\partial \pi} = -\frac{\frac{\partial^2 v(0,0)}{\partial x \partial \pi}}{\frac{\partial^2 v(0,0)}{(\partial x)^2}} = -\frac{\frac{\partial \alpha_1(0)}{\partial \pi}}{2\alpha_2(0)} = \frac{\frac{7}{45} B \frac{\bar{x}^4}{\sigma^4}}{\frac{1}{3} B \frac{\bar{x}^2}{\sigma^2}} = \frac{21}{90} \frac{\bar{x}^2}{\sigma^2}$$

Derivative of \bar{x} and of \underline{x} . Using the implicit function theorem on the smooth pasting condition $\frac{\partial v(\bar{x},0)}{\partial x} = 0$ we get

$$\frac{\partial \bar{x}(0)}{\partial \pi} = -\frac{\frac{\partial^2 v(\bar{x},0)}{\partial x \partial \pi}}{\frac{\partial^2 v(\bar{x},0)}{(\partial x)^2}}$$

Using the expression of α_2 for $\pi = 0$ in the proof of Lemma 10 and the result for the derivative of α_1 , the numerator is

$$\begin{aligned} \frac{\partial^2 v(\bar{x},0)}{\partial x \partial \pi} &= \frac{\partial \alpha_1(0)}{\partial \pi} + \frac{2\alpha_2}{3\sigma^2} \bar{x}^2 - \frac{1}{15} B \frac{\bar{x}^4}{\sigma^4} \\ &= -\frac{7}{45} B \frac{\bar{x}^4}{\sigma^4} + \frac{1}{3} B \frac{\bar{x}^4}{\sigma^4} = \frac{8}{45} B \frac{\bar{x}^4}{\sigma^4} \end{aligned}$$

Using equation (C.3) with the coefficients in Lemma 9 evaluated at $\pi = 0$ as well as the expression for α_2 for $\pi = 0$, the denominator is

$$\begin{aligned} \frac{\partial^2 v(\bar{x},0)}{(\partial x)^2} &= 2\alpha_2 - 2B \frac{\bar{x}^2}{\sigma^2} \\ &= \frac{2}{3} B \frac{\bar{x}^2}{\sigma^2} - 2B \frac{\bar{x}^2}{\sigma^2} = -\frac{4}{3} B \frac{\bar{x}^2}{\sigma^2}. \end{aligned}$$

Symmetry implies that $\frac{\partial \bar{x}(0)}{\partial \pi} = \frac{\partial \underline{x}(0)}{\partial \pi}$ Therefore,

$$\frac{\partial \bar{x}(0)}{\partial \pi} = \frac{\partial \underline{x}(0)}{\partial \pi} = \frac{2}{15} \frac{\bar{x}^2}{\sigma^2}$$

Size of price changes. Recall that the size of price increases and of price decreases is given by

$$\Delta_p^+(\pi) = \hat{x}(\pi) - \underline{x}(\pi) \text{ and } \Delta_p^-(\pi) = \bar{x}(\pi) - \hat{x}(\pi)$$

Using the previous results $\frac{\partial \Delta_p^+(0)}{\partial \pi} = \frac{\partial \hat{x}(0)}{\partial \pi} - \frac{\partial \underline{x}(0)}{\partial \pi} = \frac{21}{90} \frac{\bar{x}^2}{\sigma^2} - \frac{2}{15} \frac{\bar{x}^2}{\sigma^2} = 0.1 \frac{\bar{x}^2}{\sigma^2}$. We also know that

$\lambda_a(\pi)$ for $\pi = 0$ is $\lambda_a(0) = \frac{\sigma^2}{\bar{x}^2}$. Hence,

$$\frac{\partial \Delta_p^+(0)}{\partial \pi} = -\frac{\partial \Delta_p^-(0)}{\partial \pi} = 0.1 \frac{1}{\lambda_a(0)}$$

■

Lemma 12 *Assume that $\sigma = 0$ and $\pi > 0$ or, equivalently, that $\pi/\sigma \rightarrow \infty$ for $\sigma > 0$. Then,*

$$\hat{x}(\pi) = -\underline{x}(\pi) = \frac{1}{2} \Delta_p^+(\pi) = \left(\frac{3}{4} \frac{c}{B} \pi \right)^{1/3}$$

$$\lambda_a = \lambda_a^+ = \frac{\pi}{\Delta_p^+(\pi)} = \frac{1}{2} \left(\frac{3}{4} \frac{c}{B} \right)^{-1/3} \pi^{2/3}$$

Proof. Consider the case in which $\sigma = 0$. This is equivalent to taking the limit for $\pi \rightarrow \infty$ with $\sigma > 0$ —see Lemma 7 or Lemma 9. The state moves deterministically from \hat{x} to \underline{x} at speed π . We find closed form solutions for these thresholds. From Lemma 9 we know that for $\sigma = 0$ the value function is $v(x) = \alpha_1 \pi + \alpha_3 x^3$ with $\alpha_1 = -A/\pi$ and $\alpha_3 = \frac{B}{3\pi}$. We use smooth pasting at \hat{x} and \underline{x} , together with value matching to solve for α_1 , \hat{x} and \underline{x} . Smooth pasting, $v'(\hat{x}) = 0$ and $v'(\underline{x}) = 0$ implies that $\hat{x} = -\underline{x} = \sqrt{\frac{-\alpha_1 \pi}{B}}$. Value matching, $v(\hat{x}) - v(\underline{x}) = c$, then implies $\alpha_1 = -\left(c \frac{3}{4}\right)^{2/3} \left(\frac{B}{\pi}\right)^{1/3}$. The results of the lemma follow. ■

C.2 ANALYTICAL CHARACTERIZATION OF THE MODEL IN SECTION 3.2.2

In what follows we use $x = p - z$ for the log of the real gross mark-up. In Proposition 6 we show that when $\sigma > 0$, the inaction set is given by $\mathcal{I} = \{(p, z) : \underline{x} + z < p < \bar{x} + z\}$ and that the optimal return point is given by $\hat{\psi}(z) = \hat{x} + z$ for three constants $X \equiv (\underline{x}, \hat{x}, \bar{x})$. This is due to the combination of a assumptions of constant elasticity of demand, constant returns to scale and permanent shocks to cost while the product last. This means that it is optimal to keep the price unchanged when the real markup x is in the interval (\underline{x}, \bar{x}) . When prices are not changed, the real markup evolves according to $dx = -(\mu_z + \pi)dt + \sigma dW$. When the real markup hits either of the two thresholds, prices are adjusted so that the real markup is \hat{x} . Proposition 6 derives a system of three equations in three unknowns for X , as well as the explicit solution to the value function, as function of the parameters $\Theta \equiv (\pi, \mu_z, \sigma^2, \rho, r, \eta, c)$.

Proposition 7 derives an explicit solution for the expected number of adjustment per unit of time λ_a given a policy X and parameters $(\pi, \mu_z, \sigma^2, \rho)$. Proposition 8 characterizes the density g for invariant distribution of (p, z) implied by the policy X and the parameters $(\pi, \mu_z, \sigma^2, \rho, \eta)$.

For future reference we define \hat{c} implicitly $\zeta(z) = \hat{c} e^{z(1-\eta)}$

Proposition 6 *Assume that $\sigma > 0, c > 0$ and that equation (3.12) holds. The inaction set is given by $\mathcal{I} = \{(p, z) : \underline{x} + z < p < \bar{x} + z\}$. The optimal return point is given by $\hat{\psi}(z) = \hat{x} + z$. The value function in the range of inaction and the constants $X \equiv (\underline{x}, \hat{x}, \bar{x})$ with $\underline{x} < \hat{x} < \bar{x}$ solve*

$$V(p, z) = e^{z(1-\eta)} V(p - z, 0) \equiv e^{z(1-\eta)} v(p - z) \quad (\text{C.4})$$

$$v(x) = a_1 e^{x(1-\eta)} + a_2 e^{-x\eta} + \sum_{i=1}^2 A_i e^{\nu_i x} \quad (\text{C.5})$$

where the coefficients a_i, ν_i are given by

$$\begin{aligned} 0 &= -b_0 + b_1 \nu_i + b_2 (\nu_i)^2 \\ a_1 &= \frac{1}{b_0 - (1-\eta) b_1 - (1-\eta)^2 b_2} \text{ and } a_2 = -\frac{1}{b_0 + \eta b_1 - (\eta)^2 b_2} \text{ where} \\ b_0 &= r + \rho - \mu_z(1-\eta) - (1-\eta)^2 \frac{\sigma^2}{2}, \quad b_1 = -\left[\mu_z + \pi + 2(1-\eta) \frac{\sigma^2}{2} \right], \quad b_2 = \frac{\sigma^2}{2}. \end{aligned}$$

and where the five values A_1, A_2, X solve the following five equations:

$$\begin{aligned} \hat{c} & -a_1 (e^{\hat{x}(1-\eta)} - e^{\bar{x}(1-\eta)}) - a_2 (e^{-\hat{x}\eta} - e^{-\bar{x}\eta}) = \sum_{i=1}^2 A_i (e^{\nu_i \hat{x}} - e^{\nu_i \bar{x}}), \\ \hat{c} & -a_1 (e^{\hat{x}(1-\eta)} - e^{\underline{x}(1-\eta)}) - a_2 (e^{-\hat{x}\eta} - e^{-\underline{x}\eta}) = \sum_{i=1}^2 A_i (e^{\nu_i \hat{x}} - e^{\nu_i \underline{x}}), \\ 0 &= a_1 (1-\eta) e^{\hat{x}(1-\eta)} - a_2 \eta e^{-\hat{x}\eta} + \sum_{i=1}^2 A_i \nu_i e^{\nu_i \hat{x}}, \\ 0 &= a_1 (1-\eta) e^{\bar{x}(1-\eta)} - a_2 \eta e^{-\bar{x}\eta} + \sum_{i=1}^2 A_i \nu_i e^{\nu_i \bar{x}}, \\ 0 &= a_1 (1-\eta) e^{\underline{x}(1-\eta)} - a_2 \eta e^{-\underline{x}\eta} + \sum_{i=1}^2 A_i \nu_i e^{\nu_i \underline{x}}. \end{aligned}$$

The first two equations are linear in (A_1, A_2) , given X .

Proof. To simplify the notation we evaluate all the expressions when $\bar{p} = 0$, so the relative price and the nominal price coincide.

The Bellman equation in the inaction region $(p, z) \in \mathcal{I}$ is

$$(r + \rho)V(p, z) = e^{-\eta p} (e^p - e^z) - \pi V_p(p, z) + V_z(p, z)\mu_z + V_{zz}(p, z)\frac{\sigma^2}{2}$$

for all $p \in [\underline{p}(z), \bar{p}(z)]$. The boundary conditions are given by first order conditions for the optimal return point:

$$V_p(\hat{\psi}(z), z) = 0 \tag{C.6}$$

the value matching conditions, stating that the value at each of the two boundaries is the same as the value at the optimal price after paying the cost:

$$V(\underline{p}(z), z) = V(\hat{\psi}(z), z) - \zeta(z), \quad V(\bar{p}(z), z) = V(\hat{\psi}(z), z) - \zeta(z). \tag{C.7}$$

and the smooth pasting conditions, stating that the value function should have the same slope at the boundary than the value function in the control region (which is flat), so:

$$V_p(\underline{p}(z), z) = 0, \quad V_p(\bar{p}(z), z) = 0, \tag{C.8}$$

Under this conditions, the value function and optimal policies are homogeneous in the sense that:

$$V(p, z) = e^{z(1-\eta)}V(p - z, 0) \equiv e^{z(1-\eta)}v(p - z) \tag{C.9}$$

$$\underline{p}(z) = \underline{z} + z, \quad \bar{p}(z) = \bar{x} + z, \quad \text{and} \quad \hat{\psi}(z) = \hat{x} + z, \tag{C.10}$$

where \underline{z}, \bar{x} and \hat{x} are three constant to be determined.

Using the homogeneity of the value function in equation (C.9) we can compute the derivatives

$$V_p(p, z) = e^{z(1-\eta)}v'(p - z)$$

$$V_z(p, z) = (1 - \eta)e^{z(1-\eta)}v(p - z) - e^{z(1-\eta)}v'(p - z)$$

$$V_{zz}(p, z) = (1 - \eta)^2 e^{z(1-\eta)}v(p - z) - 2(1 - \eta)e^{z(1-\eta)}v'(p - z) + e^{z(1-\eta)}v''(p - z)$$

Replacing this derivatives in the Bellman equation for the inaction region we get

$$(r + \rho)v(p - z) = e^{-(p-z)\eta} (e^{p-z} - 1) - \pi v'(p - z) + [(1 - \eta)v(p - z) - v'(p - z)]\mu_z \\ + [(1 - \eta)^2 v(p - z) - 2(1 - \eta)v'(p - z) + v''(p - z)]\frac{\sigma^2}{2}$$

or

$$\left[r + \rho - \mu_z(1 - \eta) - (1 - \eta)^2 \frac{\sigma^2}{2} \right] v(p - z) = e^{(p-z)(1-\eta)} - e^{-\eta(p-z)} \\ - v'(p - z)[\mu_z + \pi + 2(1 - \eta)\frac{\sigma^2}{2}] + v''(p - z)\frac{\sigma^2}{2}$$

We write $x = p - z$ be the log of the gross markup, or the net markup. Consider the free boundary ODE:

$$b_0 v(x) = e^{x(1-\eta)} - e^{-\eta x} + b_1 v'(x) + b_2 v''(x) \text{ for all } x \in [\underline{x}, \bar{x}] \\ v(\underline{x}) = v(\hat{x}) - \hat{c}, \quad v(\bar{x}) = v(\hat{x}) - \hat{c}, \\ v'(\underline{x}) = 0, \quad v'(\bar{x}) = 0, \quad v'(\hat{x}) = 0,$$

where

$$b_0 = \left[r + \rho - \mu_z(1 - \eta) - (1 - \eta)^2 \frac{\sigma^2}{2} \right], \\ b_1 = - \left[\mu_z + \pi + 2(1 - \eta)\frac{\sigma^2}{2} \right], \\ b_2 = \frac{\sigma^2}{2}.$$

The value function is given by the sum of the particular solution and the solution of the homogeneous equation:

$$v(x) = a_1 e^{x(1-\eta)} + a_2 e^{-\eta x} + \sum_{i=1}^2 A_i e^{\nu_i x}$$

where ν_i are the roots of the quadratic equation

$$0 = -b_0 + b_1\nu_i + b_2(\nu_i)^2$$

and where the coefficients for the particular solution are

$$a_1 = \frac{1}{b_0 - (1 - \eta) b_1 - (1 - \eta)^2 b_2}$$

$$a_2 = -\frac{1}{b_0 + \eta b_1 - (\eta)^2 b_2},$$

since

$$b_0 a_1 e^{x(1-\eta)} = e^{x(1-\eta)} + a_1 (1 - \eta) e^{x(1-\eta)} b_1 + a_1 (1 - \eta)^2 e^{x(1-\eta)} b_2$$

$$b_0 a_2 e^{-x\eta} = -e^{-x\eta} - a_2 \eta e^{-x\eta} b_1 + a_2 (\eta)^2 e^{-x\eta} b_2 .$$

The five constants A_1, A_2 and $X \equiv (\underline{x}, \bar{x}, \hat{x})$ are chosen to satisfies the 2 value matching conditions equation (C.7), the two smooth pasting conditions equation (C.8) and the optimal return point equation (C.6). It is actually more convenient to solve the value function in two steps. First to solve for the constants $A_i(X)$ for $i = 1, 2$ using the two value matching conditions. Mathematically, the advantage of this intermediate step is that, given X , the equations for the A_1, A_2 are linear. Conceptually, the advantage is that the solution represent the value of the policy described by the triplet $X = (\bar{x}, \underline{x}, \hat{x})$. Then we solve for $(\bar{x}, \underline{x}, \hat{x})$ using the conditions for the optimality of the thresholds, namely the two smooth pasting equation (C.8) and the f.o.c. for the return point equation (C.6).

Solving A_1, A_2 for a given policy X amount to solve the following linear system:

$$\hat{c} - a_1 (e^{\hat{x}(1-\eta)} - e^{\bar{x}(1-\eta)}) - a_2 (e^{-\hat{x}\eta} - e^{-\bar{x}\eta}) = \sum_{i=1}^2 A_i (e^{\nu_i \hat{x}} - e^{\nu_i \bar{x}})$$

$$\hat{c} - a_1 (e^{\hat{x}(1-\eta)} - e^{\underline{x}(1-\eta)}) - a_2 (e^{-\hat{x}\eta} - e^{-\underline{x}\eta}) = \sum_{i=1}^2 A_i (e^{\nu_i \hat{x}} - e^{\nu_i \underline{x}})$$

Given $A_1(X), A_2(X)$ we need to solve the following three equations:

$$\begin{aligned} 0 &= a_1 (1 - \eta) e^{\hat{x}(1-\eta)} - a_2 \eta e^{-\hat{x}\eta} + \sum_{i=1}^2 A_i(X) \nu_i e^{\nu_i \hat{x}}, \\ 0 &= a_1 (1 - \eta) e^{\bar{x}(1-\eta)} - a_2 \eta e^{-\bar{x}\eta} + \sum_{i=1}^2 A_i(X) \nu_i e^{\nu_i \bar{x}}, \\ 0 &= a_1 (1 - \eta) e^{x(1-\eta)} - a_2 \eta e^{-x\eta} + \sum_{i=1}^2 A_i(X) \nu_i e^{\nu_i x}. \end{aligned}$$

■

The expected number of adjustments per unit of time is given in the next proposition:

Proposition 7 *Given a policy described by $X = (\underline{x}, \hat{x}, \bar{x})$ the expected number of adjustment per unit of time λ_a and the expected number of price increases λ_a^+ are given by*

$$\begin{aligned} \lambda_a &= 1 / \left[\frac{1}{\rho} + \sum_{i=1}^2 B_i e^{q_i \hat{x}} \right] \\ \lambda_a^+ &= 1 / \left[\frac{1}{\rho} + \sum_{i=1}^2 B_{l,i} e^{q_i \hat{x}} \right] \end{aligned}$$

where q_i are the roots of $\rho = -(\pi + \mu_z)q_i + \frac{\sigma^2}{2}(q_i)^2$ and where B_i and $B_{l,i}, B_{H,i}$ solve the following system of linear equations:

$$\begin{aligned} 0 &= \frac{1}{\rho} + \sum_{i=1}^2 B_i e^{q_i x} = \frac{1}{\rho} + \sum_{i=1}^2 B_i e^{q_i \bar{x}}, \\ \frac{1}{\rho} &= -B_{h,1} e^{q_1 \hat{x}} - B_{h,2} e^{q_2 \hat{x}}, \quad 0 = B_{h,1} (e^{q_1 \hat{x}} - e^{q_1 \bar{x}}) + B_{h,2} (e^{q_2 \bar{x}} - e^{q_2 \hat{x}}), \\ -\frac{1}{\rho} &= B_{l,1} e^{q_1 x} + B_{l,2} e^{q_2 x}, \quad 0 = B_{l,1} q_1 e^{q_1 \hat{x}} + B_{l,2} q_2 e^{q_2 \hat{x}} - B_{h,1} q_1 e^{q_1 \hat{x}} - B_{h,2} q_2 e^{q_2 \hat{x}}. \end{aligned}$$

Proof. The expected time until the next adjustment solves the following Kolmogorov equation:

$$\rho \mathcal{T}(p, z) = 1 - \pi \mathcal{T}_p(p, z) + \mathcal{T}_z(p, z) \mu_z + \mathcal{T}_{zz}(p, z) \frac{\sigma^2}{2}$$

for all p such that $\underline{p}(z) < p < \bar{p}(z)$, and all z . The boundary conditions are that time reaches

zero when it hits the barriers:

$$\mathcal{T}(\bar{p}(z), z) = \mathcal{T}(\underline{p}(z), z) = 0 .$$

Given the homogeneity of the policies we look for a function satisfying

$$\mathcal{T}(p, z) = T(p - z)$$

Given the form of the expected time we have:

$$\mathcal{T}_p(p, z) = T'(p - z), \quad \mathcal{T}_z(p, z) = -T'(p - z) \text{ and } \mathcal{T}_{zz}(p, z) = T''(p - z) ,$$

so the Kolmogorov equation becomes:

$$\rho T(x) = 1 - (\pi + \mu_z)T'(x) + T''(x) \frac{\sigma^2}{2} \text{ for all } x \in (\underline{x}, \bar{x}) .$$

The solution of this equation, given \underline{x}, \bar{x} is:

$$T(x) = \frac{1}{\rho} + \sum_{i=1}^2 B_i e^{q_i x} \text{ for all } x \in (\underline{x}, \bar{x})$$

where q_i are roots of

$$\rho = -(\pi + \mu_z)q_i + \frac{\sigma^2}{2}(q_i)^2 , \tag{C.11}$$

and where the B_1, B_2 are chosen so that the expected time is zero at the boundaries:

$$0 = \frac{1}{\rho} + \sum_{i=1}^2 B_i e^{q_i \underline{x}} \tag{C.12}$$

$$0 = \frac{1}{\rho} + \sum_{i=1}^2 B_i e^{q_i \bar{x}} \tag{C.13}$$

Given the solution of this two linear equations $B_1(\underline{x}, \bar{x}), B_2(\underline{x}, \bar{x})$ the expected number of adjustments per unit of time λ_a is given by

$$\lambda_a = \frac{1}{T(\hat{x})} = \frac{1}{\frac{1}{\rho} + \sum_{i=1}^2 B_i(\underline{x}, \bar{x}) e^{q_i \hat{x}}}$$

Finally, we derive the expression for the frequency of price increases. The time until the next price increase is the first time until x hits \underline{x} or the product dies while $\underline{x} < x < \hat{x}$. If x hits \bar{x} , or the product dies exogenously while $\bar{x} > x > \hat{x}$, then x then is returned to \hat{x} . Thus the expected time until the next increase in price solves the following Kolmogorov equation:

$$\rho \mathcal{T}(p, z) = \begin{cases} 1 - \pi \mathcal{T}_p(p, z) + \mathcal{T}_z(p, z) \mu_z + \mathcal{T}_{zz}(p, z) \frac{\sigma^2}{2} & \text{if } p < z + \hat{x} \\ 1 + \rho \mathcal{T}(z + \hat{x}, z) - \pi \mathcal{T}_p(p, z) + \mathcal{T}_z(p, z) \mu_z + \mathcal{T}_{zz}(p, z) \frac{\sigma^2}{2} & \text{if } p > z + \hat{x} \end{cases}$$

for all p such that $\underline{p}(z) < p < \bar{p}(z)$, and all z . The boundary conditions are that time reaches zero when it hits the barriers:

$$\mathcal{T}(\bar{x} + z, z) = \mathcal{T}(\hat{x} + z, z) \text{ and } \mathcal{T}(\underline{x} + z, z) = 0 .$$

We look for a solution that is continuous and once differentiable at $(p, z) = (\hat{x} + z, z)$, and otherwise twice continuously differentiable. To do so we let $\mathcal{T}(p, z) = T_h(x)$ for $x \in [\hat{x}, \bar{x}]$ and $\mathcal{T}(p, z) = T_l(x)$ for $x \in [\underline{x}, \hat{x}]$ and

$$\begin{aligned} \rho T_l(x) &= 1 - (\pi + \mu_z) T_l'(x) + T_l''(x) \frac{\sigma^2}{2} \\ \rho T_h(x) &= 1 + \rho T_h(\hat{x}) - (\mu_z + \pi) T_h'(x) + T_h''(x) \frac{\sigma^2}{2} \\ T_l(\hat{x}) &= T_h(\hat{x}) , T_l'(\hat{x}) = T_h'(\hat{x}) \\ T_h(\hat{x}) &= T_h(\bar{x}) , T_l(\underline{x}) = 0 . \end{aligned}$$

The solution for T_j for $j = h, l$ are:

$$T_l(x) = \frac{1}{\rho} + \sum_{i=1}^2 B_{l,i} e^{q_i x} \text{ and } T_h(x) = \frac{1}{\rho} + \sum_{i=1}^2 B_{h,i} e^{q_i x} + \left(\frac{1}{\rho} + \sum_{i=1}^2 B_{l,i} e^{q_i \hat{x}} \right)$$

The four boundary conditions become the following four linear equations of the constants

B' 's:

$$\begin{aligned}
\sum_{i=1}^2 B_{l,i} e^{q_i \hat{x}} &= \sum_{i=1}^2 B_{h,i} e^{q_i \hat{x}} + \sum_{i=1}^2 B_{l,i} e^{q_i \hat{x}} + \frac{1}{\rho} \\
\sum_{i=1}^2 B_{l,i} q_i e^{q_i \hat{x}} &= \sum_{i=1}^2 B_{h,i} q_i e^{q_i \hat{x}} \\
\sum_{i=1}^2 B_{l,i} e^{q_i \bar{x}} &= 0 \\
\sum_{i=1}^2 B_{h,i} e^{q_i \hat{x}} &= \sum_{i=1}^2 B_{h,i} e^{q_i \bar{x}}
\end{aligned}$$

Hence, the frequency of price increases λ_a^+ is given by

$$\lambda_a^+ = 1/T_l(\hat{x}) = 1 / \left(1/\rho + \sum_{i=1}^2 B_{l,i} e^{q_i \hat{x}} \right) . \blacksquare$$

Now we turn to the density of the invariant distribution

Proposition 8 *Given a policy described by $X = (\underline{x}, \hat{x}, \bar{x})$ the density of the invariant distribution $g(p, z)$ is given by*

$$g(p, z) = \begin{cases} e^{\phi_1 z} [U_1^+ e^{\xi_1(p-z)} + U_2^+] & \text{if } p - z \in (\hat{x}, \bar{x}] , z > 0 \\ e^{\phi_1 z} [L_1^+ e^{\xi_1(p-z)} + L_2^+] & \text{if } p - z \in [\underline{x}, \hat{x}] , z > 0 \\ e^{\phi_2 z} [U_1^- e^{\xi_2(p-z)} + L_2^-] & \text{if } p - z \in (\hat{x}, \bar{x}] , z < 0 \\ e^{\phi_2 z} [L_1^- e^{\xi_2(p-z)} + L_2^-] & \text{if } p - z \in [\underline{x}, \hat{x}] , z < 0 \\ 0 & \text{otherwise} \end{cases} \quad (\text{C.14})$$

where $\{\phi_1, \phi_2, \xi_1, \xi_2\}$ are given by

$$\begin{aligned}
\rho &= -\mu_z \phi_j + \frac{\sigma^2}{2} \phi_j^2 \text{ for each of the roots } j = 1, 2 \text{ and} \\
\xi_j &= -\frac{\pi + \mu_z - 2\phi_j \frac{\sigma^2}{2}}{\sigma^2/2}
\end{aligned}$$

and where the coefficients $\{U_i^+, L_i^+, U_i^-, L_i^-\}_{i=1,2}$ solve 8 linear equations:

$$\begin{aligned}
0 &= U_1^+ e^{\xi_1 \bar{x}} + U_2^+ = L_1^+ e^{\xi_1 \underline{x}} + L_2^+ \\
0 &= U_1^- e^{\xi_2 \bar{x}} + U_2^- = L_1^- e^{\xi_2 \underline{x}} + L_2^- \\
\frac{\phi_1 \phi_2}{\phi_1 - \phi_2} &= \frac{L_1^+}{\xi_1} [e^{\xi_1 \hat{x}} - e^{\xi_1 \underline{x}}] + L_2^+ [\hat{x} - \underline{x}] + \frac{U_1^+}{\xi_1} [e^{\xi_1 \bar{x}} - e^{\xi_1 \hat{x}}] + U_2^+ [\bar{x} - \hat{x}] \\
\frac{\phi_1 \phi_2}{\phi_1 - \phi_2} &= \frac{L_1^-}{\xi_2} [e^{\xi_2 \hat{x}} - e^{\xi_2 \underline{x}}] + L_2^- [\hat{x} - \underline{x}] + \frac{U_1^-}{\xi_2} [e^{\xi_2 \bar{x}} - e^{\xi_2 \hat{x}}] + U_2^- [\bar{x} - \hat{x}] \\
U_1^+ e^{\xi_1 \hat{x}} + U_2^+ &= L_1^+ e^{\xi_1 \hat{x}} + L_2^+ \\
U_1^- e^{\xi_2 \hat{x}} + U_2^- &= L_1^- e^{\xi_2 \hat{x}} + L_2^- .
\end{aligned}$$

Proof. The density of the invariant distribution for (p, z) solves the forward Kolmogorov p.d.e.:

$$\rho g(p, z) = \pi g_p(p, z) - \mu_z g_z(p, z) + g_{zz}(p, z) \frac{\sigma^2}{2} \quad (\text{C.15})$$

for all $(p, z) \neq (\hat{x} + z, z) = (\hat{\psi}(z), z)$ and all $p : \underline{p}(z) = \underline{x} + z \leq p \leq \bar{x} + z = \bar{p}(z)$ and all z . The pde does not apply at the optimal return point, since local consideration cannot determine g there. The other boundary conditions are zero density at the lower and upper boundaries of adjustments, and that g integrates to one:

$$\begin{aligned}
g(\underline{x} + z, z) &= g(\bar{x} + z, z) = 0 \text{ for all } z \\
1 &= \int_{\infty}^{\infty} \int_{\underline{x}+z}^{\bar{x}+z} g(p, z) dp dz
\end{aligned}$$

We will show that g can be computed dividing the state space in four regions given by whether (p, z) is such that $z > 0$ and $z < 0$ and given z whether $p \in [\underline{x} + z, \hat{x} + z]$ and $p \in [\hat{x} + z, \bar{x} + z]$.

As a preliminary step we solve for the marginal on z of the invariant distribution, which we denote by \tilde{g} . This is the invariant distribution of the process $\{z\}$ which with intensity ρ is re-started at zero and otherwise follows $dz = \mu_z dt + \sigma dW$. It can be shown that \tilde{g} is given by

$$\tilde{g}(z) \equiv \int_{\underline{x}+z}^{\bar{x}+z} g(p, z) dp = \begin{cases} \frac{\phi_1 \phi_2}{\phi_1 - \phi_2} e^{\phi_2 z} & \text{if } z < 0 \\ \frac{\phi_1 \phi_2}{\phi_1 - \phi_2} e^{\phi_1 z} & \text{if } z > 0, \end{cases} \quad (\text{C.16})$$

where $\phi_1 < 0 < \phi_2$ are the two real roots of the characteristic equation

$$\rho = -\mu_z \phi + \frac{\sigma^2}{2} \phi^2. \quad (\text{C.17})$$

We conjecture that g can be written as follows:

$$g(p, z) = \begin{cases} e^{\phi_1 z} k(p-z) & \text{if } z < 0 \\ e^{\phi_2 z} k(p-z) & \text{if } z > 0, \end{cases} \quad (\text{C.18})$$

In this case we compute the derivatives as:

$$\begin{aligned} g_p(p, z) &= e^{\phi z} k'(p-z), \\ g_z(p, z) &= e^{\phi z} \phi k(p-z) - e^{\phi z} k'(p-z), \\ g_{zz}(p, z) &= e^{\phi z} \phi^2 k(p-z) - 2e^{\phi z} \phi k'(p-z) + e^{\phi z} k''(p-z). \end{aligned}$$

where $\phi = \phi_1$ for $z < 0$ and $\phi = \phi_2$ for $z > 0$. The p.d.e. then becomes:

$$\rho k(p-z) = \pi k'(p-z) - \mu_z [\phi k(p-z) - k'(p-z)] + [\phi^2 k(p-z) - 2\phi k'(p-z) + k''(p-z)] \frac{\sigma^2}{2}$$

for $p-z \neq \hat{x}$ or

$$\left[\rho + \phi \mu_z - \phi^2 \frac{\sigma^2}{2} \right] k(p-z) = \left(\pi + \mu_z - 2\phi \frac{\sigma^2}{2} \right) k'(p-z) + k''(p-z) \frac{\sigma^2}{2}$$

Thus

$$k(p-z) = \begin{cases} U_1 e^{\xi_{1j}(p-z)} + U_2 e^{\xi_{2j}(p-z)} & \text{if } p-z \in (\hat{x}, \bar{x}] \\ L_1 e^{\xi_{1j}(p-z)} + L_2 e^{\xi_{2j}(p-z)} & \text{if } p-z \in [\underline{x}, \hat{x}] \end{cases}$$

where ξ_{1j}, ξ_{2j} solves the quadratic equation:

$$\left[\rho + \phi_j \mu_z - \phi_j^2 \frac{\sigma^2}{2} \right] = \left(\pi + \mu_z - 2\phi_j \frac{\sigma^2}{2} \right) \xi + \frac{\sigma^2}{2} \xi^2 \quad (\text{C.19})$$

for each $j = 1, 2$ corresponding to $\phi = \phi_1$ and $\phi = \phi_2$, i.e., the positive and negative values of z . We note that, by definition of ϕ in equation (C.17), the left hand side of equation (C.19) equal zero, and hence one of the two roots is always equal to zero. Thus we label $\xi_{2j} = 0$ for

$j = 1, 2$. The remaining root equals:

$$\xi_{1j} = -\frac{\pi + \mu_z - 2\phi_j \frac{\sigma^2}{2}}{\sigma^2/2} \quad \text{and} \quad \xi_{2j} = 0 \quad \text{for } j = 1, 2. \quad (\text{C.20})$$

We integrate $g(p, z)$ over p and equate it to $\tilde{g}(z)$ to obtain a condition for coefficients C . First we consider the case of $z > 0$:

$$\begin{aligned} \tilde{g}(z) &= \int_{\underline{x}+z}^{\hat{x}+z} e^{\phi_1 z} \sum_{i=1}^2 L_i^+ e^{\xi_{i1}(p-z)} dp + \int_{\hat{x}+z}^{\bar{x}+z} e^{\phi_1 z} \sum_{i=1}^2 U_i^+ e^{\xi_{i1}(p-z)} dp \\ &= e^{\phi_1 z} \sum_{i=1}^2 \frac{L_i^+ e^{-\xi_{i1} z}}{\xi_{i1}} [e^{\xi_{i1}(\bar{x}+z)} - e^{\xi_{i1}(\hat{x}+z)}] + e^{\phi_1 z} \sum_{i=1}^2 \frac{U_i^+ e^{-\xi_{i1} z}}{\xi_{i1}} [e^{\xi_{i1}(\hat{x}+z)} - e^{\xi_{i1}(\underline{x}+z)}] \\ &= e^{\phi_1 z} \left(\sum_{i=1}^2 \frac{L_i^+}{\xi_{i1}} [e^{\xi_{i1}\bar{x}} - e^{\xi_{i1}\hat{x}}] + \sum_{i=1}^2 \frac{U_i^+}{\xi_{i1}} [e^{\xi_{i1}\hat{x}} - e^{\xi_{i1}\underline{x}}] \right) \\ &= e^{\phi_1 z} \left(\frac{L_1^+}{\xi_{11}} [e^{\xi_{11}\hat{x}} - e^{\xi_{11}\underline{x}}] + L_2^+ [\hat{x} - \underline{x}] + \frac{U_1^+}{\xi_{11}} [e^{\xi_{11}\bar{x}} - e^{\xi_{11}\hat{x}}] + U_2^+ [\bar{x} - \hat{x}] \right) \end{aligned}$$

where the last line uses that $\xi_{2,1} = 0$. The analogous expression for $z < 0$ is

$$\tilde{g}(z) = e^{\phi_2 z} \left(\frac{L_1^-}{\xi_{12}} [e^{\xi_{12}\hat{x}} - e^{\xi_{12}\underline{x}}] + L_2^- [\hat{x} - \underline{x}] + \frac{U_1^-}{\xi_{12}} [e^{\xi_{12}\bar{x}} - e^{\xi_{12}\hat{x}}] + U_2^- [\bar{x} - \hat{x}] \right)$$

The value of the density at the boundary of the range of inaction is given by

$$g(\bar{x} + z, z) = \begin{cases} e^{\phi_1 z} \sum_{i=1}^2 U_i^+ e^{\xi_{i1}\bar{x}} & \text{for } z > 0 \\ e^{\phi_2 z} \sum_{i=1}^2 U_i^+ e^{\xi_{i2}\bar{x}} & \text{for } z < 0 \end{cases}$$

$$g(\underline{x} + z, z) = \begin{cases} e^{\phi_1 z} \sum_{i=1}^2 L_i^+ e^{\xi_{i1}\underline{x}} & \text{for } z > 0 \\ e^{\phi_2 z} \sum_{i=1}^2 L_i^+ e^{\xi_{i2}\underline{x}} & \text{for } z < 0 \end{cases}$$

If the density g at $(p, z) = (\hat{\psi}(z), z) = (z + \hat{x}, z)$ is continuous on p for a given z we have:

$$g(\hat{x}, z) = \begin{cases} e^{\phi_1 z} [\sum_{i=1}^2 U_i^+ e^{\xi_{i1}\hat{x}}] = e^{\phi_1 z} [\sum_{i=1}^2 L_i^+ e^{\xi_{i1}\hat{x}}] & \text{if } z > 0, \\ e^{\phi_2 z} [\sum_{i=1}^2 U_i^- e^{\xi_{i2}\hat{x}}] = e^{\phi_2 z} [\sum_{i=1}^2 L_i^- e^{\xi_{i2}\hat{x}}] & \text{if } z < 0. \end{cases} \quad (\text{C.21})$$

We summarize the results for the invariant density g here

$$g(p, z) = \begin{cases} e^{\phi_1 z} [U_1^+ e^{\xi_1(p-z)} + U_2^+] & \text{if } p - z \in (\hat{x}, \bar{x}), z > 0 \\ e^{\phi_1 z} [L_1^+ e^{\xi_1(p-z)} + L_2^+] & \text{if } p - z \in [\underline{x}, \hat{x}), z > 0 \\ e^{\phi_2 z} [U_1^- e^{\xi_2(p-z)} + L_2^-] & \text{if } p - z \in (\hat{x}, \bar{x}), z < 0 \\ e^{\phi_2 z} [L_1^- e^{\xi_2(p-z)} + L_2^-] & \text{if } p - z \in [\underline{x}, \hat{x}), z < 0 \end{cases} \quad (\text{C.22})$$

where $\{\phi_1, \phi_2\}$ are the two roots of the quadratic equation (C.17) and where the use $\xi_1 \equiv \xi_{11}, \xi_2 \equiv \xi_{12}$ are given by the non-zero roots equation (C.20).

The 8 values for $\{U_i^+, L_i^+, U_i^-, L_i^-\}_{i=1,2}$ solve two system of 4 linear equations, one for $\{U_i^+, L_i^+\}_{i=1,2}$ and one for $\{U_i^-, L_i^-\}_{i=1,2}$. The upper and lower boundary of the range of inaction has zero density for both positive and negative values of z :

$$0 = U_1^+ e^{\xi_1 \bar{x}} + U_2^+ = L_1^+ e^{\xi_1 \underline{x}} + L_2^+ \quad (\text{C.23})$$

$$0 = U_1^- e^{\xi_2 \bar{x}} + U_2^- = L_1^- e^{\xi_2 \underline{x}} + L_2^- \quad (\text{C.24})$$

The marginal distribution of the z computed using g coincides with \tilde{g} for positive and negative values of z :

$$\frac{\phi_1 \phi_2}{\phi_1 - \phi_2} = \frac{L_1^+}{\xi_1} [e^{\xi_1 \hat{x}} - e^{\xi_1 \underline{x}}] + L_2^+ [\hat{x} - \underline{x}] + \frac{U_1^+}{\xi_1} [e^{\xi_1 \bar{x}} - e^{\xi_1 \hat{x}}] + U_2^+ [\bar{x} - \hat{x}] \quad (\text{C.25})$$

$$\frac{\phi_1 \phi_2}{\phi_1 - \phi_2} = \frac{L_1^-}{\xi_2} [e^{\xi_2 \hat{x}} - e^{\xi_2 \underline{x}}] + L_2^- [\hat{x} - \underline{x}] + \frac{U_1^-}{\xi_2} [e^{\xi_2 \bar{x}} - e^{\xi_2 \hat{x}}] + U_2^- [\bar{x} - \hat{x}] \quad (\text{C.26})$$

The density is continuous at $(p, z) = (\hat{\psi}(z), z) = (z + \hat{x}, z)$. Thus

$$U_1^+ e^{\xi_1 \hat{x}} + U_2^+ = L_1^+ e^{\xi_1 \hat{x}} + L_2^+ \quad (\text{C.27})$$

$$U_1^- e^{\xi_2 \hat{x}} + U_2^- = L_1^- e^{\xi_2 \hat{x}} + L_2^-. \quad (\text{C.28})$$

■

We are now ready to characterize the marginal density of prices, denoted by h .

$$h(p) \equiv \int_{-\infty}^{\infty} g(p, z) dz = \int_{p-\bar{x}}^{p-\underline{x}} g(p, z) dz = \int_{p-\bar{x}}^{p-\hat{x}} g(p, z) dz + \int_{p-\hat{x}}^{p-\underline{x}} g(p, z) dz \quad (\text{C.29})$$

$$h(p) = \begin{cases} \int_{p-\bar{x}}^{p-\hat{x}} e^{\phi_1 z} [U_1^+ e^{\xi_1(p-z)} + U_2^+] dz + \int_{p-\hat{x}}^{p-\underline{x}} e^{\phi_1 z} [L_1^+ e^{\xi_1(p-z)} + L_2^+] dz & \text{if } p > \bar{x} \\ \int_{p-\bar{x}}^0 e^{\phi_2 z} [U_1^- e^{\xi_2(p-z)} + U_2^-] dz + \int_0^{p-\hat{x}} e^{\phi_1 z} [U_1^+ e^{\xi_1(p-z)} + U_2^+] dz \\ + \int_{p-\hat{x}}^{p-\underline{x}} e^{\phi_1 z} [L_1^+ e^{\xi_1(p-z)} + L_2^+] dz & \text{if } p \in [\hat{x}, \bar{x}) \\ \int_{p-\bar{x}}^{p-\hat{x}} e^{\phi_2 z} [U_1^- e^{\xi_2(p-z)} + U_2^-] dz + \int_{p-\hat{x}}^0 e^{\phi_2 z} [L_1^- e^{\xi_2(p-z)} + L_2^-] dz \\ + \int_0^{p-\underline{x}} e^{\phi_1 z} [L_1^+ e^{\xi_1(p-z)} + L_2^+] dz & \text{if } p \in [\underline{x}, \hat{x}) \\ \int_{p-\bar{x}}^{p-\hat{x}} e^{\phi_2 z} [U_1^- e^{\xi_2(p-z)} + U_2^-] dz + \int_{p-\hat{x}}^{p-\underline{x}} e^{\phi_2 z} [L_1^- e^{\xi_2(p-z)} + L_2^-] dz & \text{if } p < \underline{x} \end{cases} \quad (\text{C.30})$$

Using that

$$\begin{aligned} \int_a^b e^{\phi z} [D_1 e^{\xi(p-z)} + D_2] dz &= D_1 e^{\xi p} \int_a^b e^{(\phi-\xi)z} dz + D_2 \int_a^b e^{\phi z} dz \\ &= D_1 \frac{e^{\xi p + (\phi-\xi)b} - e^{\xi p + (\phi-\xi)a}}{\phi - \xi} + D_2 \frac{e^{\phi b} - e^{\phi a}}{\phi} \end{aligned}$$

we can solve the integrals for each of the branches of h where $D_i \in \{U_i^+, U_i^-, L_i^+, L_i^-\}$ for $i = 1, 2$ and $(\phi, \xi) \in \{\phi_1, \xi_1, \phi_2, \xi_2\}$ and a and b take different values accordingly.

For completeness we give the expression for the case with $\sigma = 0$, a version of Sheshinski and Weiss (1977) model.

Proposition 9 *Assume that $\sigma = 0, c > 0, \pi + \mu_z > 0$ and equation (3.12) holds. The inaction set is given by $\mathcal{I} = \{(p, z) : \underline{x} + z < p < \bar{x} + z\}$. The optimal return point is given by $\hat{\psi}(z) = \hat{x} + z$. The value function in the range of inaction and the constants $X \equiv (\underline{x}, \hat{x})$ solve*

$$\begin{aligned} V(p, z) &= e^{z(1-\eta)} V(p - z, 0) \equiv e^{z(1-\eta)} v(p - z) \\ v(x) &= a_1 e^{x(1-\eta)} + a_2 e^{-x\eta} + A e^{\nu x} \end{aligned}$$

where the coefficients a_i , ν are given by

$$\begin{aligned}\nu &= \frac{b_0}{b_1}, \quad a_1 = \frac{1}{b_0 - (1-\eta)b_1} \quad \text{and} \quad a_2 = -\frac{1}{b_0 + \eta b_1} \quad \text{where} \\ b_0 &= r + \rho - \mu_z(1-\eta), \quad b_1 = -[\mu_z + \pi].\end{aligned}$$

and where the three values $A, X \equiv (\underline{x}, \hat{x})$ solve the following three equations:

$$\begin{aligned}\hat{c} &= -a_1 (e^{\hat{x}(1-\eta)} - e^{\underline{x}(1-\eta)}) - a_2 (e^{-\hat{x}\eta} - e^{-\underline{x}\eta}) = A (e^{\nu\hat{x}} - e^{\nu\underline{x}}), \\ 0 &= a_1 (1-\eta)e^{\hat{x}(1-\eta)} - a_2 \eta e^{-\hat{x}\eta} + A \nu e^{\nu\hat{x}}, \\ 0 &= a_1 (1-\eta)e^{\underline{x}(1-\eta)} - a_2 \eta e^{-\underline{x}\eta} + A \nu e^{\nu\underline{x}}.\end{aligned}$$

Furthermore:

$$\lambda_a = \frac{\rho}{1 - \exp\left(-\frac{\rho}{\pi + \mu_z}(\hat{x} - \underline{x})\right)}$$

Proof. The expressions for the value function are obtained by setting $\sigma = 0$ and imposing the sS policy between the bands \underline{x}, \hat{x} . This problem is identical to the one in Sheshinski and Weiss (1977), where the discount rate is $r + \rho$. For the frequency of price adjustment we need to include the death and replacement of the products. For this we let \mathcal{T} the expected time until an adjustment:

$$\begin{aligned}\rho\mathcal{T}(p, z) &= 1 + \mathcal{T}_z(p, z)\mu_z - \mathcal{T}_p(p, z)\pi \\ \rho T(x) &= 1 - T'(x)(\mu_z + \pi).\end{aligned}$$

with boundary conditions $T(\underline{x}) = 0$. So the solution is $T(x) = 1/\rho + B \exp\left(-\frac{\rho}{\pi + \mu_z}x\right)$ with $B = -\exp\left(\frac{\rho}{\pi + \mu_z}\underline{x}\right)/\rho$ so

$$T(x) = \frac{1}{\rho} \left[1 - \exp\left(-\frac{\rho}{\pi + \mu_z}(x - \underline{x})\right) \right]$$

and hence

$$1/\lambda_a = T(\hat{x}) = \frac{1}{\rho} \left[1 - \exp\left(-\frac{\rho}{\pi + \mu_z}(\hat{x} - \underline{x})\right) \right] \quad \blacksquare$$

C.2.1 SOME GRAPHICAL ILLUSTRATIONS OF THE MODEL IN SECTION 3.2.2

Figure C.1 plots the size of the “regular” price increases and “regular” price decreases, for different inflation rates. Imitating the empirical literature, we define as regular price changes those not triggered by the jump shock that reset the value of z to zero. The figure shows that for low inflation rates the size of price increases and of price decreases is very similar as predicted by part (iii) of Proposition 3. The figure also shows that the slope of the size of price changes with respect to inflation is very small, consistent with the prediction that $\frac{\partial \Delta_p^+(0, \sigma^2)}{\partial \pi} + \frac{\partial \Delta_p^-(0, \sigma^2)}{\partial \pi} = 0$. As inflation rises, the size of both price increases and decreases becomes larger. Figure 3.10 in Section 3.5 is the empirical counterpart of Figure C.1.

Figure C.2 plots the unconditional variance of relative prices as a function of inflation. For low values of inflation the unconditional variance is insensitive to inflation as predicted by the theory. For large enough inflation rates the width of the inaction range should swamp the effect of the variation of z on σ . In our numerical examples, however, it takes inflation rates even much higher than the ones observed in the peak months in Argentina for this to happen.

C.3 CLASSIFICATION OF GOODS SERVICES AND DATA COLLECTION

Goods/services in the dataset are classified according to the MERCOSUR Harmonized Index of Consumer Price (HICP) classification. The HICP uses the first four digit levels of the Classification of Individual Consumption According to Purpose (COICOP) of the United Nations plus three digit levels based on the CPI of the MERCOSUR countries. The *goods/services* in the database are the seven digit level of the HICP classification; six digit level groups are called *products*; five digit level groups are called *sub-classes*; four digit level categories are called *classes*; three digit level categories are called *groups* and two digit level groups are called *divisions*. Table C.1 shows two examples of this classification.

For most cases, the brand chosen for the product is the one most widely sold by the outlet, or the one that occupies more space in the stands, if applicable (hence brands can change from month to month or from two-weeks to two-weeks). For some cases, the brand is part of the attributes, the product is defined as one from a “top brand”.

The weight of each good in the CPI is obtained from the 1986 National Expenditure Sur-

Figure C.1: Average size of price increases Δ_p^+ and decreases Δ_p^-

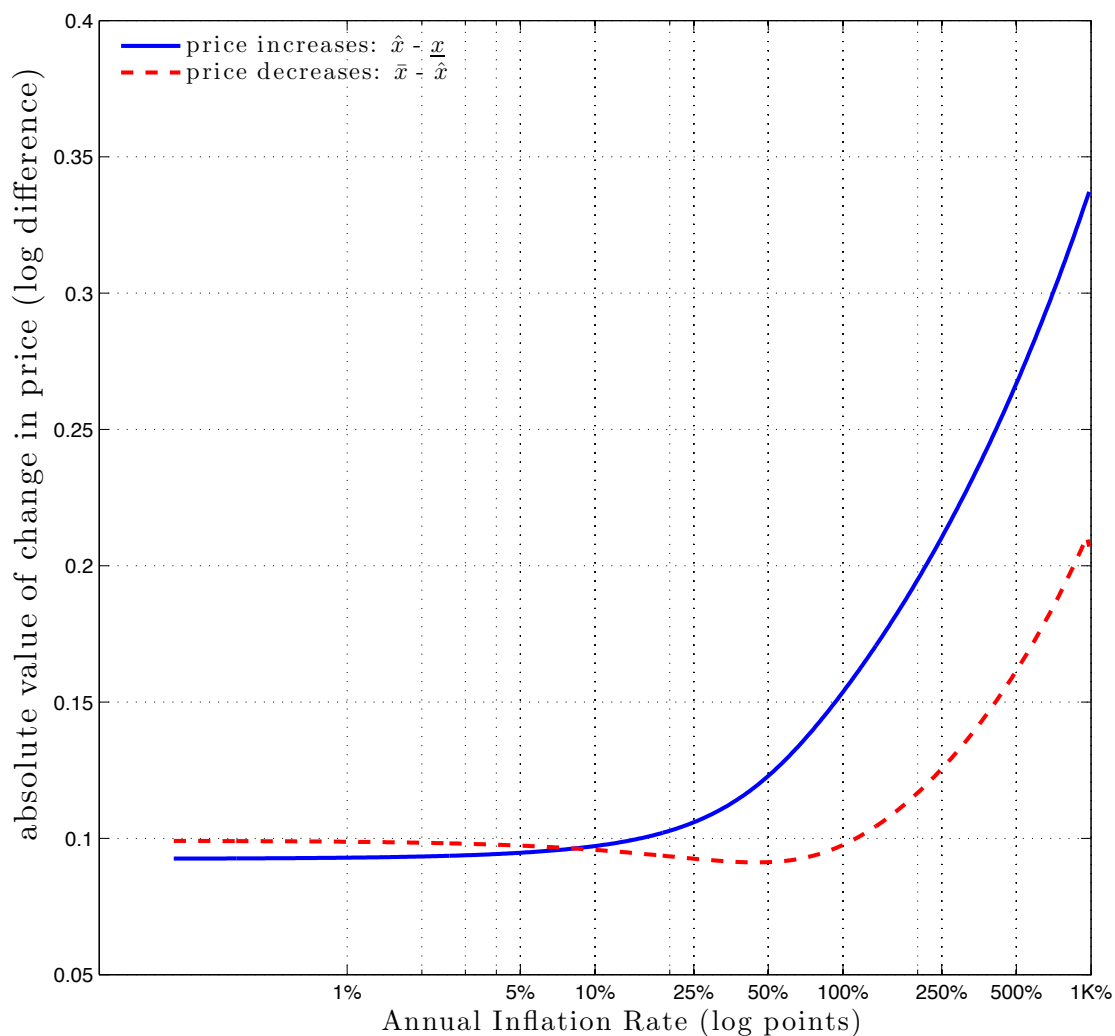



Table C.1: Example of the Harmonized Index of Consumer Price Classification

Classification	Example 1	Example 2
Division	Food and Beverages	Household equipment and maintenance
Group	Food	Household maintenance
Class	Fruits	Cleaning tools and products
Sub-Class	Fresh Fruits	Cleaning products
Product	Citric Fruits	Soaps and detergents
Good	Lemons	Liquid soap

Figure C.2: Unconditional standard deviation of log relative prices $\bar{\sigma}$



vey (Encuesta Nacional de Gasto de los Hogares). Weights are computed as the proportion of the households expenditure on each good over the total expenditure of the households. We re-normalize these weights to reflect the fact that our working dataset represents only 84% of expenditure. The weight of a particular good is proportional to the importance of its expenditure with respect to the total expenditure without taking into account the percentage of households buying it.

Table C.2 shows the top 20 goods, in terms of the importance of their weights in our sample. As it can be seen from the table, most goods whose prices are gathered twice a month are represented by food and beverages while goods whose prices are gathered monthly include services, apparel and other miscellaneous goods and services.

Table C.3 shows the weight structure in our database classified by divisions. The table shows goods in terms of their weight with respect to the total weight in the sample (Total column), and with respect to the total weight of their belonging category (one or two visit goods) Food and non-alcoholic beverages represent almost 43% of the total weight in the

Table C.2: Goods ordered by weight whose prices are gathered once and twice per month

Differentiated)	Weight (%)	Homogeneous	Weight (%)
Lunch	2.02	Whole chicken	1.51
Lunch in the workplace	1.67	Wine	1.49
Car	1.52	French bread (less than 12 pieces)	1.38
Housemaid	1.48	Fresh whole milk	1.31
Monthly union membership	1.22	Blade steaks	0.96
Snack	0.91	Standing rump	0.93
Medical consultation	0.83	Eggs	0.87
Gas bottle	0.58	Short ribs (Roast prime ribs)	0.85
Laides hairdresser	0.56	Striploin steaks	0.78
Labor for construction	0.54	Apple	0.73
Adult cloth slippers	0.50	Oil	0.72
Color TV	0.50	Rump steaks	0.72
Funeral expenses	0.49	Potatos	0.71
Men's dress shirt	0.49	Soda (coke)	0.70
Dry cleaning and ironing	0.48	Cheese (quartirollo type)	0.70
Sports club fee	0.47	Tomatoes	0.63
Movie ticket	0.47	Minced meat	0.60
Men's denim pants	0.45	Sugar (white)	0.59
Disposable diapers	0.43	Coffee (in package)	0.57
Shampoo	0.40	Yerba mate	0.55

sample and 82% of the total weight of goods whose prices are gathered twice a month. On the other hand, weights of one visit goods are less concentrated. Almost 12% of the total weight in the sample corresponds to furniture and household items and around 9% correspond to apparel. These percentages are around 24% and 19%, respectively, when computing percentages over the total weight in the one visit goods category.

Table C.3: Weights for Divisions of the Harmonized Index of Consumer Price

Divisions	Weight		Weight	
	Total	One visit	Total	Two visits
Food and non-alcoholic beverages	4.94	9.91	42.67	82.39
Apparel	9.24	18.53		
Conservation and repair of housing plus gas in bottle	0.99	1.99		
Furniture and household items	11.86	23.79		
Medical products, appliances and equipment plus external medical services	3.39	6.79		
Transportation	2.60	5.22		
Recreation and culture	4.50	9.02		
Education	2.42	4.85		
Miscellaneous goods and services (Toiletries, haircut, etc.)	5.07	10.16		
Jewelry, clocks, watches plus other personal belongings	4.86	9.75		
Alcoholic beverages			3.94	7.62
Non-durable household goods			2.40	4.64
Other items and personal care products			2.77	5.36

C.3.1 INSTRUCTIONS TO CPI'S POLLSTERS

Pollsters record item's prices. Remember that an item is a good/service of a determined brand sold in a specific outlet in a specific period of time. Prices are transactional, meaning that the pollster should be able to buy the product in the outlet. Goods are defined by their attributes. For the majority of the goods, the brand is not an attribute. The brand of a specific item is determined the first time the pollster visits an outlet. The brand is the most sold/displayed by the outlet. Once the item is completely defined, the pollster collects the price of that item next time she visits the same outlet. After the first visit, in the following visits, the pollsters arrive to each outlet with a form that includes all items for which prices are to be collected.

For example, assume the good is *soda-cola top brand* and the attributes are *package: plastic bottle and weight: 1.5 liters*. The first time the pollster goes to, say, outlet *A* she ask for the cola top brand most sold in that outlet. Assume that *Coca-Cola* is the most sold soda in outlet *A*. Then the item is completely defined: *Coca-Cola* in plastic bottle of 1.5 liters in outlet *A*. Next time the pollster goes to outlet *A* she records the price of that item.

All prices are in Argentine pesos. Our dataset does not contain flags for indexed prices. In traditional outlets pollsters ask for the price of an item even when, for example it is displayed in a blackboard, because the good has to be available in order to record its price. In supermarket chains the price recorded is collected from the shelf/counter display area.

There are a number of special situations to be taken into account:

1. *Substitutions*: every time there is a change in the attributes of the good the pollster has to replace that particular good for another one. In this case, the pollster mark the price collected with a flag indicating a substitution has occurred. The goods that are substituted should be similar in terms of the type of brand and or quality.
2. *Stockouts*: every time the pollster could not buy the item, either because the good is out of stock in the outlet or because the same good or a similar one is not sold by outlet at the time the price has to be collected, she has to mark the item with a flag of stockout and she has to assign a missing price for that item. Stock-outs include what we label "pure stock-outs", the case where the outlet has depleted the stock of the good, including end of seasonal goods/services. Stock outs also include the case where the outlet no longer carries the same good/service and it does not offers a similar

good/service of comparable quality/brand. Examples of stock-outs include many fruits and vegetables not available off-season, as well as clothing such as winter coats and sweaters during summer.

3. *Sales*: every time the pollster observes a good with a sale flag in an outlet, she has to mark the price of that particular item with a sale flag.

All pollsters were supervised at least once a month. Supervisors visit some of the outlets, visited earlier the same day by the pollster, and collect a sample of the prices that the pollster should have to collect. Then at the National Statistic Institute, another supervisor compared both forms.

C.4 ROBUSTNESS

In this section we conduct a battery of robustness checks to evaluate the sensitivity of the main results in the paper. The first set of checks consists in dealing with recurrent issues when analyzing micro-price datasets such as missing observations and price changes due to substitutions or sales. Secondly, we discuss issues of aggregation across products. Third, we address biases resulting from discrete sampling. Fourth, we present results using a measure of expected inflation instead of contemporaneous inflation. Finally, we address the possibility that the theoretical propositions which hold in the steady state are a poor description of the Argentinean experience in the high inflation period leading to the stabilization plan in 1991 where agents are likely to have anticipated the strong disinflation that followed. To summarize, the empirical findings at low inflation go through intact. At high inflation, we observe some quantitative but not qualitative differences. Most notable, depending on the estimator used, the elasticity of the frequency of price changes can range from approximately $1/2$ to the theoretical $2/3$.

C.4.1 MISSING DATA, SUBSTITUTION, SALES AND AGGREGATION

This subsection reports the sensitivity of the estimates of the elasticity γ , the semielasticity $\Delta\lambda$ and the duration at low inflation obtained with the simple estimator to different treatments of missing data, sales, product substitution and broad aggregation levels. The results obtained from using the different methodologies for estimating the frequencies of

price adjustment are presented in Table C.4. In Appendix C.4.1.1 we describe in detail the methodologies and the definitions of all estimators in Table C.4.

We report estimates of the three parameters of equation (3.15) for the sample of differentiated goods (monthly), for the sample of homogeneous goods (bi-weekly) and for the aggregate. The latter is obtained by averaging the estimated λ s with their expenditure shares after converting the bi-weekly estimates to monthly ones.

The first and second columns show the elasticities at high and low inflation. The third block of columns shows the implied duration of price spells when inflation is low (below the threshold) under the assumption that the frequency of price adjustment is constant.

The first line in Table C.4 corresponds to the pooled simple estimator reported in Figure 3.6. The estimates of the elasticity of the frequency of price adjustment with respect to inflation, γ , are very similar for the λ s in the two samples and for the aggregate λ . The estimates for the semi-elasticity and expected duration at low inflation are markedly higher for differentiated goods in comparison to homogeneous goods. The other lines in the table provide estimates of the three parameters of interest for different aggregation methods and for the different treatments of missing observations, product substitutions and sales. The values for the elasticity γ across all these estimation techniques ranges from approximately 1/2 to 2/3. The variation in $\Delta\% \lambda$ estimates is much smaller across methodologies. Both differences in the estimators result from alternative aggregation methodologies and not from the treatment of sales, substitutions and missing values. For instance, the elasticity at high inflation when using the simple estimator with pooled data climbs from 0.53 to 0.68 when using the median estimate across industries.

The treatment of sales and substitutions does seem to have an effect on the estimates of the expected duration of price spells when inflation is low, as in other papers in the literature (see Klenow and Malin (2011)). For example, durations increase from 4.5 months to 5.7 months when sales price quotes are replaced by the price quote of the previous regular price. In Klenow and Kryvtsov (2008) durations go from 2.2 months to 2.8 after the sales treatment and in Nakamura and Steinsson (2008) they go from 4.2 months to 3.2. Time series for frequency of substitution, sales and missing values in the sample can be seen in Figure 3.4.

What accounts for the differences in the estimates implied duration at low inflation between the sample of differentiated and homogeneous goods? Expected durations are much higher for differentiated goods than for homogeneous goods. In principle, we believe that

Table C.4: The Frequency of Price Adjustment and Inflation: Robustness Checks.

Aggregation	Elasticity			Semi-Elasticity			Expected Duration		
	γ			at zero $\Delta\% \lambda$			at $\pi = 0$ (months)		
	Diff.	Hom.	Agg	Diff.	Hom.	Agg	Diff.	Hom.	Agg
A. Simple Estimator (No information from missing price quotes)									
Pooled	0.51	0.5	0.53	0.08	-0.01	0.04	9.1	5.8	4.5
Weighted Average	0.52	0.48	0.52	0.07	-0.02	0.04	8.2	5.7	4.4
Median	0.64	0.64	0.68	0.1	0.03	0.05	15	14	9.3
Weighted Median	0.65	0.64	0.68	0.09	0.02	0.04	12.2	11	7.8
Pooled (excluding sales)	0.5	0.47	0.52	0.08	0.03	0.05	10	7.5	5.7
B. All price quotes									
Pooled	0.51	0.5	0.52	0.08	-0.01	0.04	8.8	5.9	5.1
Weighted Average	0.52	0.45	0.49	-0.05	-0.02	0.04	8.9	5.8	4.6
Median	0.62	0.58	0.65	0.09	-0.21	0.02	15.3	8.9	10.2
Weighted Median	0.62	0.65	0.65	0.09	0.02	0.04	12.9	10.8	7.5
C. Excluding substitution quotes									
Pooled	0.55	0.5	0.52	0.09	-0.01	0.03	7	6.3	6.2
Weighted Average	0.52	0.45	0.51	0.06	-0.02	0.03	10.7	6.1	6.7
Median	0.66	0.65	0.68	0.13	0.02	0.02	18.8	17.9	12.4
Weighted Median	0.66	0.62	0.66	0.07	-0.06	0.02	16.5	12.2	10.4
D. Excluding substitution spells									
Pooled	0.52	0.5	0.52	0.07	0	0.05	10.7	6	5.5
Weighted Average	0.53	0.44	0.49	-0.1	-0.02	0.04	8.6	5.8	4.8
Median	0.62	0.64	0.64	0.09	0.04	-0.02	18.4	16.6	10.9
Weighted Median	0.63	0.6	0.66	0.09	-0.13	-0.05	15.4	9.4	8.2
E. Excluding substitution and sales quotes									
Pooled	0.5	0.47	0.52	0.08	0.06	0.05	9.9	5.2	6.7

Note: *Diff.* denotes differentiated goods, which are samples once a month. *Hom.* denotes homogeneous goods, which are sampled twice a moth. *Agg* denotes the weighted average of the Differentiated and homogeneous goods, with weights given by the expenditure shares and where the homogeneous goods have been aggregated to monthly frequencies. For each case we use NLLS to fit: $\log \lambda_t = a + \epsilon \min \{ \pi_t - \pi^c, 0 \} + \nu (\min \{ \pi_t - \pi^c, 0 \})^2 + \gamma \max \{ \log \pi_t - \log \pi^c, 0 \} + \omega_t$. The semi-elasticity at zero $\Delta\% \lambda$ is the percentage change in λ when inflation goes from 0 to 1%. A. estimates λ with the simple estimator in equation (3.14) discarding information from missing prices, B. is the full information maximum likelihood estimator described in Appendix C.4.1.1, C. replaces price quotes with a product substitution by missing data, D. replaces price spells ending in a product substitution by missing data and E. replaces sales quotes by the previous price and product substitutions by a missing quote.

this discrepancy can be attributed to two features: an intrinsic difference between the type of goods or due to the fact that the prices of homogeneous goods are sampled bi-monthly and prices for differentiated goods once a month.

C.4.1.1 MISSING DATA, SUBSTITUTIONS, SALES AND AGGREGATION

Here we describe in detail the methodologies and the definitions of all estimators in Table C.4 in Section C.4.1.

We start by describing the assumptions used to estimate the probability of a price change when we observe missing prices. If between two observed prices there are some missing prices we use the following assumption. If the two observed prices are exactly equal we assume there has been *no* changes in prices in any times between these dates. This is the same assumption as in Klenow and Kryvtsov (2008). Instead, if the two observed prices are different we assume there has been *at least one change* in prices in between. The first assumption allow us to complete the missing prices in between two observed prices that are equal. From here on, assume that the missing prices in such string of prices have been replaced.

We will refer from now on to the sequence of prices between two different observed prices as a *spell of constant prices*, or for short a *spell of prices*. Without any missing prices, a spell of constant prices is just a sequence of repeated prices ending with a different price. Notice that the last (observed) price in a spell of constant prices is the first price of the next spell.

Next we describe the possible patterns of prices, and its implications for the estimation of the probability of a price change. After following the procedure described above, all spell of prices and missing observations have only two possible patterns. The first pattern is a spell of prices ending with a price change, but with no missing observations. We consider the second pattern in the following section, where we deal with the effect of missing prices. Consider the following example for a spell for an outlet i that contains no missing prices nor substitutions:

The braces on top of the values of λ are meant to remind the reader that λ_t refers to the constant probability of change in prices between $t - 1$ and t . The indicator I_{it} adopts the value one if, in outlet i , the price in period t is different from the price at period $t - 1$, and zero otherwise, except for the first temporal period of the first string of prices where it is missing¹. In this example we have exactly no changes for the first four periods and at least

¹We also include the indicator γ_t , which we explain below, for completeness.

Table C.5: Example of a spell of constant prices without missing prices

p_t	10	10	10	10	10	15
	t	$t + 1$	$t + 2$	$t + 3$	$t + 4$	$t + 5$
	$\overbrace{\lambda_{t+1}}$	$\overbrace{\lambda_{t+2}}$	$\overbrace{\lambda_{t+3}}$	$\overbrace{\lambda_{t+4}}$	$\overbrace{\lambda_{t+5}}$	
I_t	y	0	0	0	0	1
γ_t	y	1	1	1	1	1

one change in the next period. The probability of observing this completed spell of constant prices is thus:²

$$P = e^{-\lambda_{t+1}} \times e^{-\lambda_{t+2}} \times e^{-\lambda_{t+3}} \times e^{-\lambda_{t+4}} \times (1 - e^{-\lambda_{t+5}}) \quad (\text{C.31})$$

It follows that in this simple case, assuming all the outlets selling the same good have the same λ_j , the likelihood function of prices observed for product j is

$$L_j = \prod_{i \in O_j} [e^{-\lambda_{j,t}}]^{(1-I_{it})} \times [1 - e^{-\lambda_{j,t-\tau}}]^{I_{it}} \quad (\text{C.32})$$

The maximum likelihood estimator of the arrival rate of price changes for product j between times t and $t+1$ in the simple case without missing prices and without substitutions is

$$\lambda_{j,t+1} = \log \left(\frac{\sum_{i \in O_j} 1}{\sum_{i \in O_j} (1 - I_{it})} \right) = -\log(1 - f_{jt}) \quad (\text{C.33})$$

where we let O_j denote the set of the outlets i of the product j and f_{jt} is the fraction of outlets of good j that changed prices in period t . In words, $\lambda_{j,t+1}$ is the log of the ratio of the number of outlets to the number of outlets that have not changed the price between t and $t + 1$. Thus $\lambda_{j,t+1}$ ranges between zero, if no outlets have change prices, and infinite if all outlets have changed prices. The probability of at least one change in prices in period t for product j is $1 - e^{-\lambda_{jt}} = f_{jt}$.

²We assume that the number of price changes between $t - 1$ and t follows a homogeneous Poisson process with arrival rate λ_t per unit of time. The probability of k occurrences is $e^{-\lambda} \lambda^k / k!$ and the waiting time between occurrences follows an exponential distribution.

C.4.1.2 INCORPORATING INFORMATION ON MISSING PRICES

Now we consider the case where there are some missing prices before the price change, but we postpone the discussion of the effect of price substitutions.

In general, a spell of constant prices is a sequence of $n + 1$ prices that starts with an observed price p_t , possibly followed by a series of prices all equal to p_t , then followed, possibly, by a series of missing prices, that finally ends with an observed price at p_{t+n} that differs from the value of the initial price p_t . Notice that while we also refer to this sequence of prices as a spell of constant prices, it can include more than one price change if there were missing observations, a topic to which we return in detail below.

To deal with missing prices, the interesting patterns for a spell of constant prices are those which end with a price change, but that contain some missing price(s) just before the end of the spell. For example, consider a spell of prices for an outlet, i , in which there are exactly no changes in the first four periods and at least one change sometime during the next three periods.

Table C.6: Example of a spell of constant prices with observed and missing prices

p_t	10	10	10	10	10	m	m	15
	t	$t + 1$	$t + 2$	$t + 3$	$t + 4$	$t + 5$	$t + 6$	$t + 7$
		$\underbrace{\hspace{1.5cm}}$	$\underbrace{\hspace{1.5cm}}$	$\underbrace{\hspace{1.5cm}}$	$\underbrace{\hspace{1.5cm}}$	$\underbrace{\hspace{1.5cm}}$	$\underbrace{\hspace{1.5cm}}$	$\underbrace{\hspace{1.5cm}}$
I_t	y	0	0	0	0	x	x	1
γ_t	y	1	1	1	1	x	x	3
χ_t	0	0	0	0	0	1	1	0

In Table C.6 an x means that the variable is not defined for that case, the m denotes a missing/imputed price and the y denotes that for indicator I and counter γ the first observation in the spell of prices is missing because it depends on the prices in period $t - 1$ which are not in the information of the table. The probability of observing this spell is:

$$P = e^{-\lambda_{t+1}} \times e^{-\lambda_{t+2}} \times e^{-\lambda_{t+3}} \times e^{-\lambda_{t+4}} \times (1 - e^{-\lambda_{t+5}} \times e^{-\lambda_{t+6}} \times e^{-\lambda_{t+7}}) \quad (\text{C.34})$$

The first four products are the probability of no change during the first four periods, and the last term is the probability of at least one change during the last three periods. The second term is the complement of the probability of no change in prices during the last three periods.

The likelihood of the sample of T periods (with $T + 1$ prices) of all the outlets for the good j –denoted by the set O_j – is the product over all outlets i of the product of all spells for outlet i of the probability equation (C.34). To write the likelihood we define an indicator, χ_{it} , and a counter γ_{it} . The indicator χ_{it} adopts the value one if a price is missing for outlet i in period t , and zero otherwise. The value of γ_{it} counts the number of periods between two non-missing prices. The counter γ_{it} is Klenow and Kryvtsov (2008) duration clock . Then the likelihood function of the prices observed for product j is:

$$L_j = \prod_{i \in O_j} \prod_{t=1}^T [e^{-\lambda_{j,t} \gamma_{it}}]^{(1-I_{it})(1-\chi_{it})} \times \left[1 - e^{-\sum_{\tau=0}^{\gamma_{it}-1} \lambda_{j,t-\tau}} \right]^{I_{it}(1-\chi_{it})} \quad (\text{C.35})$$

Since the λ 's are the probability of a price change and they are indexed at the end of a period, the first temporal observation of prices at $t = 0$ does not enter the likelihood. The log likelihood is:

$$\ell_j = \sum_{i \in O_j} \left(\sum_{t=1}^T (1 - \chi_{it})(1 - I_{it}) \times (-\lambda_{j,t} \gamma_{it}) + \sum_{t=1}^T (1 - \chi_{it}) I_{it} \ln \left[1 - e^{-\sum_{\tau=0}^{\gamma_{it}-1} \lambda_{j,t-\tau}} \right] \right) \quad (\text{C.36})$$

To compute the contribution to the likelihood of a given value of $\lambda_{j,t}$ for $t = 1, \dots, T$ we find convenient to introduce two extra counters: κ_{it} and η_{it} for any period t in which prices are missing/imputed. The variable κ_{it} counts the number of periods since the beginning of a string of missing/imputed prices. The variable η_{it} counts the number of periods of missing/imputed prices until the next price is observed. For example, consider the string of prices in Table C.7.

Table C.7: Example of a spell of constant prices, w/ missing prices and counters

p	10	10	m	m	m	m	m	15
	t	$t + 1$	$t + 2$	$t + 3$	$t + 4$	$t + 5$	$t + 6$	$t + 7$
	λ_{t+1}	λ_{t+2}	λ_{t+3}	λ_{t+4}	λ_{t+5}	λ_{t+6}	λ_{t+7}	
I	y	0	x	x	x	x	x	1
χ	0	0	1	1	1	1	1	0
γ	y	1	x	x	x	x	x	6
κ	x	x	1	2	3	4	5	x
η	x	x	5	4	3	2	1	x

Table C.7 shows an example of a spell of constant prices for a given variety and a given

outlet. In equation (C.37) we highlight the contribution to the log-likelihood of the value of λ_t for a given outlet i :

$$\begin{aligned} \ell_j = & \dots + (1 - \chi_{it})(1 - I_{it}) \times (-\lambda_{j,t}\gamma_{it}) + (1 - \chi_{it})I_{it} \ln \left[1 - e^{-\sum_{\tau=0}^{\gamma_{it}-1} \lambda_{j,t-\tau}} \right] \\ & + \chi_{it} \ln \left[1 - e^{-(\sum_{\tau=0}^{\kappa_{it}-1} \lambda_{j,t-\tau} + \sum_{\tau=1}^{\eta_{it}} \lambda_{j,t+\tau})} \right] + \dots . \end{aligned} \quad (\text{C.37})$$

The first two terms have the contribution to the likelihood of $\lambda_{j,t}$, if the price at time t is not missing. The first case corresponds to a price at the beginning of the spell of constant prices. The second to a case where the price is the last one of the spell, and uses γ_{it} to be able to write the corresponding probability. The third term, correspond to the contribution of $\lambda_{j,t}$, if the price at time t is missing, and uses κ_{it} , and η_{it} to write the corresponding probability.

Using equation (C.37) it is easy to write the FOC with respect to $\lambda_{j,t}$ of the sample as:

$$\begin{aligned} \frac{\partial \ell_j}{\partial \lambda_{j,t}} = & \sum_{i \in O_j} (1 - \chi_{it}) (1 - I_{it}) \times (-\gamma_{it}) + \sum_{i \in O_j} (1 - \chi_{it}) I_{it} \frac{1}{e^{(\sum_{\tau=0}^{\gamma_{it}-1} \lambda_{j,t-\tau})} - 1} \\ & + \sum_{i \in O_j} \chi_{it} \frac{1}{e^{(\sum_{\tau=0}^{\kappa_{it}-1} \lambda_{j,t-\tau} + \sum_{\tau=1}^{\eta_{it}} \lambda_{j,t+\tau})} - 1} = 0 , \end{aligned} \quad (\text{C.38})$$

for $t = 1, \dots, T$.

Roughly speaking the estimator for $\lambda_{j,t}$ computes the ratio of the number of outlets i that a time t have change the prices with those that have not change prices or that have missing prices. This approximation is exact if no outlet has a missing/imputed price at period t or before. In this case $\chi_{is} = 0$ for $s = 1, 2, \dots, t$ for all $i \in O_j$, then equation (C.38) becomes

$$\sum_{i \in O_j} (1 - I_{it}) \times \gamma_{it} = \sum_{i \in O_j} \left[\frac{I_{it}}{e^{\sum_{\tau=0}^{\gamma_{it}-1} \lambda_{j,t-\tau}} - 1} \right] ,$$

which, if we make $\lambda_{j,t} = \lambda_{j,t-1} = \dots = \lambda_{j,t-\gamma_{it}+1}$ is the same expression than in Klenow and Kryvtsov (2008). In the case of no missing observations, and where all the λ 's are assumed to be the same, this maximum likelihood estimator coincide with the simple estimator introduced in equation (C.33).

C.4.1.3 INCORPORATING MISSING PRICES AND SALES

Our data contains a flag indicating whether an item was on sale. We consider a procedure that disregards the changes in prices that occur during a sale. The idea behind this procedure is that sales are anomalies for the point of view of some models of price adjustment, and hence they are not counted as price changes. To explain this assumption we write an hypothetical example:

Table C.8: Example of a spell of constant prices removing sales

p_t	10	10	10	m	8	15
	t	$t + 1$	$t + 2$	$t + 3$	$t + 4$	$t + 5$
		λ_{t+1}	λ_{t+2}	λ_{t+3}	λ_{t+4}	λ_{t+5}
I_t	y	0	0	0	0	1
a_t	y	0	0	0	1	0

In Table C.8 the indicator a_t takes the value of one if the good is on sale on period t and if the price at the time t is smaller than the previous recorded price. In this case, the price in periods $t + 3$ and $t + 4$ is changed to 10, the value of the previous recorded price. The general principle is to consider a string of recorded prices, possible missing values, and a price that has a sale flag, and replace the price of the string of missing values and the period with a sale flag for the previous recorded price. In other words, we replace the price at the period with a sale flag for the previous recorded price and then using our first assumption on missing prices we complete the missing price in between two prices that are equal. When the sales are disregarded the number and duration of price spells can change. Once this procedure is implemented, the likelihood is the one presented above using the modified price series -indeed in Table C.8 the indicator I_t is the one that corresponds to the modified price string. We refer to the corresponding estimates as those that *excludes sale quotes*. This is a procedure used by many, e.g., Klenow and Kryvtsov (2008). By construction with this method the estimates for λ will be smaller.

C.4.1.4 INCORPORATING MISSING PRICES AND PRICE SUBSTITUTIONS

In this section we discuss different assumptions on the treatment of missing data and good substitutions that allow us to construct four estimators of the frequency of price changes that we report later on.

We use the indicators \tilde{e}_{it} , e_{it}^* , \bar{e}_{it} to consider two different assumptions on how to treat a price spell that ends in some missing values or price substitutions. In particular, consider the case of a substitution of a product or a missing price. As explained above, our data set contains the information of whether the characteristic of the product sold at the outlet has changed and was subsequently substituted by a similar product. We also have information on whether the price is missing (mostly due to a stock-out). To be concrete, consider the following example of a spell of constant prices in Table C.9. In this table, and $s_{it} = 1$ denotes the period where a price substitution has occurred. Thus, the example has a spell of 9 prices, with two periods (3 prices) with no change in prices, then 5 periods with missing prices, and finally in the last period there is an observed price that correspond to a substitution of the good.

Table C.9: Example: spell of constant prices w/ counters for substitutions

p	10	10	10	m	m	m	m	15
	t	$t + 1$	$t + 2$	$t + 3$	$t + 4$	$t + 5$	$t + 6$	$t + 7$
	λ_{t+1}	λ_{t+2}	λ_{t+3}	λ_{t+4}	λ_{t+5}	λ_{t+6}	λ_{t+7}	
I	y	0	x	x	x	x	x	1
χ	0	0	0	1	1	1	1	0
s	y	0	0	0	0	0	0	1
\tilde{e}	y	0	0	0	0	0	0	0
e^*	y	1	1	1	1	1	1	1
\bar{e}	y	0	0	1	1	1	1	1

As in the previous examples, for the first observation an y indicates that the value of the indicator cannot be decided based on the information in the table.

The issue is the interpretation of when and whether there has been a price change in the previous price spell. One interpretation is that there has been a price change somewhere between periods $t+2$ and $t+7$. A different interpretation is that, because the price spell ends with the substitution of the good, the price has not changed. The idea for this interpretation

is that if the good would have not changed, the price could have been constant beyond $t + 7$. The next three cases explain how to implement the first interpretation, and two ways to implement the second one. The last two cases present two simple estimators, one that treats quotes with substitutions as regular price changes, and one that exclude them.

1. We disregard the information of the substitution of a good, and proceed as we have done so far: including all price quotes as if the good have been not changed. In the previous example, it consists on assuming that the price has changed between periods $t + 2$ and $t + 7$. In this case we say that the probability of observing the spell in the table is given by:

$$P = e^{-\lambda_{t+1}} \times e^{-\lambda_{t+2}} \times (1 - e^{-\lambda_{t+3} - \lambda_{t+4} - \lambda_{t+5} - \lambda_{t+6} - \lambda_{t+7}}) \quad (\text{C.39})$$

In this case we set $\tilde{e}_{it} = 0$ for all periods, since we don't want to exclude any part of this price spell. We refer to these assumptions as *including all price quotes*.

2. We follow Klenow and Kryvtsov (2008) and others and exclude completely any spell of prices that ends with a substitution of a product. We implement this by defining the indicator $e_{it}^* = 1$ for any price corresponding to this spell, i.e., we exclude all the observations –with the exception of the first price, which is relevant for the previous string. In this case we have no associated probability for this spell. The idea behind this treatment is that if there would have been no substitution of the good the price could have stayed constant even beyond $t + 7$. We refer to these assumptions as *excluding substitution spells* for short.
3. We introduce a new way to handle this information based on the following underlying assumption: if a spell of constant prices ends up in a substitution we interpret that the product has changed, and hence we cannot infer from the observed price whether the price has changed or not, as in the previous case. Yet, unlike the previous case, we do not discard the information at the beginning of the string, where the product was the same and its price was not changing. In this case the associated probability for this spell is:

$$P = e^{-\lambda_{t+1}} \times e^{-\lambda_{t+2}} \quad (\text{C.40})$$

In this case we use the indicator $\bar{e}_{it} = 0$ for the first two observations, among which we know that there was no price change, and $\bar{e}_{it} = 1$ for the rest of the observations where we can't conclude if there was a price change for the *same* product. We refer to these assumptions as *excluding substitution quotes* for short.

4. We present an alternative estimator to the maximum likelihood, which has the advantage of being simpler to describe and understand. This estimator imitates the one for the case where there are no missing price quotes in equation (C.33), and simply excludes the values of the missing. In this case, this estimator is:

$$\lambda_{j,t+1} = \log \left(\frac{\sum_{i \in O_j} (1 - \chi_t) (1 - \chi_{t+1})}{\sum_{i \in O_j} (1 - \chi_t) (1 - \chi_{t+1}) (1 - I_{it})} \right) \quad (\text{C.41})$$

In words, $\lambda_{j,t+1}$ is a non-linear transformation of the probability of the change of prices. This probability is estimated as the ratio of the outlets that have changed prices over all the outlets, including only those price quotes that are simultaneously not missing at t and $t + 1$. While this estimator is simpler than the maximum likelihood, it does not use all the information of the missing values efficiently. We refer to this estimator as the *simple* estimator.

Finally, to completely state the notation for the likelihood function, we use the indicator ζ to deal with missing prices at the beginning of the sample. In particular, if for an outlet i the sample starts with $n + 1$ missing prices, we exclude these observations from the likelihood since we cannot determine the previous price. We do this by setting the indicator $\zeta_{it} = 1$ for $t = 0, 1, \dots, n$. Thus, depending of the assumption, the exclusion indicator e takes the values given by \tilde{e} , \bar{e} or e^* , besides the value of 1 for all the missing observations at the beginning of the sample. The log likelihood function is then:

$$\ell_j = \sum_{i \in O_j} \sum_{t=1}^T (1 - e_{it}) \left((1 - \chi_{it})(1 - I_{it}) (-\lambda_{j,t} \gamma_{it}) + (1 - \chi_{it}) I_{it} \ln \left[1 - e^{-\sum_{\tau=0}^{\gamma_{it}-1} \lambda_{j,t-\tau}} \right] \right) \quad (\text{C.42})$$

The first order condition for $\lambda_{j,t}$, using the counters κ and η is:

$$\begin{aligned} \frac{\partial \ell_j}{\partial \lambda_{j,t}} &= \sum_{i \in O_j} (1 - \chi_{it}) (1 - e_{it}) (1 - I_{it}) \times (-\gamma_{it}) + \\ &\sum_{i \in O_j} (1 - \chi_{it}) (1 - e_{it}) I_{it} \frac{1}{e^{(\sum_{\tau=0}^{\gamma_{it}-1} \lambda_{j,t-\tau})} - 1} \\ &+ \sum_{i \in O_j} \chi_{it} (1 - e_{it}) \frac{1}{e^{(\sum_{\tau=0}^{\kappa_{it}-1} \lambda_{j,t-\tau} + \sum_{\tau=1}^{\eta_{it}} \lambda_{j,t+\tau})} - 1} = 0, \end{aligned} \quad (\text{C.43})$$

for $t = 1, \dots, T$.

C.4.1.5 COMPUTATION OF ESTIMATES AND STANDARD ERRORS OF λ

We compute the maximum likelihood estimator by using an iterative procedure. In this discussion we fix a good or item. We denote the iterations by superindex j . The initial guess for λ_t^0 is the log of the ratio of the outlets that change the price between $t - 1$ and t . Then we solve for λ_t^{j+1} in equation (C.43) for the foc of λ_t taking as given the values of λ_{t-i}^j and λ_{t+i}^j for all $i \neq 0$. Notice that this equation has a unique solution. Also notice that the solution can be ∞ , for instance if all prices change, or zero.

We have not derived expression for standard error for the different estimator of λ . Nevertheless to give an idea of the precision of our estimator, we note that for the simple pooled estimator, assuming that missing values are independent, we can use the expression for the standard error of a binomial distributed variable for the probability of a price change π , while using the estimate $\hat{\pi}_t = \text{number of outlets with a price change} / N_t$ where N_t is the number of outlet with a price quote between t and t_1 . In this case $se(\hat{\pi}) = \sqrt{\hat{\pi}(1 - \hat{\pi})/N}$. Using the delta method and $\lambda(p) = -\log(1 - p)$ then we have:

$$se(\hat{\lambda}) = \frac{\sqrt{\exp(\hat{\lambda}) - 1}}{\sqrt{N}} \quad \text{and} \quad se(\log \hat{\lambda}) = \frac{1}{\hat{\lambda}} \frac{\sqrt{\exp(\hat{\lambda}) - 1}}{\sqrt{N}}. \quad (\text{C.44})$$

To give an idea of the magnitudes for our estimated parameters we use some round numbers for both homogeneous and differentiated goods. For the case where we pool *all* the differentiated goods we can take $N_d = 230 \times 80 \times 0.80 = 14720 \approx \# \text{ diff. goods} \times \text{avg. \# outlets diff. goods} \times \text{fraction of non-missing quotes}$. Pooling *all* the homogeneous goods

we $N_h = 60000 = 300 \times 250 \times 0.80 \approx \# \text{ homog. goods} \times \text{avg. \# outlets homog. goods} \times$ fraction of non-missing quotes. So for high value of λ such as for $\log \hat{\lambda} = \log 5$, we have that the standard errors are $se(\log \hat{\lambda}_d) = 0.02$ and $se(\log \hat{\lambda}_h) = 0.010$. Instead for low values of λ such as for $\log \hat{\lambda} = \log 1/5$, we have that the standard errors are $se(\log \hat{\lambda}_d) = 0.019$ and $se(\log \hat{\lambda}_h) = 0.009$. Instead if we estimate λ at the level of each good, the standard errors should be larger by a factor $\sqrt{230} \approx 15$ and $\sqrt{330} \approx 17$ for differentiated and homogeneous goods respectively.

C.4.1.6 AGGREGATION: WEIGHTED AVERAGE, MEDIAN, WEIGHTED MEDIAN AND POOLED MAXIMUM LIKELIHOOD

In this section we deal with the issue of aggregation. So far we have described how to estimate the frequency of price adjustment for each good category separately.

Remember that those goods that fall in the homogeneous goods category are sampled bi-monthly and so will be our estimates. In this way, the first step in order to aggregate all categories is to convert them into monthly estimates, which is done simply by adding the two estimates in any given month (this results from the exponential assumption for our likelihood function).

Next, we compute three aggregated estimations. First, we calculate the weighted average of all monthly estimates (both differentiated and homogeneous goods), where the weights are the corresponding expenditure shares of each good category.

$$\lambda_t = \sum_{i=1}^N \omega_i \lambda_{it}$$

For future reference, we will call this *weighted average* or simply *WA*, followed by the specific treatment of missing, sales and substitutions. For example, *Weighted average excluding sales*.

The other two estimates are the median and weighted median of all monthly estimates (both differentiated and homogeneous goods). The aggregated median estimation (*median* for short) consists in taking the median λ of all products at each time period. The aggregate weighted median estimation is computed by sorting, at each time period, the expenditure weights of each product by the value of their associated λ from lowest to highest. Then we compute the accumulated sum of the weights until reaching 0.5. The aggregate weighted

median of the frequency of price adjustment is the associated λ of the product whose weight makes the accumulated sum equal to or greater than 0.5. We refer to this estimate as *weighted median*.

Finally we consider a last aggregated estimation. As mentioned, the estimates for the frequency of price changes presented allow the value of λ to depend on the time period and the good. We now consider an estimate based on the assumption that the frequency price changes is common for all goods, but that it can change between time periods. This simply puts together the outlets for all goods in our sample. Thus, the log likelihood is:

$$\ell(\lambda_1, \dots, \lambda_T) = \sum_{j=1}^N \ell_j(\lambda_1, \dots, \lambda_T)$$

where the $\ell_j(\lambda_1, \dots, \lambda_T)$ corresponding to the log likelihood for each assumption about missing prices and or price substitutions as it has been introduced in the previous sections, and where N is the number of goods in our sample. We refer to this estimator as the *pooled maximum likelihood* or for short PML. Likewise, when we assume that all goods have the same frequency of price changes but use the simple estimator for λ , we refer to it as *simple pooled* estimator.

C.4.2 ESTIMATION OF THE ELASTICITIES AT A MORE DISAGGREGATED LEVEL

The theory predicts, and the data confirms, that the frequency of price adjustment is very different across goods. In fact, in Table 3.1 we observe a considerable heterogeneity in the *simple* estimator of λ across 5-digit level industries. In this subsection we explore the robustness of the parameter estimates for the elasticity of the frequency of price changes with respect to inflation at high and low inflation rates by fitting equation (3.15) for each of the 5-digit industries, using the simple estimator of λ .³

Table C.10 presents statistics describing the distribution of the coefficient estimates derived from equation (3.15) across 5-digit industries. The elasticity estimates confirm our previous findings: (i) the elasticity of the frequency of price changes at high inflation, γ , varies between 1/2 and 2/3; and (ii) the semi-elasticity $\Delta\% \lambda$ is approximately zero regardless of the industry. Consistent with the results in Table 3.1, there is large variation in the expected duration at low inflation; particularly so for differentiated goods.

³We performed the same exercise at a 6-digit level obtaining qualitatively similar results. See for Table C.1 disaggregation levels

Table C.10: Distribution of fitted coefficients at 5 digit level

	Elasticity		Semi-elasticity		Duration	
	γ		at zero $\Delta\% \lambda$		at $\pi = 0$	
	Diff.	Hom.	Diff.	Hom.	Diff.	Hom.
Mean	0.58	0.56	0.03	0.01	18.5	7.5
Median	0.58	0.55	0.03	0	15	6.4
Perc 75	0.7	0.6	0.1	0.02	25.6	9.3
Perc 25	0.48	0.5	-0.03	-0.02	6.9	4.7
Std. Dev.	0.14	0.08	0.1	0.07	14.2	5.2

Note: *Diff.* denotes differentiated goods. *Hom.* denotes homogeneous goods. For each 5-digit industry we use NLLS to fit: $\log \lambda_t = a + \epsilon \min \{\pi_t - \pi^c, 0\} + \nu (\min \{\pi_t - \pi^c, 0\})^2 + \gamma \max \{\log \pi_t - \log \pi^c, 0\} + \omega_t$. The semi-elasticity at zero $\Delta\% \lambda$ is the percentage change in λ when inflation goes from 0 to 1%. λ is estimated with the simple estimator in equation (3.14).


C.4.3 SAMPLING PERIODICITY

So far we have been using the estimator of the theoretical frequency λ_a that has been proposed in the literature, $\hat{\lambda}_t = -\ln(1 - f_t)$, where f_t is the fraction that changed outlets in $(t - 1, t]$. If price changes follow a Poisson process this is the maximum likelihood estimator of λ_a . However, since we only observe f_t at discrete times, a well known bias may arise if prices change more than once within the time interval and these changes are not independent. In particular, we would expect the bias to become larger as inflation increases and prices change more frequently.

In this section, we consider an alternative estimator $\lambda_t^{SW} = f_t$ and compare it to $\hat{\lambda}_t$. In the Sheshinski-Weiss model with no idiosyncratic shocks (or in the limit as inflation becomes very large compared to the volatility of the shocks) this is a maximum likelihood estimator of λ_a .

In Figure C.3 we present the results of Monte-Carlo simulations using data generated by the model in Section 3.2. We sample observations every two-weeks and calculate both estimators of λ_a for different inflation rates. The true frequency of price adjustment, λ_a , is represented by the red line, the frequency, f , by the green line and the simple estimator $\hat{\lambda}_t$ by the blue line. The figure points to the existence of an upward bias in $\hat{\lambda}_t$ for high inflation rates, and as such, in the elasticity of the frequency of price adjustments to inflation.

Figure C.3: Estimators $\hat{\lambda}$ and f , sampled every two weeks




To reduce the incidence of such bias in our empirical estimates we proceed by re-estimating the elasticity of the frequency of price changes to inflation by excluding observations corresponding to inflation above some threshold. Figure C.4 illustrates the results for a threshold inflation of 200 percent using simple estimators of λ and f . The estimated elasticities are 0.63 and 0.48 respectively, much in line with our benchmark estimates.⁴

C.4.4 EXPECTED INFLATION

The robustness check in this section pertains to the measure of expected inflation as opposed to our estimates of the frequencies of price adjustment. Theoretical models of price setting

⁴Excluding observations corresponding to inflation below 50 percent and above 200 percent or below 50 percent and above 100 percent result in estimated elasticities of 0.64 and 0.76 when using $\hat{\lambda}$; when using f instead, these are 0.44 and 0.59

Figure C.4: Estimators λ and f , homogeneous goods sampled twice a month



behavior, such as the menu cost model presented in Section 3.2, predict that firms decision to change prices, and the magnitude of such change, depend on *expected* inflation between adjustments. Our measure of expected inflation so far was the average realized inflation for the expected duration of the price set in period t , $1/\lambda_t$.

We now consider a more flexible form for expected inflation as an average of the actual inflation rate of the following k_t months, where $k_t = [n/\lambda_t]$ and $[x]$ is the integer part of x ; that is

$$\pi_t^e = \frac{1}{k_t} \sum_{s=t}^{t+k_t} \pi_s \quad (\text{C.45})$$

We refer to n as the *forward looking factor*. Thus, as inflation falls (and implied durations rise) in our sample agents put an increasing weight on future inflation. When $n = 0$ expected and actual inflation are the same.

Table C.11: Elasticities and implied duration for different estimates of expected inflation.

Forward looking factor	Elasticity γ	Semi-elasticity at zero $\Delta\% \lambda$	Duration at $\pi = 0$	R^2
$n = 0$	0.57	0.01	4.4	0.959
$n = 0.5$	0.54	0.03	4.4	0.949
$n = 1$	0.53	0.04	4.5	0.953
$n = 1.5$	0.48	0.04	4.6	0.946
$n = 2$	0.46	0.04	4.7	0.944

Table C.11 shows that the results presented in Section 3.5.1 are not very sensitive to estimating equation (3.15) using different forward looking factors in equation (C.45). The first row of the table shows the results when we use the actual rate of inflation⁵. The following four rows shows that as expectations are more forward looking the threshold inflation and the elasticity fall slightly and the implied duration at the threshold remains constant. The R^2 of the regression is maximized for $n = 1$, so that the best fit is when we assume agents set expected inflation equal to the average of the perfect foresight future inflation rates for the expected duration of the prices. All the estimates of the elasticity γ are consistent with the theoretical prediction in Section 3.2.

C.5 BACKGROUND ON INFLATION AND ECONOMIC POLICY

Here we give a brief chronology of economic policy to help readers understand the economic environment in our sample period, which goes from December 1988 to September 1997. The beginning of our sample coincides with decades of high inflation culminating in two years of extremely high inflation (typically referred as two short hyperinflations) followed by a successful stabilization plan, based on a currency board, started in April of 1991 which brought price stability in about a year, and stable prices until at least three more years after the end of our sample.


The years before the introduction of the currency board witnessed several unsuccessful stabilization plans, whose duration become shorter and shorter, and that culminated in the two short hyperinflations, all of these during a period of political turmoil. Several sources

⁵We use the inflation rate estimated by INDEC for the whole CPI instead of the inflation rate in our sample because we need forward looking values of inflation at the end of the sample that otherwise we cannot obtain.

describe the inflation experience of Argentina since the 1970, such as Kiguel (1991) and Alvarez and Zeldes (2005) for the period before 1991 and Cavallo and Cottani (1997) for descriptions right after 1991. For a more comprehensive study see Buera and Nicolini (2010).

Argentina had a very high average inflation rate since the beginning of the 1970s. Institutionally, the Central Bank has been part of the executive branch with no independent powers, and typically has been one of the most important sources of finance for a chronic fiscal deficit. Figure C.5 plots inflation, money growth and deficits between 1960 and 2010 with our sample period highlighted in yellow. The deficit as a percentage of GDP was on average well above 5% from 1975 to 1990, see Figure C.5. At the beginning of the 1980s fiscal deficit and its financing by the Central Bank were large even for Argentinean standards. These years coincide with a bout of high inflation that started in the second half of the last military government, 1980 to 1982, and continued during the first two years of the newly elected administration of Dr. Alfonsín, 1983 to 1984.

Figure C.5: Inflation, Money Growth and Deficits



In June of 1985 there was a serious attempt to control inflation by a new economic team which implemented what it is referred to as the Austral stabilization plan (the name comes from the introduction of the “Austral” currency in place of the “Argentine Peso”). The core of this stabilization plan was to fix the exchange rate, to control the fiscal deficit and its financing from the Central Bank, and to introduce price and wage controls. While the Austral plan had some initial success, reducing the monthly inflation rate from 30% in June of 1985 to 3.1% in August of 1985, by mid 1986 the exchange rate was allowed to depreciate every month and inflation reached about 5% per month. By July of 1987 the monthly inflation rate was already above 10%. The same economic team started what is referred to as the “Primavera” stabilization plan in October 1988, when the inflation rate was again around 30% per month, at a time when the Alfonsin administration was becoming politically weak. The primavera plan was a new short lived exchange rate based stabilization plan that was abandoned in February of 1989⁶. Our data set starts right after the beginning of the Primavera stabilization plan, in December of 1988.

After the collapse of the “Plan Primavera” the economy lost its nominal anchor and a perverse monetary regime was in place. Legal reserve requirements for banks were practically 100% and the Central Bank paid interest on reserves (most of the monetary base) printing money. Thus a self fulfilling mechanism for inflation was in place. High inflationary expectations, led to high nominal deposit rates, which turned into high rates of money creation that validated the inflationary expectations.

Figure C.6 displays the yearly percent continuously compounded inflation rate and interest rate for the first years of our sample, together with references to some of the main changes in economic policy during the period. Observe how interest rates and inflation skyrocketed after the plan Primavera’s crawling peg was abandoned. In May 1989 a presidential election took place where the opposition candidate, Dr. Menem, was elected. The finance minister and the central bank president that carried the Primavera plan resigned in April 1989. Thereafter, Dr. Alfonsin’s administration had two different finance ministers and two different central bank presidents, in the midst of a very weak political position and rampant uncertainty about the policies to be followed by the next administration. During the campaign for the presidential election Dr. Menem proposed economic policies that can be safely characterized as populist, with a strong backing from labor unions. Indeed, the core

⁶The peg started at about 12 units of Argentine currency (“the Austral”) per US dollar. To put it in perspective, at the beginning of the Austral plan, the peg was 0.8 units of Argentine currency per US dollar.

of his proposed economic policy was to decree a very large generalized wage increase, “el salariazó”.⁷ In July 1989 the elected president, Dr Menem, took office, several months before the stipulated transition date, in the midst of uncontrolled looting, riots and extreme social tension. The inflation rate at this time was the highest ever recorded in Argentina, almost 200% per month—2.3% per day.


The beginning of the Menem administration started with a large devaluation of the Argentine currency, in what is known as the BB stabilization plan, for the name of the company Bunge and Born, where the two first secretaries of the treasury came from. Indeed these appointments made by the Menem administration were a surprise to most observers, given the promises made in the campaign. The announcement of tight control of the fiscal deficit, and the management of the exchange rate of this plan were also surprising for most observers. During this time inflation transitorily fell.

In December 1989, amid large losses in the value of the Argentine peso, a new finance minister was appointed, Dr. Erman Gonzalez, who started yet a new “stabilization plan” (referred to as Plan Bonex). The core of this plan was a big compulsory open market operation by which the central bank exchanged all time deposits in the Banking system (mostly peso denominated time deposits with maturities of less than a month) for 10 year US Dollar denominated government bonds (Bonex 1989). This big open market operation changed the monetary regime and allowed the Central Bank to regain control of the money supply, as the government no longer had to pay interest on money (reserves on time deposits) by printing money. During Dr. Gonzalez tenure there were several fiscal measures aimed at controlling the fiscal deficit. In March 1990 a renewed version of the stabilization plan was launched, with a stricter control of the money supply and of the fiscal deficit. The actual percentage inflation rates during 1989 and 1990 were 4924% and 1342% respectively!

In January of 1991, Dr. Gonzalez resigned and Dr. Cavallo was appointed as finance minister. During the first two months of his tenure there was a large devaluation of the currency and a large increase in the prices of government owned public utility firms. In April 1st of 1991 there was a regime shift that lasted until 2001. The new regime was a currency board that fixed the exchange rate and enacted the independence of the Central Bank, first by means of presidential decrees, and then by laws approved by congress. At this time the Argentinean currency, the Austral, was pegged to the US dollar at 10,000 units per

⁷The slogan to summarize his proposed economic policy was “el salariazó”, i.e., “the huge wage increase” in Spanish.

Figure C.6: Inflation and Depreciation Rate during HyperInflations



USD.⁸

On January 1st 1992 there was a currency reform that introduced a new currency (the Peso Argentino) to replace the Austral, chopping four zeros of the latter (so that one peso was pegged to one dollar, and to 10,000 australes).

There are a host of changes that were introduced at this time, both in term of deregulation and in terms of fiscal arrangements (broadening of the value added tax's base, sale of state owned firms, etc.), which in the first years reduced the size of the fiscal deficit. There was also a renewed access to the international bond markets, and a constant increase in public debt. There were also an acceleration of the trade liberalization that started in the mid 80s and a liberalization of all price and wage controls. During the years covered in our sample, GDP grew substantially, despite the short and sharp recession during 1995, typically associated

⁸To have an idea of the average inflation rate until 1991, notice that the exchange rate when the Austral was introduced in June of 1985 was 0.8 austral per US dollars.

with the balance of payment crisis in Mexico. The exchanged rate remained fixed until January of 2002, where the exchange rate was depreciated in the midst of a banking run that started in the last quarter of 2001, a recession that started at least a year prior, and the simultaneous default of the public debt.

C.6 COMPARISON TO OTHER STUDIES

Table C.12 provides a succinct view of the samples used in other studies of price dynamics under high inflation. It shows the countries analyzed, the product coverage of each sample, and the range of inflation rates in each study. The most salient features of this paper are the broader coverage of our sample and the larger variation in inflation rates.

Table C.12: Comparison with other Studies in Countries with High Inflation

Country	Authors	Sample product coverage	Observ. per month	Sample	Inflation (%. a.r.)	Monthly freq. (%)
Argentina	This paper	506 goods/servicies, representing 84% of Argentinian consumption expenditures	81,305 on average	1988-1997	$0 - 7.2 \times 10^6$	16-99
Brazil	Barros et al. (2009)	70% of Brazilian consumption expenditures	98,194 on average	1997-2010	2-13	39-50
Israel	Llach and Tsiddon (1992)	26 food products (mostly meat and alcoholic beverages)	250	1978-1979	64	41
Israel	Llach and Tsiddon (1992), Eden (2001),	26 food products (mostly meat and alcoholic beverages)	530	1981-1982	118	61
Israel	Baharada and Eden (2004)	up to 390 narrowly-defined products from the Israeli CPI	2800	1991-1992	16	24
Mexico	Gagnon (2007)	227 product categories, representing 54.1 percent of Mexican consumption expenditures	31,500	1994-2002	3.5-45	27-45
Poland	Konieczny and Skrzypacz (2005)	52 goods, including 37 grocery items, and 3 services	up to 2400	1990-1996	18-249	59-30