Essays in Macroeconomics

Citation

Permanent link

Terms of Use

Share Your Story

Essays in Macroeconomics

A dissertation presented

by

David Rezza Baqaee

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

April 2015

Dissertation Advisors: *Author:*

Professor Emmanuel Farhi   **David Rezza Baqaee**

Professor Marc Melitz

Essays in Macroeconomics

Abstract

This dissertation focuses on three prominent areas of macroeconomic policy: fiscal stimulus, bail-outs and industrial policy, and monetary policy. In each case, I analyze the nature of the problem without intervention first before turning to why and how policy can be used to improve outcomes. In the first chapter, I study how relative demand shocks for different goods and services propagate through the economy to affect aggregate employment – and I use these insights to show how fiscal stimulus should be designed to achieve the greatest bang for buck in terms of employment. In the second chapter, I study how firm entry and exit in one industry can affect other industries and the economy as a whole through input-output relationships. I characterize which firms and industries are systemically important, show that the equilibrium is generically inefficient, and study when and how bailouts can be used to improve welfare. In the final paper, I provide a new microfoundation for downward wage rigidity, show that this microfoundation yields predictions that are consistent with the data, and study how monetary policy should behave given this microfoundation.

Contents

List of Tables

List of Figures

Acknowledgments

To my parents

Introduction

In the first chapter, I address how supply chains affect the intensity with which an industry uses labor. I derive the *network-adjusted labor intensity* as the answer to this question. The network-adjusted labor intensity measures not just the direct labor intensity of a given industry, but also takes into account the labor intensity of all its inputs, its inputs' inputs, and so on. I show that this measure is the relevant sufficient statistic determining labor's share of income, the propagation of demand shocks, the relative rankings of government employment multipliers, and the composition of optimal fiscal policy. I use network-adjusted labor shares to decompose labor's share of income into disaggregated industrial components. Using a sample of 34 countries from 1995 to 2009, I find that labor's share of income has declined primarily due to a universal decrease in labor-use by all industries, rather than changes in households' consumption demands or firms' input demands. This is in contrast to the popular value-added decomposition, which gives a much larger role to industrial composition.

In the second chapter, I show how the extensive margin of firm entry and exit can greatly amplify idiosyncratic shocks in an economy with a production network. I show that canonical input-output models, which lack the extensive margin of firm entry and exit, have some crucial limitations. In these models, the systemic importance of a firm does not respond to productivity shocks, depends only on the firm's role as a supplier, and is equal to or well-approximated by the firm's size. This means that for every canonical input-output model, there exists a non-interconnected model that has the same aggregate response to productivity shocks. I show that when we allow for entry and exit, the systemic

importance of a firm responds endogenously to productivity shocks, depends on a firm's role not just as a supplier but also as a consumer, and a firm's systemic influence is no longer well-approximated by its size. Furthermore, I show that non-divisibilities in systemically important industries can cause one failure to snowball into a large-scale avalanche of failures. In this sense, shocks can be amplified as they travel through the network, whereas in canonical input-output models they cannot.

In the third and final chapter, I show that household expectations of the inflation rate are more sensitive to inflation than to disinflation. To the extent that workers have bargaining power in wage determination, this asymmetry in their beliefs makes wages respond quickly to inflationary forces but sluggishly to deflationary ones. I microfound asymmetric household expectations using ambiguity-aversion: households, who do not know the quality of their information, overweight inflationary news since it reduces their purchasing power, and underweight deflationary news since it increases their purchasing power. I embed asymmetric beliefs into a general equilibrium model and show that, in such a model, monetary policy has asymmetric effects on employment, output, and wage inflation in ways consistent with the data. I show that although wages are downwardly rigid in this environment, optimal monetary policy need not have a bias towards using inflation to grease the wheels of the labor market.

Chapter 1

Labor Intensity in an Interconnected Economy

1.1 Introduction

How labor intensive is a production process given the existence of supply chains? With constant returns to scale, the labor intensity of producing a good, if there are no intermediate inputs, is clear: we simply divide the wage bill by total revenue. This is the gross labor share of a firm. However, when there are intermediate inputs, it is insufficient to simply look at the gross labor share since some fraction of revenues is spent on intermediate inputs. A popular measure used in the literature to account for this is labor's share of value-added. This is a firm's total wage bill divided by its value-added (revenues minus intermediate input costs).

In a neoclassical model with independent industries, the value-added labor share is a key statistic that answers many important questions. For instance, an industry's value-added labor share converts demand for goods into demand for labor. As such, it determines how demand shocks or government spending shocks to different industries can move employment. Furthermore, a weighted-average of value-added labor shares determines labor's share of aggregate income. The aggregate labor share of income is a crucial object

with implications for long-run growth, inequality, and macroeconomic dynamics.

This paper argues that in the presence of non-trivial firm-to-firm connections, the value-added labor share is the wrong measure of how much labor an industry uses. This is because the value-added labor share does not incorporate any information about the nature of an industry's supply chain. The key insight is that the labor intensity of an industry is *not* solely determined by how much of its revenues, or even its value-added, it spends on labor. To know how labor intensive an industry is, we also need to take into account the labor intensity of its entire supply chain. I derive an industry-level measure of labor intensity, the *network-adjusted labor intensity*, that takes these considerations into account. I show that the network-adjusted labor intensity is the key statistic determining labor's share of income, the employment multipliers from different kinds of government spending, and the propagation of demand shocks.

The network-adjusted labor intensity is conceptually distinct from the value-added labor share commonly seen in the literature, for example in Estrada and Valdeolivas (2012), Elsby *et al.* (2013), or Neiman and Karabarbounis (2014). Whereas the network-adjusted labor intensity counts the contributions of labor to the production of a given good throughout its supply chain, the value-added labor share divides the wage bill by revenues net of intermediate inputs. This means that the value-added measures do not take into account the labor/capital mix of the intermediate inputs of an industry.

Network-adjusted labor intensity, and the closely related network-adjusted labor share, are key statistics for understanding how industry-level changes affect aggregate outcomes. For example, the consumption-weighted average of network-adjusted labor shares is equal to labor's share of income. Writing the labor share of income as a consumption weighted average of network-adjusted labor shares allows us to decompose fluctuations in labor's share of income into more disaggregated parts. In particular, we can decompose changes in aggregate labor share into changes in consumption patterns, changes in supply chains (including trade), and changes in gross labor shares at the industry level for low, medium, and high-skill labor. This decomposition is related to the seminal work of Berman *et al.*

4

(1994), but improves upon their work by explicitly accounting for intermediate inputs and trade in intermediate inputs. The resulting decomposition is more detailed, more closely tied to theory, and more stable than the value-added decomposition of labor share common in the literature. Furthermore, by explicitly accounting for imports, it allows us to distinguish between some competing theories of why the aggregate labor share is moving.

Using this decomposition and a sample of 34 countries over 15 years, I find that, on average over the sample, the overwhelming culprit behind the decline in aggregate labor's share of aggregate income is the decline in the gross labor share of all industries, rather than compositional effects across industries. That is, it is not the case that labor intensive industries are getting smaller and capital intensive ones are getting larger – instead, all industries are using less labor. Furthermore, the decomposition casts doubt on theories of the decline of the aggregate labor share that rely directly on increased imports of intermediate and final consumption goods. These findings are in contrast to the conclusion one would reach if one relied on the popular (but misleading) value-added decomposition of the same data. The value-added decomposition attributes most of the changes to compositional effects between industries.

One of the advantages of using network-adjusted labor shares is that we can analyze changes to an industry's labor share taking into account its entire supply chain. For example, we can analyze labor's share of manufacturing income, taking into account manufacturing's reliance on non-manufacturing labor. This contrasts with the manufacturing's value added labor share, used for example in Oberfield and Raval (2012), which ignores the nature of manufacturing's supply chain. Using my approach, I show that for the United States, changes to the composition of industries is responsible for the decline in labor's share of manufacturing income. This is consistent with a story where increased imports are responsible for the drop in manufacturing's labor share. However, as stated previously, this is not a significant driver of the decline of the aggregate labor share.

Using network-adjusted labor shares by skill level, I also find that substitution of income across different types of labor, emphasized by Goldin and Katz (2009), dwarfs the

5

substitution of income between labor and capital. To the extent that factor income shares are important determinants of inequality, this suggests that substitution across different labor types has been more important than substitution from labor to capital. I show that there is a near-universal trend of industries substituting from low and medium-skilled labor towards high-skilled labor in almost all countries in the sample. Once again, the culprits are the movements of the gross labor-shares of all industries, rather than changes in the supply chains or consumption patterns of households.

Not only are network-adjusted labor intensities important for studying long-run patterns in labor's share of income, but they are also important for analyzing short-run fluctuations. Network-adjusted labor intensities determine the relative ranking of employment multipliers from demand shocks. In particular, they show how fiscal policy should be targeted to maximize its impact on output and employment. I show that in a model with involuntary unemployment, the network-adjusted labor intensity is a key determinant of the composition of optimal countercyclical fiscal policy. Furthermore, the network adjustment allows us to compute the fraction of each dollar of government spending that is eventually paid out to different kinds of workers. I find that federal government consumption expenditures (defense and nondefense) are overwhelmingly tilted towards spending on high-skilled workers with at least 4 years of college education. On the other hand, state and local government investment and private investment spend much more on low-skilled workers without college degrees. These results speak to how government fiscal policy can be targeted to stimulate specific parts of the labor market.

The structure of the paper is as follows: in section 1.2, I develop a one-factor model where the network structure of the economy is irrelevant both in terms of labor's share of income, and the relative employment multipliers from government spending and demand shocks. This shows that in a very general sense, simply having a network structure is not enough to generate interesting answers to our questions; we need a second factor. In section 1.3, I introduce the benchmark model used throughout the rest of the paper that breaks the irrelevance of section 1.2 by adding capital, and I define the network-adjusted labor intensity.

6

In section 1.4, I decompose labor's share of income into disaggregated components and show how each component has varied over time for a sample of 34 countries. In section 1.5, I characterize the relative employment multipliers from government expenditures in terms of network-adjusted labor intensities. I add a nominal rigidity and show that the network-adjusted labor intensity pins down the industrial composition of optimal fiscal policy when the zero lower bound constrains the central bank. I conclude in section 2.7.

1.2 An Irrelevance Result

In this section, I sketch a competitive constant returns to scale model where labor is the only non-constant returns to scale factor. I prove an irrelevance result in this environment showing that the network structure does not affect equilibrium employment in any meaningful way. Specifically, I show that the network does not affect labor's share of income nor the size of government multipliers. This drives home the point that models without a second factor, like Long and Plosser (1983) and Acemoglu *et al.* (2012), are uninformative about the determinants of labor's share of income or the composition of fiscal stimulus, despite having production networks. The intuition is that without profits or capital, all income is ultimately spent on labor. Therefore, the details of how industries are interconnected do not matter in terms of how shocks affect aggregate employment.

I use a dynamic framework so that the results can be directly compared to the later sections. Let the representative household maximize

$$\max_{c_{it},l} \quad \sum_{t=0}^{\infty} \rho_t U(c_{1t},\ldots,c_{nt},l_t) \tag{1.1}$$

such that

$$\sum_{i=1}^{N} p_{it}c_{it} = w_t l_t + \Pi_t + (1+i_{t-1})B_{t-1} - B_t - \tau_t,$$

where ρ_t is the discount factor in period t, p_{it} is the price of good i and c_{it} is the quantity of good i consumed in period t, the wage is w_t, labor is l_t, lump sum taxes are τ_t, nominal

government bonds are B_t with interest rate i_t, and Π_t denotes firm profits.

Assume that the representative firm in industry i maximizes profits

$$\max \quad p_{it}y_{it} - \sum_j p_{jt}x_{ijt} - w_t l_{it}$$

such that

$$y_{it} = F_i(l_{it}, x_{i1t}, \ldots, x_{int}),$$

where F_i is a constant returns to scale function and x_{ijt} are units of good j used by firm i in period t. The set of functions $\{F_i\}$ defines the network structure of this economy.

Let g_{it} be government consumption of good i, and assume that the government runs a balanced budget

$$\sum_i p_{it}g_{it} = \tau_t,$$

and sets the net supply of nominal bonds to be zero. Suppose that the distribution of government spending is given by the vector δ:

$$\delta_{it} = \frac{p_{it}g_{it}}{\sum_j p_{jt}g_{jt}}.$$

Note that in this basic setup, government consumption is socially wasteful, and is not consumed by the household.

Definition 1.2.1. A competitive equilibrium is a collection of prices $\{p_{it}\}_{i=1}^N$, sequence of wages w_t and interest rates i_t, and quantities $\{x_{ijt}, c_{it}\}_{ijt}$, and labor supplies l_t and labor demands $\{l_{it}\}_i$ such that for any given government policy $\{g_{it}\}_{i=1}^N$, and τ_t,

(i) Each firm maximizes its profits given prices,

(ii) the representative household chooses consumption and labor supply to maximize utility,

(iii) the government runs a balanced budget,

(iv) and markets for each good, labor, and bonds clear.

I focus on the steady-state equilibrium of this model, so time-subscripts are suppressed.

8

Definition 1.2.2. Labor's share of income is the wage bill wl divided by total expenditures on final goods $\sum_i p_i c_i + \sum_i p_i g_i$.

Trivially, labor's share of income, in this model, is always equal to one, regardless of the network structure. This follows from the fact that firms have constant returns to scale and make zero profits. More interestingly, we have the following result.

Theorem 1.2.1. *In the absence of profits or capital, the distribution of government expenditures has no effect on equilibrium employment. That is, equilibrium employment is not a function of $\boldsymbol{\delta}$.*

The intuition is that constant-returns-to-scale at the firm level mean that relative prices do not respond to δ. This allows us to use the Hicks-Leontief composite commodity theorem, see for example Woods (1979), to represent this economy as having only one aggregate consumption good. Therefore, the only way δ can change employment is through labor supply, or in other words, through the marginal utility of wealth. However, the marginal utility of wealth only depends on the amount the government taxes the household, not on how those taxes are spent because the household does not derive utility from government consumption. Therefore, it is only the total size of the government's budget, not its distribution, that matters.

Another way to see this is to note that constant returns to scale firms make zero profits in a competitive equilibrium. Therefore, all revenues are spent either on intermediate inputs or on labor. The portion of revenues spent on intermediaries is in turn either spent on other intermediate inputs or on labor. Ultimately, all firm revenues must be spent on labor, which means that changing the composition of government expenditures has no effect for a fixed amount of total spending. Of course, as stated above, this requires that the *composition* of government spending not affect labor supply directly.

Proof. See Appendix I.   ∎

1.3 Benchmark Model

To break the equivalence between GDP and compensation by labor, we need a second source of earnings. This could be profits or returns from land and capital. A second source of earnings is also necessary for the composition of government spending to affect equilibrium employment. Once a second "sink" for revenues is introduced in the model, the composition of government spending matters for labor demand, and labor's share of income is no longer equal to one. Then, in order to maximize employment, the government should concentrate its spending in a way that minimizes the amount of money being spent on the other factors as it travels through the supply chain. This is equivalent to the government varying the composition of its expenditures to boost labor's share of income.

In this section, the second "sink" is inelastically supplied capital rented out by households to firms in a spot market. In appendix IV, I detail how profits, rather than inelastically supplied capital, can also play the role of a second sink. Ultimately, it does not matter for the results whether the second sink is returns to capital or profits. Either way, it is a notion of the network-adjusted labor intensity that acts as the relevant sufficient statistic.

The model in this section is neoclassical. Therefore, although fiscal policy can affect equilibrium employment, interventions are socially harmful. Nevertheless, in section 1.5, I show that the intuition from the neoclassical model carries over to models with involuntary unemployment and nominal rigidities.

1.3.1 Household's problem

The household chooses

$$\max_{c_{it}, l_t, B_t} \sum_{t=0}^{\infty} \rho^t \left(\log(C_t) - \frac{l_t^\theta}{\theta} \right),$$

where

$$C_t = u(c_{1t}, \ldots, c_{nt}),$$

with u having symmetric and constant elasticity of substitution across different consumption goods. The household's budget constraint is

$$\sum (1 + \tau_{it}) p_{it} c_{it} + q_t B_t = w_t l_t + r_t K + B_{t-1} + \Pi_t - \tau_t,$$

where p_{it} is the price of good i in period t, B_t is a nominal bond (in zero net supply), Π_t is firm profits, τ_t is lump sum taxes in period t. The new ingredients in this section are r_t, the rental rate of capital, and τ_{it}, an ad valorem consumption tax on good i. So that the problem is well-defined, suppose that there is a physical limit on the number of hours that can be worked

$$l_t \leq \bar{l},$$

although we assume that this is always non-binding. This simply allows for the inclusion of inelastic labor supply as a special case. To keep the exposition clear, for now, I assume a homogenous labor market. In section 1.3.9 I consider the extension with heterogenous labor markets. Assume that capital is inelastically supplied at K.

1.3.2 Firms' problem

Firms rent capital and labor on spot markets from the household, and reoptimize every period. Therefore, their problems are static, so I suppress time-subscripts. Since in a competitive equilibrium with constant returns to scale, firm size is indeterminate, I simply state the problem of the representative firm in industry i:

$$\max_{y_i, l_i, x_{ij}} \quad p_i y_i - \sum_j p_j x_{ij} - w l_i - r k_i$$

subject to the constant-returns to scale production function

$$y_i = F_i(l_i, k_i, G_i(x_{i1}, \ldots, x_{in})),$$

where l_i is labor, k_i is capital, and x_{ij} are inputs from industry j. Let F_i and G_i have constant and symmetric elasticities of substitution between their arguments σ_F and σ_G respectively. For simplicity, I assume that $\sigma_F = \sigma_G$, thought this can be relaxed. Once again, the network

11

structure of the economy is captured by the production functions. In particular, note that if all industries used no labor, the model in this section is a special case of the model in section 1.2.

1.3.3 Government behavior

The government runs balanced budgets every period

$$\sum p_i g_i = \tau_t + \sum_i \tau_{it} p_{it} c_{it},\tag{1.2}$$

and the fraction of government expenditures on industry i is

$$\frac{\delta_i}{\sum_j \delta_j} = \frac{p_i g_i}{\sum_i p_i g_i}.$$

Before solving for the equilibrium, we need a a few key definitions that will serve us throughout the rest of the paper.

1.3.4 Network-adjusted labor intensity

Now we can define the network-adjusted labor intensity. It turns out that network-adjusted labor intensities play a key role in the determination of equilibrium employment. Recall that the representative firm in industry i has the following production function

$$y_i = F_i(l_i, k_i, G_i(x_{i1}, \ldots, x_{in})),$$

where F_i and G_i have symmetric elasticities of substitution between their arguments σ_F and σ_G respectively. Now define the following generalized elasticities of production with respect to inputs by

$$\hat{\omega}_{ij} := \frac{\mathrm{d}y_i}{\mathrm{d}x_{ij}} \left(\frac{x_{ij}}{y_i} \right)^{1/\sigma_G},$$

and

$$\alpha_i := \frac{\mathrm{d}y_i}{\mathrm{d}l_i} \left(\frac{l_i}{y_i} \right)^{1/\sigma_F}.$$

Let

$$\hat{\Omega} := \left[\hat{\omega}_{ij} \right]_{ij},$$

be the matrix of ω_{ij}'s and α be the column vector of α_i's. Define

$$\Psi := I + \hat{\Omega} + \hat{\Omega}^2 + \hat{\Omega}^3 + \cdots = (I - \hat{\Omega})^{-1},$$

to be the influence matrix. If we think of $\hat{\Omega}$ as defining a weighted directed graph, then the influence matrix is its inverse Laplacian. The ijth element of Ψ can be interpreted as the total intensity with which i uses inputs from j, taking into account both direct and indirect connections. Finally, the vector

$$\boxed{\tilde{\alpha} = \Psi \alpha}$$

is **network-adjusted labor intensity.** Intuitively, $\tilde{\alpha}_i$ captures both the direct and indirect uses of labor by industry i. Computationally, $\tilde{\alpha}_i$ is a weighted sum of the labor intensities of i, and i's suppliers, and i's suppliers' suppliers, and so on.[1]

A closely related object of interest is the network-adjusted labor share. To define this, let p_i denote the price of good i and w the wage. Then, let

$$\hat{w}_{ij} := \frac{p_j x_{ij}}{p_i y_i},$$

be industry i's expenditure share on input j, and

$$a_i := \frac{w l_i}{p_i y_i},$$

be its expenditure share on labor (gross labor share). Let

$$\hat{W} := \left[\hat{w}_{ij} \right]_{ij},$$

be the matrix of w_{ij}'s and a be the column vector of a_i's. Define

$$\boxed{\tilde{a} = (I - \hat{W})^{-1} a}$$

---

[1]We can also think of $\tilde{\alpha}$ and $1 - \tilde{\alpha}$ as the dominant eigenvectors of the matrix defined by $\begin{pmatrix} \hat{\Omega} & [\alpha \ \eta] \\ \mathbf{0} & I_2 \end{pmatrix}$, where I_2 is the 2×2 identity matrix.

to be the ***network-adjusted labor share.*** Intuitively, $\tilde{a}_i$ captures the total fraction of industry i's income that is eventually paid out to labor, whether directly by that industry itself, or through its supply chain. This fact is also noted by Valentinyi and Herrendorf (2008) who use it to measure factor income shares for the four major sectors of the US economy. The network-adjusted labor share also has an interpretation from the input-output literature pioneered by Leontief (1936). If we fix prices, and assume that production functions have a Leontief form, then the network-adjusted labor share is the total amount of labor required in order to produce a unit of a good.[2]

An important observation is that when F_i and G_i have Cobb-Douglas forms, the network-adjusted labor share and the network-adjusted labor intensity coincide. This makes Cobb-Douglas a very convenient modelling assumption, since it allows us to identify network-adjusted labor intensities from only expenditures data, and it makes the structural objects of interest $\tilde{\alpha}$ coincide with accounting objects of interest $\tilde{a}$. In Appendix III, I show how the network-adjusted labor intensities affect labor's share of income in a CES economy.

Intuitively, the network-adjusted labor intensity is always weakly greater than the gross labor intensity for every industry. This follows from the fact that taking into account the supply chain of a given industry can only increase the intensity with which labor is used.

Proposition 1.3.1. *For every industry i, we have*

$$\tilde{\alpha}_i \geq \alpha_i.$$

Proof. This follows from the non-negativity of Ω and α. ∎

Just as with the value-added labor and capital shares, the network-adjusted labor and network-adjusted capital shares always sum to one.

Lemma 1.3.2. *Let a be the gross labor shares and c be the gross capital shares of industries. Then,*

---

[2]As we shall see, when prices are being set flexibly, it is the network-adjusted labor intensity and not the network-adjusted labor share that determines equilibrium responses to shocks. So, with flexible prices and Leontief production functions, the usual input-output estimates of the impact of a demand shock are invalid.

the sum of the network-adjusted labor and capital share equals 1 *for every industry.*

$$(I - W)^{-1}a + (I - W)^{-1}c = \mathbf{1}.$$

Proof. See Appendix I. ∎

This means that the network-adjusted labor share of industry i is *labor's share of income* from industry i. We can aggregate this observation up to get labor's share of total income.

Proposition 1.3.3. *Labor's share of aggregate income is equal to the final-consumption weighted average of network-adjusted labor shares. Specifically,*

$$\frac{wl}{GDP} = \frac{(H + G)'\tilde{a}}{GDP},$$

where H is the vector of household spending net of consumption taxes, and G is the vector of government spending by industry.

Proof. See Appendix I. ∎

It is crucial to note that proposition 1.3.3 is an *accounting identity*, and it will hold for all production and utility functions. Proposition 1.3.3 will serve as the foundation for a decomposition of labor's share of income into disaggregated components in section 1.4.

1.3.5 Response to Demand Shocks

In this subsection, I show that network-adjusted labor intensities allow us to trace out the aggregate effect of shocks to an industry's demand. For clarity, I assume that production and utility functions have Cobb-Douglas forms and leave the more general case to Appendix III. Competitive equilibrium is defined in the usual way. I focus on the steady-state equilibrium of this model and therefore, suppress time subscripts.

Definition 1.3.1. The employment multiplier of a taste shock to industry i is defined as $dl/d\beta_i$, where l is equilibrium employment and β_i is the Cobb-Douglas taste of the household for goods from industry i.

The following proposition shows how the network-adjusted labor intensities allow us to translate a change in final demand for goods into changes in equilibrium employment:

Proposition 1.3.4. *Employment multipliers of taste shocks satisfy*

$$\frac{\mathrm{d}l/\mathrm{d}\beta_i}{\mathrm{d}l/\mathrm{d}\beta_j} = \frac{\tilde{\alpha}_i - \beta'\tilde{\alpha}}{\tilde{\alpha}_j - \beta'\tilde{\alpha}}.$$

Proof. See Appendix I. ∎

This shows that multipliers for industry i's demand are proportional to industry i's network-adjusted labor intensity minus the average labor intensity of consumption. This is quite intuitive since increasing the household's taste for good i will reorient household expenditures from all other industries towards industry i. To the extent that household preferences change at business cycle frequencies, this proposition shows the relative importance of preference shocks for aggregate employment. In section 1.5, I show that fiscal stimulus affects equilibrium employment in much the same way, and I derive the optimal industrial composition of fiscal policy when monetary policy is passive.

This result can be extended, without change, to the case where consumption and production functions have constant and symmetric elasticity of substitution, with β being the CES share parameters of consumption.

1.3.6 Comparison to labor share's share of value-added

The network-adjusted labor share is different from the value-added labor share commonly seen in the literature, for example in Estrada and Valdeolivas (2012), Elsby *et al.* (2013), Neiman and Karabarbounis (2014), or Oberfield and Raval (2012). Whereas the network-adjusted labor share counts the contributions of labor to the production of a given good up the supply chain, the value-added labor share divides the wage bill by revenues net of intermediate inputs.

To see the difference, consider figure 1.1. There, we see how the same production process being broken from a single aggregate firm into two firms affects the network-adjusted and value-added labor shares. In the first panel, an aggregate firm provides a good to the

(a) *Single firm serving the household.*   **(b)** *Fragmented production process.*
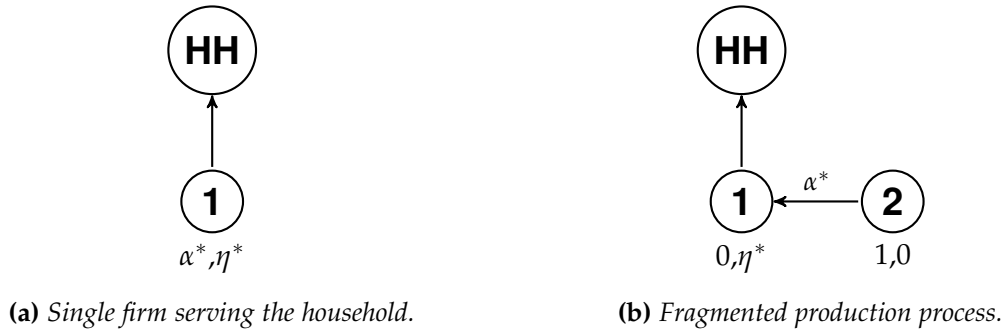
Figure 1.1: *Fragmentation of the same production process. Each node represents a firm and the numbers under each node denote the gross labor and capital share of that firm. Edges denote the flow of goods and services. The labels on the edges denote the transaction's share of the downstream firm's total expenditures.*

household. In the second panel, the first firm supplies the capital part and the second firm the labor part of production. Intuitively, the activities of firm 1 in panel (b) are just as labor intensive as firm 1 in panel (a). However, the value-added labor share and the gross labor share of firm 1 in panel (b) are both equal to zero. On the other hand, the network-adjusted labor intensity of firm 1 is the same in both cases. So the network-adjusted factor shares and factor intensities are robust to changes in accounting rules or ownership structure, while the value-added measures are not.

Theoretically, there is no reason to expect a tight connection between the value-added labor share and the network-adjusted labor share. To see this, consider an industry whose inputs are made purely from labor but does not use labor directly. This industry's value-added labor share will be zero, but its network-adjusted labor share can be arbitrarily close to one. Reversing the roles of labor and capital produces the opposite result, with value-added labor share equal to one, and a network-adjusted labor share that can be arbitrarily close to zero. In section 1.3.7, I compare observed network-adjusted, gross, and value-added measures of labor intensity for the US economy.

In practice, the difference between these two measures becomes most apparent when considering primary industries with low-margins. The value-added approach will assign high labor shares of around 90% to primary industries like "Soybean and other oilseed processing," "Fiber, yarn, and thread mills," and "poultry processing" since their capital

17

share (revenues minus labor and intermediate inputs) is close to zero. However, the network-adjusted labor share of these industries is quite low since their supply chains are not very labor-intensive.

The value-added labor share and the network-adjusted labor share will coincide when the supply chain of an industry uses the same capital/labor mix as the industry itself. The leading case of this is when an industry buys its intermediate inputs exclusively from itself (i.e. a degenerate input-output matrix with only diagonal elements). This intuition makes clear why the level of aggregation will be crucially important for whether or not the value-added labor share is a useful statistic. Once we aggregate the economy into a single sector, the input-output matrix is always diagonal (since it is a scalar), and so, the network-adjusted labor share and the value-added labor share coincide. However, at this level of aggregation, the value-added labor share is no longer informative about the industrial composition of the economy.

1.3.7 Calibration of $\tilde{\alpha}$ for the US

Now that we have defined and interpreted network-adjusted labor intensity $\tilde{\alpha}$, we turn to calibrating it for the United States. If we assume Cobb-Douglas functional forms, then we can measure $\tilde{\alpha}$ using national accounts data from the Bureau of Economic Analysis. There are two reasons to assume Cobb-Douglas. First, as discussed earlier, network adjustments are only interesting in cases when the underlying data is disaggregated. Assuming Cobb-Douglas allows us to use much more disaggregation since we only need expenditures data to calibrate the model, whereas a non-unitary elasticity requires both price and quantity data. Second, for Cobb-Douglas, all network-adjusted labor intensities can also be interpreted as network-adjusted labor shares, which are accounting objects of independent interest. Once we deviate from Cobb-Douglas, the relevance of the computations will depend on how well we choose the elasticity of substitution, which is a controversial question outside of the scope of this paper.

I use the detailed 2007 benchmark use-tables at purchaser values. This measures the
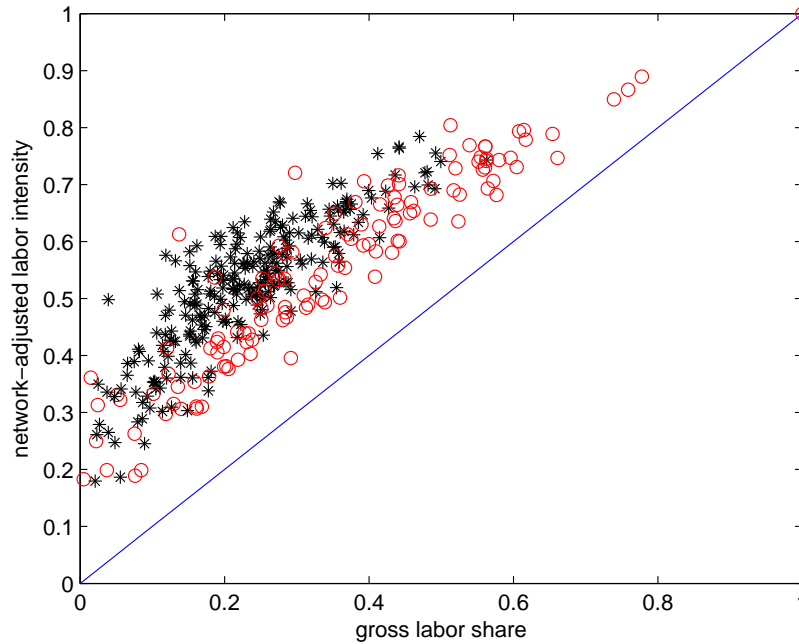
Figure 1.2: *Labor intensities plotted against network-adjusted labor intensities using the detailed BEA input-output table. There are 381 industries in this plot. The black asterisks are manufacturing industries, while red circles represent non-manufacturing industries. Non-monotonicities represent cases where industries are ranked differently according to the different measures of labor use.*

dollar expenditures of a given industry on inputs. Since the Cobb-Douglas parameters are equal to the shares of expenditures, it is easy to calibrate the production functions of the various industries, as well as the utility function of the household using this data. I let the capital intensity of a given industry equal one minus the labor share and the intermediate input share. The calculations in this section abstract from world trade in production and assume all imports are final goods. The reason for this is that data on the breakdown of imports between intermediate and final use are not available from the Bureau of Economic Analysis' statistical data sources. In section 1.4, I explain how we can account for trade in this model by using other data sources. The results are plotted in figure 1.2. As implied by proposition 1.3.1, all points in figure 1.2 lie above the 45-degree line.

Industries with the largest and smallest $|\tilde{\alpha}_i - \alpha_i|$ are listed in table 1.1. Generally speaking, manufacturing industries are much more labor intensive than their gross labor shares would

Table 1.1: *Industries with the largest and smallest differences between their network-adjusted labor intensity and their labor share.*

Industry	$\tilde{\alpha} - \alpha$	$\tilde{\alpha}$	α
Funds, trusts, and other financial vehicles	0.476	0.613	0.137
Light truck and utility vehicle manufacturing	0.459	0.498	0.039
Heavy duty truck manufacturing	0.457	0.576	0.119
Railroad rolling stock manufacturing	0.434	0.566	0.132
Motor vehicle seating and interior trim manufacturing	0.431	0.593	0.162
Automobile manufacturing	0.43	0.582	0.152
Other financial investment activities	0.423	0.721	0.298
Sawmills and wood preservation	0.422	0.609	0.187
News syndicates, libraries, archives and all other information services	0.104	0.396	0.292
Support activities for agriculture and forestry	0.105	0.682	0.577
Postal service	0.108	0.866	0.759
Civic, social, professional, and similar organizations	0.11	0.85	0.739
Office administrative services	0.111	0.889	0.778
Accounting, tax preparation, bookkeeping, and payroll services	0.111	0.635	0.524
Animal production, except cattle and poultry and eggs	0.113	0.189	0.076
Oil and gas extraction	0.114	0.199	0.085

indicate, whereas service industries like "the postal service" or "office administrative services" are about as labor intensive as their gross labor share suggests. This is intuitive, since service industries have shorter, less labor-intensive, supply chains. The key exception to this general rule are some financial industries like "Funds, trusts, and other financial vehicles," which also have much higher labor intensities than one might infer from their gross labor share. The calculations here indicate that once supply chains are properly taken into account, the manufacturing sector is very labor intensive.

The alternative popular measure of labor intensity at the industry level is the value-added labor share. As discussed earlier, there is no reason to theoretically expect the network-adjusted labor share to be related to the value-added labor share. In the data, the correlation between the value-added and network-adjusted labor shares is 0.90, which is only slightly higher than the correlation between the gross and network-adjusted labor share at 0.87. The value-added and network-adjusted labor intensities are plotted in figure 1.3. Unlike figure 1.2, there are data points above and below the 45-degree line. Furthermore,
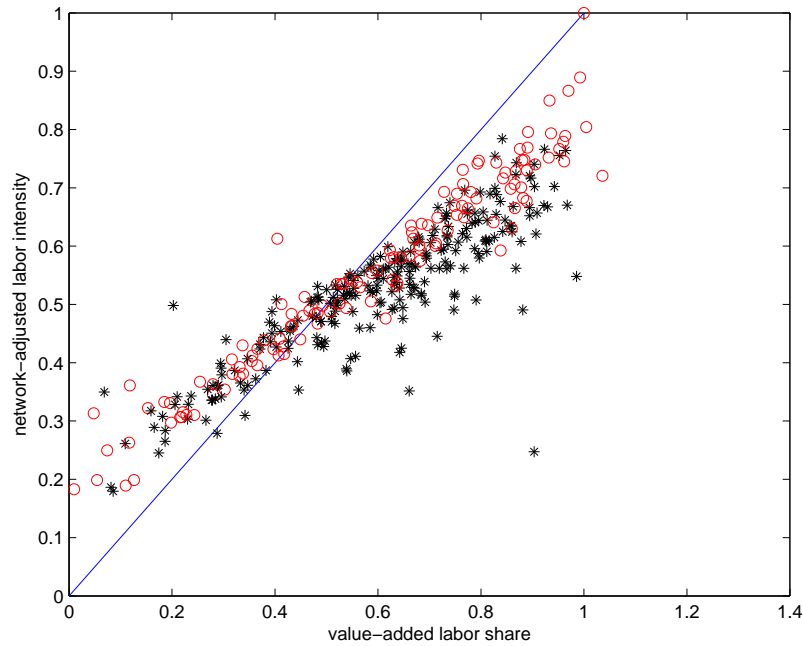
Figure 1.3: *Value-added labor shares plotted against network-adjusted labor intensities using the detailed BEA input-output table. There are 381 industries in this plot. The black asterisks are manufacturing industries, while red circles represent non-manufacturing industries. Non-monotonicities represent cases where industries are ranked differently according to the different measures of labor use.*

unlike figure 1.2, where a gross labor share close to 1 implied that the network-adjusted labor intensity must also be close to 1, no such pattern need hold now.

The fact that the slope of the line of best fit in figure 1.3 is less than 1 implies that an industry's capital-labor mix is negative correlated with its supply chain's capital-labor mix. This is because network-adjusted labor shares are higher (lower) when value-added labor shares are low (high), meaning that accounting for the supply chain properly increases (decreases) the labor share.

Another advantage of network-adjusted labor shares over value-added labor shares is that they are more stable in time series. Value-added labor shares can move around violently at a high frequency if an industry's profits fluctuate. Network-adjusted labor shares, since they are weighted averages of many industries' labor and capital shares, are more stable over time. The increased time series stability suggests that secular changes in network-adjusted

labor shares are more likely to reveal meaningful patterns. Furthermore, since they are averages over many industries, they are less badly affected by measurement error.

1.3.8 Upstream and Downstream Influence

Proposition 1.3.4 shows that $\tilde{\alpha}$ is the relevant influence measure of how industry-specific demand shocks move aggregate employment. Acemoglu *et al.* (2012) derive an alternative influence measure that they show maps supply (labor-augmenting productivity) shocks to aggregate output. To clarify $\tilde{\alpha}$'s network-theoretic properties, it helps to compare it with the alternative influence measure of Acemoglu *et al.* (2012). In this model, the Acemoglu *et al.* (2012) measure corresponds to $\tilde{\beta} \equiv \beta'\Psi$. This can be thought of as a network-adjusted consumption share. It takes into account both direct sales to households, as well as sales to industries who to sell to households, and sales to industries who sell to industries who sell to households, and so on. Acemoglu *et al.* (2012) show that $\tilde{\beta}$ is the key statistic determining how output responds to supply (labor-augmenting productivity) shocks. Since they are interested in the propagation of productivity shocks, Acemoglu *et al.* (2012) abstract away from capital and assume that labor is inelastically supplied. Since we are interested in the effect of demand shocks on employment, we need both of these ingredients, because as we saw in section 1.2, without them the model would give trivial answers to our questions.

To see the difference between the $\tilde{\alpha}$ and $\tilde{\beta}$, consider the example in figure 2.3.
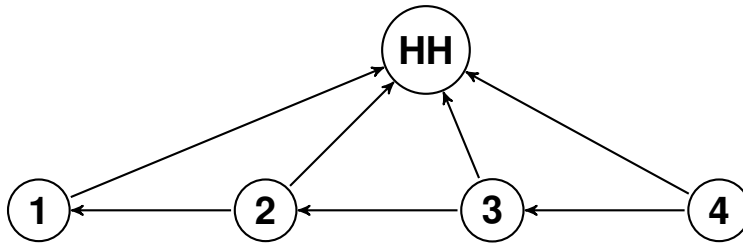
Figure 1.4: *The arrows represent the flow of goods and services.*

In figure 2.3, the network-adjusted labor intensity of firm (1) is

$$\tilde{\alpha}_1 = \alpha_1 + \omega_{12}\alpha_2 + \omega_{12}\omega_{23}\alpha_3 + \dots,$$

while the network-adjusted labor intensity of the final firm (4) is simply equal to the regular labor intensity

$$\tilde{\alpha}_4 = \alpha_4.$$

On the other hand, in figure 2.3, the network-adjusted consumption share of firm (1) is the same as its regular consumption share

$$\tilde{\beta}_1 = \beta_1,$$

while the network-adjusted consumption share of firm (4) is

$$\tilde{\beta}_4 = \beta_4 + \beta_3 \omega_{34} + \beta_2 \omega_{23} \omega_{34} + \dots.$$

This simple example makes the difference clear: the network-adjusted labor share $\tilde{\alpha}$ is a downstream centrality measure, while the network-adjusted consumption share of Acemoglu *et al.* (2012) is an upstream centrality measure. This is because demand shocks travel upstream from consumers of inputs to producers of inputs, while supply shocks travel downstream from suppliers of inputs to consumers of inputs.[3]

The difference between the two measures can be seen most clearly by setting $(\beta_i, \alpha_i) = (\alpha, \beta)$ for all i and considering two different star economies in figure 1.5.

(a) *Firm 1 is a star consumer*   **(b)** *Firm 1 is a star supplier*
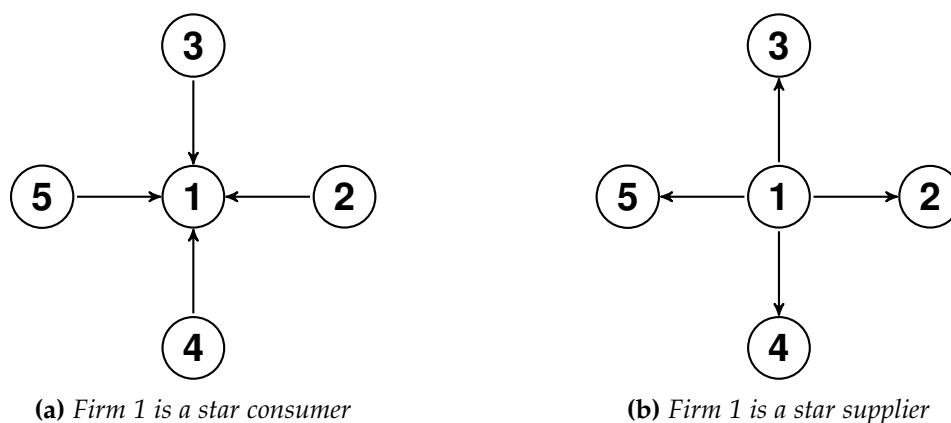
Figure 1.5: *The arrows indicate the flow of goods and services.*

---

[3]See Baqaee (2014b) for more details about the class of models where demand and supply shocks travel only in one direction.

Firm (1) in figure 1.5a will have the highest network-adjusted labor share and the lowest network-adjusted consumption share. The situation is exactly reversed in figure 1.5b. In the former case, firm (1) is an important conduit for the transmission of demand shocks and a poor conduit for the transmission of supply shocks, whereas in the latter, the opposite is true.

1.3.9 Heterogenous Labor Markets

We can easily extend the model to cover heterogenous labor markets. This allows us to analyze and craft policies to target specific parts of the labor market. To keep the notation clean, I suppress time subscripts. Suppose there are M different types of labor indexed by m and the production function of industry i is given by

$$y_i = \left(\prod_m^M l_{im}^{\iota_{im}} \right)^{\alpha_i} k_i^{\eta_i} \prod_{j=1}^N x_{ij}^{\omega_{ij}}.$$

Now, labor market clearing for type m labor is given by

$$w_m l_m = \sum_{i=1}^M p_i y_i \alpha_i \iota_{im},$$

$$= (H + G)' \Psi(\alpha \circ \iota_m),$$

where $\circ$ denotes the element-wise product and ι_m is the column vector of ι_{im}'s for different industries i. Now define the *network-adjusted type-m labor intensity* by

$$\tilde{\alpha}_m = \Psi(\alpha \circ \iota_m).$$

The network-adjusted type-m labor intensity, for a Cobb-Douglas economy, is also that labor type's share of an industry's income. In the next section, this will allow us to break up labor income into labor income by skill level. Furthermore, we can now speak of demand-side interventions to specific labor markets. So for instance, if policy-makers wish to use fiscal policy to boost low-skill employment because the low-skill labor market is failing to clear, then they can tailor policy towards increasing demand for goods with high network-adjusted low-skill labor intensity.

Unfortunately, neither the BLS nor the BEA publish statistics about the intensity with which different types of labor are used by different industries. Therefore, I use the American Community Survey from 2007 to construct ι_m for each industry group in the detailed benchmark US input-output table of 2007. I divide labor into 11 types by educational attainment. In table 1.2, I report network-adjusted and gross labor shares for each type of labor for the industries that move the most when we take the input-output structure into account. For lower skill levels, manufacturing sectors gain the most from the network adjustment, while for higher skill levels, financial industries move the most.

In the next section, I build on these results to show how and why labor's share of income has changed over the past 15 years.

Table 1.2: *Industries with the largest difference between their network-adjusted labor share and gross labor share for each labor type. This table combines data from the 2007 American Community Survey from IPUMS-USA with the detailed Benchmark Input-Output table using purchaser prices for 2007 published by the BEA.*

Education level	Industry	Network-adjusted labor share	Gross labor share
N/A or no schooling	Coffee and tea manufacturing	0.0048	0.0005
Nursery school to grade 4	Coffee and tea manufacturing	0.007	0.000
Grade 5, 6, 7, or 8	Coffee and tea manufacturing	0.023	0.001
Grade 9	Sawmills and wood preservation	0.014	0.004
Grade 10	Sawmills and wood preservation	0.015	0.004
Grade 11	Sawmills and wood preservation	0.015	0.005
Grade 12	Light truck and utility vehicle manufacturing	0.169	0.019
1 year of college	Light truck and utility vehicle manufacturing	0.067	0.006
2 years of college	Light truck and utility vehicle manufacturing	0.039	0.003
4 years of college	Funds, trusts, and other financial vehicles	0.219	0.057
5+ years of college	Funds, trusts, and other financial vehicles	0.156	0.030

Table 1.3: *Industries with the largest network-adjusted labor share for each labor type. This table combines data from the 2007 American Community Survey from IPUMS-USA with the detailed Benchmark Input-Output table using purchaser prices for 2007 published by the BEA.*

Education level	Industry	Network-adjusted labor share
N/A or no schooling	Support activities for agriculture and forestry	0.011953
Nursery school to grade 4,	Greenhouse, nursery, and floriculture production	0.018467
Grade 5, 6, 7, 8	Support activities for agriculture and forestry	0.057632
Grade 9	Civic, social, professional, and similar organizations	0.024194
Grade 10	Civic, social, professional, and similar organizations	0.026585
Grade 11	Electronic and precision equipment repair and maintenance	0.029375
Grade 12	Private households	0.407733
1 year of college	Private households	0.173831
2 years of college	Residential mental retardation, mental health, substance abuse and other facilities	0.115352
4 years of college	Management of companies and enterprises	0.360452
5+ years of college	Accounting, tax preparation, bookkeeping, and payroll services	0.385613

1.4 Labor's share of income

As lemma 1.3.2 and proposition 1.3.3 show, network-adjusted labor shares are labor's share of an industry's income. Therefore, they can help us to decompose aggregate labor shares into disaggregated industrial components. Research on the evolution of labor's share of income has recently exploded. Piketty (2014) places labor's share of income at the heart of his theory of inequality. A variety of papers have been written on the causes and consequences of the decline in labor's share of income (see Neiman and Karabarbounis, 2014; Oberfield and Raval, 2012; Elsby *et al.*, 2013). Popular theories for why labor's share of income has trended down include: increased globalization, increases in the capital stock, decreases in the price of investment goods, and increased automation in production.

This paper's contribution to this debate is to provide a coherent accounting framework for decomposing labor's share of income into disaggregated components. Using this framework, I find that the decline in labor's share of income is due primarily to a decrease in the gross labor share of all industries, and not changes to the composition of industries. In particular, I do not find strong evidence for the idea that labor's share of aggregate income has decreased due to substitution of imported inputs or imported consumption goods for domestic labor. I also find similar results for the income share of different labor types by education. Furthermore, consistent with the skill-biased technical change hypothesis of Goldin and Katz (2009), I find that changes in income share within labor types dwarfs changes between labor and capital's share of income, and that the source of these changes is within industries.

As already discussed, the network-adjusted labor share of an industry is precisely labor's share of that industry's income, and the GDP-weighted average of network-adjusted labor shares is equal to labor's share of aggregate income. This accounting identity must hold regardless of the underlying production and utility functions. This allows for a decomposition of labor's share of income into disaggregate industrial components.

For this section, I add international trade in intermediate and final goods to the model in section 1.3. This not only brings the model closer to the data I use, but it also allows

us to account for the effects of globalization on labor's share of aggregate income. The key assumption of the model is that labor and capital are immobile, but other goods and services are traded with the rest of the world in consumption and production. The results in this section complement the work of Trefler and Zhu (2010) who account for intermediate inputs in computing the factor content of trade. They show that adjusting for the role of intermediate inputs significantly improves the Heckscher-Ohlin model's fit to the data.

The key result for this section, proposition 1.4.1, does not depend on structural assumptions, and relies only on accounting identities. Let W^* be the matrix whose ijth element is industry i's share of expenditures on the domestic industry j, and let b be the column vector whose ith element is the share of final-use expenditures on domestic industry i. Finally, let a be the column vector whose ith element is the gross labor share of industry i. Then, analogous to the closed-economy proposition 1.3.3, the following proposition holds with international trade.

Proposition 1.4.1. *Labor's share of income is*

$$\frac{wl}{GDP} = b'(I - W^*)^{-1}a = b'Pa = b'\tilde{a}, \tag{1.3}$$

where $P = (I - W^*)^{-1}$, and $\tilde{a}$ is the vector of network-adjusted domestic labor shares.

Proof. Let s_i denote the sales of industry i. Then, labor market clearing implies that

$$wl = s'a.$$

Furthermore, market clearing for good i gives

$$s' = (b)'\text{GDP} + s'W^*,$$

where W^* is the domestic input-output expenditure share matrix. Then,

$$s' = (b)'(I - W^*)^{-1}\text{GDP},$$

and so

$$wl = (b)'(I - W^*)^{-1}a\text{GDP}.$$

27

■

Note that this proposition holds for *any* structural model of the economy because it only makes use of accounting identities. A convenient way to give proposition 1.4.1 a structural interpretation is presented in Appendix II. There, I show that in an Armington model of trade, with unitary elasticity of substitution in consumption and production, the network-adjusted domestic labor share coincides with the network-adjusted domestic labor intensity. Adding trade to the model in this way does not change the model's qualitative properties. Analogues of all of the propositions in section 1.3 exist in the model with international trade. The presence of traded goods simply means that we must adjust the influence matrix for the fact that some fraction of expenditures on each good purchased was imported.

Using proposition 1.4.1, decompose labor's share of income into changes in its constituent parts

$$\Delta \frac{w_t l_t}{GDP_t} = (\Delta b)'_t P_t a_t + b'_{t-1}(\Delta P_t)\alpha_t + b'_{t-1}P_{t-1}(\Delta a_t), \tag{1.4}$$

where Δ is the time difference operator. Observe that if we assume Cobb-Douglas functional forms, b, P, and a will correspond Cobb-Douglas parameters. Summing equation (1.4) over N time periods gives

$$\frac{w^i_{t+N} l^i_{t+N}}{GDP_{t+N}} - \frac{w_t l_t}{GDP_t} = \underbrace{\sum_t^{t+N} \Delta b'_t P_t a_t}_{\text{consumption}} + \underbrace{\sum_t^{t+N} b'_{t-1}\Delta P_t a_t}_{\text{supply chain}} + \underbrace{\sum_t^{t+N} b'_{t-1}P_{t-1}\Delta a_t}_{\text{gross labor share}}. \tag{1.5}$$

Equation (1.5) decomposes changes to labor's share of income over time period t through $t + N$ into changes due to three different components: (1) changes to the composition of final goods consumption, (2) changes to the supply chain, including increased use of imported inputs, and (3) changes to the fraction of expenditures on labor by each industry.

We can see that the first summand is the effect of changes in consumption because Δb_t is the change in what final goods are demanded in the economy. To turn the effect of a change in final good demand into a change in aggregate labor share, we must multiply Δb_t by the network-adjusted labor shares $P_t a_t$ to capture the flow-on effects of changes in final good

28

demand on intermediate industries.

We can see that the second summand is the effect of changes in supply chains because ΔP_t captures the changes in the input-output matrix. To turn the effect of a change in the domestic input-output matrix into a change in aggregate labor share, we must multiply it by final goods demand b_{t-1} and gross labor shares a_t.

Finally, we can see that the third summand is the effect of changes in gross labor shares because Δa_t captures changes in industry-level gross labor shares. To turn the effect of a change in gross labor shares into a change in aggregate labor share, we must multiply Δa_t by the total size of the industries $b'_{t-1}P_{t-1}$.

Of course, in practice, all of these terms will be moving together at the same time. However, changes are still interpretable. As an example, consider the case where industry i reduces its expenditures on labor, so a_{it} falls. This means that either that industry's gross capital share must be increasing or its intermediate input share must be increasing. If only the gross capital share increases, then we would observe a drop in the third component of the summand and no change in the second and first component, since $\Delta b = \Delta P_t = 0$. However, if the intermediate input share rises instead, then it depends on what intermediate input is being purchased. If that intermediate input uses a lot of labor, then we observe a drop in the third component and an increase in the second component. If that intermediate input uses a lot of capital or was imported, then we observe a drop in the third component and no change in the first and second components.

To summarize, the first component of (1.5) captures changes in how final goods consumption has changed across industries. This would capture changes in labor's share of income due to changes in household consumption patterns (either across different industries or between domestic/foreign production). The second component of (1.5) captures how changing supply chains are affecting labor's share of income. This would include either changes in the interconnections between industries, or increased use of imported intermediate inputs. The first two components capture changes in labor's share of income due to the changing composition of industries. This means that globalization-driven changes to

the labor share, most recently emphasized by Elsby *et al.* (2013), should show up in the first two components. This is because if labor's share of income is falling due to households and firms buying more labor-intensive goods from overseas, this should show up in the first or second component of (1.5).

A further breakdown is possible if we assume that labor inputs consist of high skill, medium skill, and low skill labor. Then we can further decompose the changes in the high, medium, and low skill labor share as

$$\Delta \frac{w_t^i l_t^i}{GDP_t} = \underbrace{(\Delta b_t') P_t (a \circ \iota_t^i)}_{\text{consumption}} + \underbrace{b_{t-1}' (\Delta P_t)(a_t \circ \iota_t^i)}_{\text{supply chain}} + \underbrace{b_{t-1}' P_{t-1} \Delta (a_t \circ \iota_t^i)}_{\text{gross labor share}} + \underbrace{b_{t-1}' P_{t-1} (a_{t-1} \circ \Delta \iota_t^i)}_{\text{intralabor}},$$

(1.6)

where ι_t^i is the vector of the shares of type i labor as a fraction of total labor used by the different industries in period t, and $\circ$ is the element-wise product. This formula allows us to decompose changes in labor type i's share of income into four components. The first three are the same as before, but now we have a fourth term capturing substitution within labor. The primary reason to suspect that the fourth term has changed is skill-biased technical change, emphasized by Goldin and Katz (2009).

For this section, I use data from the World Input-Ouput Database (WIOD). Using the WIOD, we can compute labor's share of income, and the decomposition of labor's share of income, implied by input-output tables of 34 different countries from 1995 to 2009. One of the great advantages of the WIOD over national input-output tables is that the WIOD includes data on trade in intermediate inputs. Whereas, many national data sources, like the BEA, do not provide this information. The downside to using the WIOD is that rather than having 381 industries, there are only 35 industries. For more information on the sources and construction of the WIOD see Timmer *et al.* (2012).

In figure 1.6, I plot the cross-country average (weighted by GDP) of equation (1.5) for the entire sample. This can be interpreted as a decomposition of labor's share of average income. We can see that the labor share at the industry level explains the majority of changes in labor's share of income. Changing consumption patterns also contribute, but

their contribution is more than 3 times smaller. Changing supply chains, on average, are not causing any trends in labor's share of income.
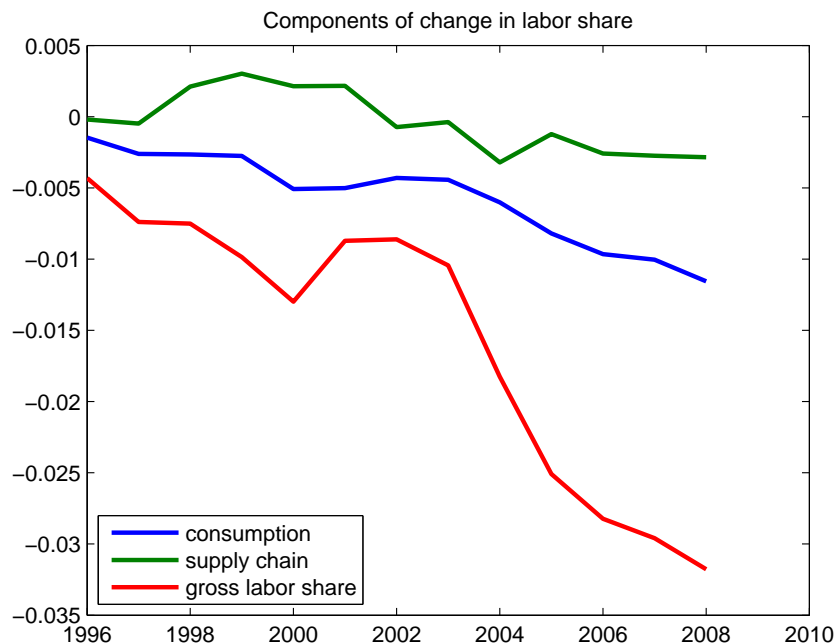
Figure 1.6: *Cross-country average of cumulative changes in labor's share of income according to decomposition in equation (1.5). The data are from the World Input-Output Database. The data is in percentages.*

The fact that the "consumption" and "supply chain" lines do not move very much in figure 1.6 is evidence against the idea that changes in the nature of supply chains or changes in imports of foreign goods for domestic goods have caused aggregate labor's share of income to drop in the sample (on average). However, by averaging over many countries, we are losing interesting variation within countries. If it turns out that in some countries globalization is increasing labor's share of income and in others it is decreasing labor's share, then we lose this by averaging. To get a sense of magnitudes, in figure 1.7 I show the total change, in absolute values, of each component of equation (1.5) over the sample for each country. The figure shows that by and large the largest movements in the labor's share of income are in the gross labor shares of all industries.

In figure 1.8, I break down the cross-country average into its effect for different labor
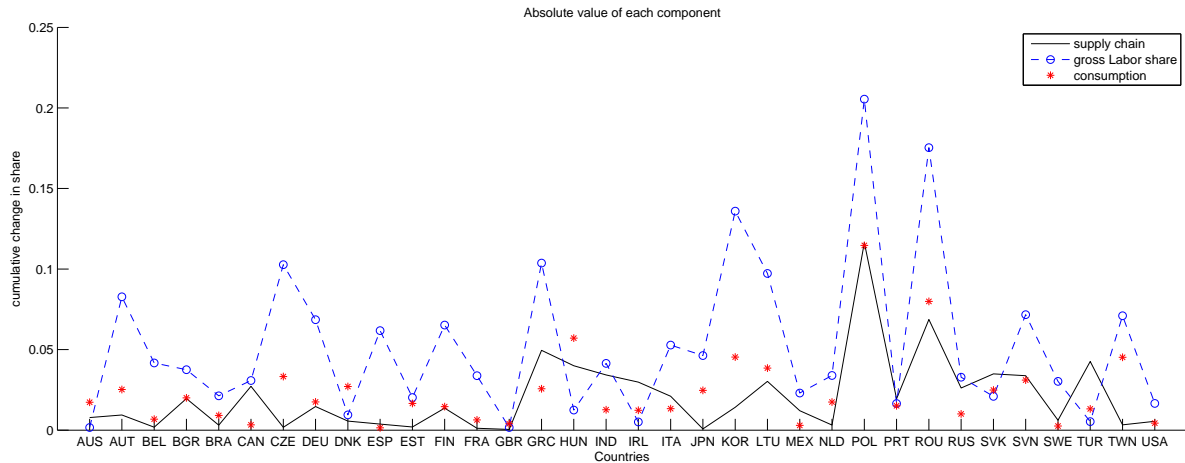
Figure 1.7: *The absolute value of the total change in each component of equation (1.5) for each country from 1995 to 2009. The data are from the World Input-Output Database*

types by equation (1.6). We see that the largest changes are attributable to the change in the composition of labor use from low-skill to high-skill. Adding the three lines in each plot gives the evolution of that labor-type's share of income. While low-skill and medium-skill labor shares have declined since 1995, high-skill labor's share of income has increased, primarily due to increased reliance on high-skilled labor relative to other types of labor.

In terms of substitution within labor's share of income, that is substitution between differently skilled labor, the global picture is much more homogenous. To show this, in table A.1, I report the total change owing to each of these components for all the countries in the sample for which all the data is available. The numbers in table A.1 show that, within labor-types, there are very strong and near-universal compositional effects, with high-skill labor's share increasing and low-skill labor's share dropping. Column 5 shows that all countries in the sample, with the exception of Denmark and Estonia, feature declining low-skill labor share. Column 4 shows that large fractions of this decline are due to changes in the fraction of low-skill labor as a fraction of total labor. Similarly, all countries except Mexico feature increasing high-skill labor share of income. Therefore, the change between labor types is happening within industries and not across them. Adding the 5th, 10th, and 15th columns gives the overall change in the labor share.
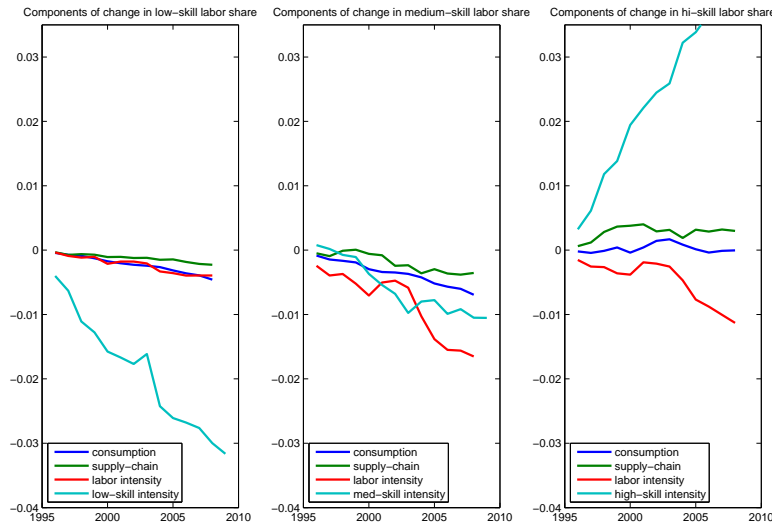
Figure 1.8: *Cross-country average of changes in components of labor's share of income according to decomposition in equation (1.6). The data are from the World Input-Output Database.*

These findings update and strengthen the results ofBerman *et al.* (1994) who found that realloaction from unskilled to skilled labor in manufacturing over the 1980s in the US occurred within industries rather than between them. My work improves upon their decomposition by explicitly accounting for intermediate inputs and trade. My findings are also strongly supportive of the "skill-biased technical change" thesis of Goldin and Katz (2009) and Katz and Murphy (1992). The pattern of intra-labor substitution is particularly pronounced for the United States (see figure A.2), where the labor share has remained roughly constant from 1995 to 2009, but the relative shares of the different skill levels have changed drastically. This suggests that, at least for the United States, increases in income inequality are more likely linked to substitution from low-skill and medium-skill labor to high-skill labor, rather than increased use of imports or capital.

Figure A.2, in the Appendix, plots the decomposition of the labor share into its constituent high-skill, medium-skill, and low-skill components, as well as a decomposition of the changes in the each type of labor share for the US. Once again, the break-down shows that the largest component of the decline in low-skill labor intensity is the final component:

33

low-skilled labor as a share of total labor. Although changing consumption patterns and supply chains contribute to fluctuations, there are no strong universal trends. This rules out theories of the decline in low and medium skilled labor's share of income that rely on substitution from labor towards imported intermediate inputs, or changing composition of industries. It also rules out the possibility that changing supply lines, for example an increase of IT services in production, is driving the trend. If these factors were driving the trends, we should expect the "consumption" and "supply-chain" lines to be trending downwards. The trends we observe are consistent with skill-biased technical change.

1.4.1 Comparison to value-added decomposition

An alternative decomposition of labor's share of income common in the literature follows from the following accounting identity:

$$\frac{wl}{GDP} = \sum_i \frac{VA_i}{GDP} \hat{\alpha}_i,$$

where $VA_i = p_i y_i - \sum_j p_j x_{ij}$ is sales net of intermediate input costs, and $\hat{\alpha}_i$ is the value-added labor share defined as labor costs divided by value-added. This identity gives rise to

$$\Delta \frac{w_t l_t}{GDP_t} = \underbrace{\Delta \frac{VA}{GDP}_t{}' \hat{\alpha}_t}_{\text{between-industries}} + \underbrace{\frac{VA}{GDP}_{t-1}{}' \Delta \hat{\alpha}_{t-1}}_{\text{within-industries}}.$$

This decomposition, if the input-output matrix were diagonal, would have the following interpretation. The first summand captures changes to the composition of industries due to changing final-use expenditure patterns. For instance, households are spending a larger fraction on foreign-made goods or they are spending a larger fraction of their income on less-labor intensive sectors. The second summand, on the other hand, captures the labor/capital mix of each industry holding fixed sectoral compositions. These terms are commonly, but misleadingly, referred to as the between-industries and within-industries changes.

Since the input-output matrix is non-diagonal, these interpretations are not technically appropriate. In the empirically relevant case where the input-output matrix is non-diagonal,

the value-added decomposition is difficult to interpret. This is because it is a non-linear transformation of the data that is not tightly connected with the theory. In figure 1.9, I plot the cross-country decomposition using value-added measures. Figure 1.9 is the value-added decomposition of the same data as figure 1.6. The two figures tell drastically different stories. The value-added decomposition gives the misleading impression that the composition term is much more important than within industry changes in explaining the decline in labor's share of aggregate income. This would incorrectly suggest that theories that affect industrial composition, could be the significant driver of the effect for the aggregate labor's share of income in this sample.

To see why the two decompositions may paint differing pictures, consider a simple example of an industry that uses labor, capital, and imported intermediate inputs to produce. Now suppose that this industry experiences capital-biased technological change, so that it changes its input mix, but it does not expand or shrink its sales. This industry substitutes away from intermediate inputs and labor towards using more capital. This would seem like a textbook case of within-industry change and under the network-adjusted decomposition, it would show up as purely a within-industry change. However, under the value-added decomposition, this would show up as both a change in the composition of industries (since the industry's value added would go up) and within-industries (since the industry's labor's share of value added would go down).
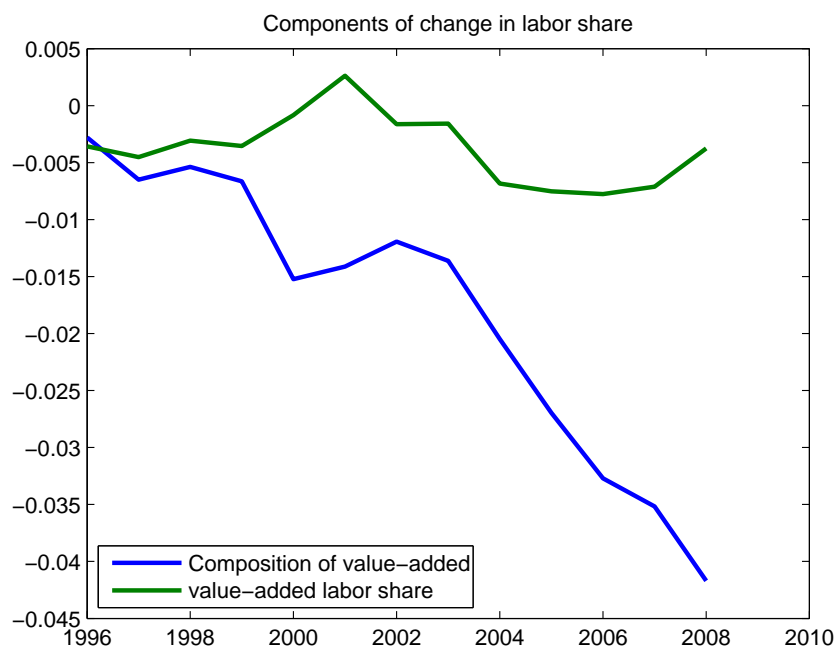
Figure 1.9: *Cross-country average of changes in labor's share of income using the value-added decomposition. The data are from the World Input-Output Database.*

1.4.2 Labor's Share of Manufacturing Income

A great advantage of this disaggregated approach is that we can zoom in on individual industries in a well-defined sense (without discarding the changes in their supply chains) and see which component is driving the change in their network-adjusted labor intensities for various labor types. In this subsection, manufacturing provides a good case-study.

The decline in manufacturing's labor share in the US has attracted much attention, for example, it forms part of the story behind the decline in labor's share of income in Elsby *et al.* (2013), and is the focus of Oberfield and Raval (2012). The papers in this literature focus on how manufacturing's value-added labor share has evolved over time. In figure 1.10, I plot manufacturing sector's labor share, as measured by compensation of employees as a fraction of revenues ("gross labor share"), compensation of employees as a fraction of value added ("value-added labor share"), and the final consumption weighted network-adjusted labor share of the manufacturing sector for the US from 1995 to 2009. The network-adjusted

labor share should be interpreted as the fraction of each dollar spent in manufacturing that is eventually paid to workers (even if they are non-manufacturing workers).

Formally, the value-added measure is defined as

$$\sum_{i \in M} \left(\frac{VA_i}{\sum_{j \in M} VA_j} \right) \left(\frac{\alpha_i}{\alpha_i + \eta_i} \right),$$

where M is the set of manufacturing industries, α_i is gross labor share, and η_i is gross capital share. Note that the value-added measure ignores how the supply chains of manufacturing are changing. The Network-adjusted labor share is given by

$$\sum_{i \in M} \left(\frac{\beta_i^*}{\sum_{j \in M} \beta_j^*} \right) \tilde{\alpha}_i,$$

where β^* is final-use expenditure shares. The supply chain of each industry is encapsulated in $\tilde{\alpha}_i$. Economically, the network-adjusted labor share of manufacturing captures the fraction of each dollar of expenditures in manufacturing eventually spent on labor (either directly or indirectly through intermediate inputs). This is precisely labor's share of manufacturing income.

From 1995-2009, the manufacturing sector's value-added labor share fell by 12.2%. The network-adjusted labor intensity, however, dropped by 6.5%. The gross labor share fell by only 0.9%. The network-adjusted and value-added measures are highly correlated (correlation of 85%) but they're far from identical, either in levels or in changes. For instance, the drop in the network-adjusted measure is almost half as large as the one in the value-added share.

Crucially, we can go one step further and decompose the share in figure 1.10 according to equation (1.5). The results are plotted in figure 1.11. We see that in manufacturing, globalization and changing industrial composition have played a much larger role in labor's share of income than for the US economy as a whole. Unlike the the aggregate labor share, labor's share of US manufacturing income has been significantly affected by changing supply chains and consumption patterns. This finding is consistent with the idea that the changing composition of industries, which includes increased import competition, are
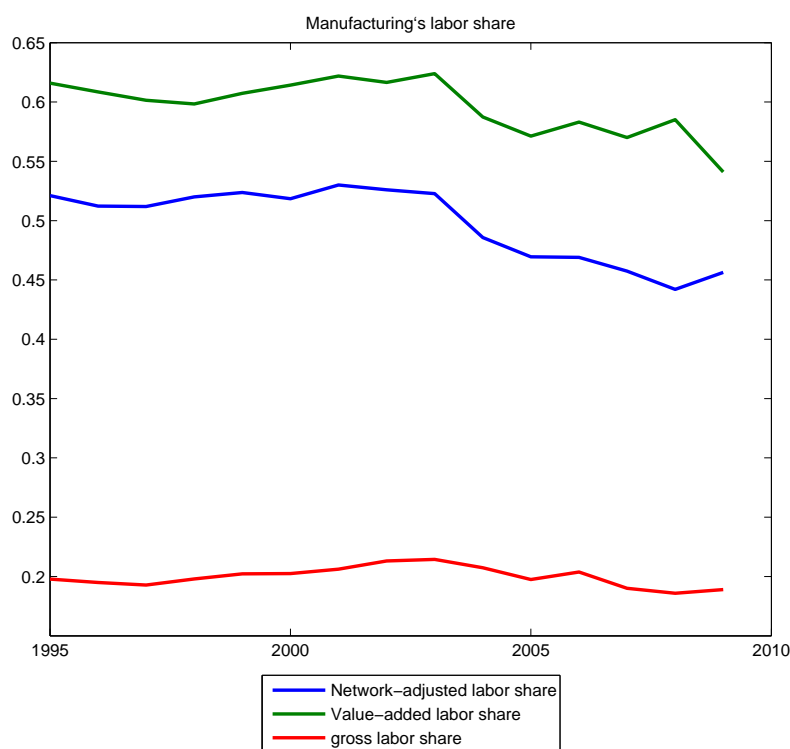
Figure 1.10: *Evolution of labor use by the manufacturing industries of the US using the WIOD data from 1995-2009.*

responsible for the decline in labor's share of manufacturing income in the US, even if this does not aggregate up to be important for the economy as a whole. The idea that import competition has been important to workers involved in manufacturing is consistent with recent findings of Acemoglu *et al.* (2013).

Figure A.3 in Appendix VI shows that the value-added decomposition is not misleading for US manfacturing. In this case, the composition effect *is* picking up the trend in aggregate labor shares in a way that's consistent with the results of figure 1.11. This suggests that once we aggregate over all manufacturing industries, assuming a block-diagonal input-output matrix, where manufacturing industries only use inputs from other manufacturing industries is not a bad assumption.

We can further decompose the network-adjusted labor intensity of manufacturing into network-adjusted high-skill, medium-skill, and low-skill intensity. These are plotted in
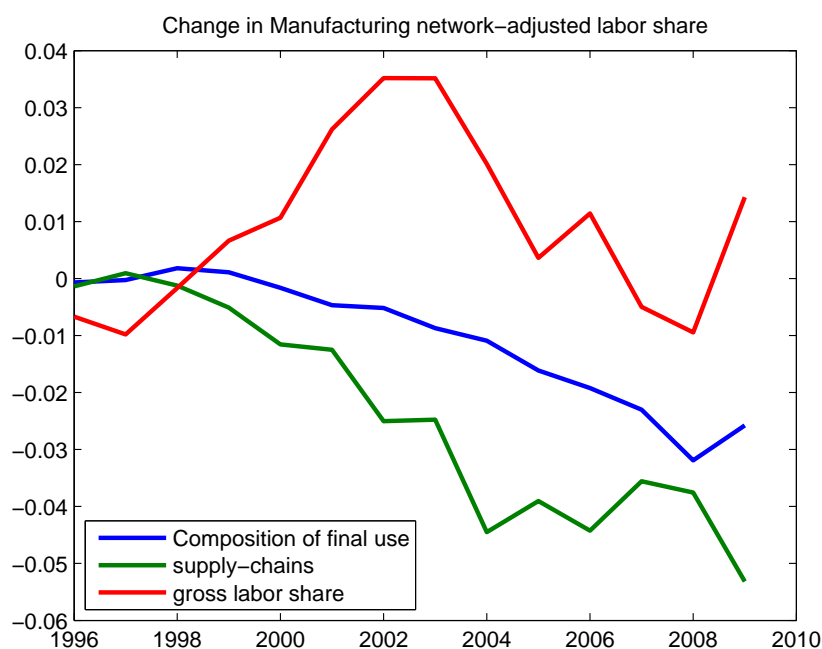
Figure 1.11: *Decomposition of network-adjusted labor intensity of the manufacturing industries of the US using WIOD data from 1995-2009.*

figure 1.12. Here, we see the same pattern as in the rest of the data: high-skill use has trended upwards as medium and low-skill use has trended down. We see that the compositional changes, which include increased import competition, have had their biggest impact on medium-skill labor use of manufacturing.
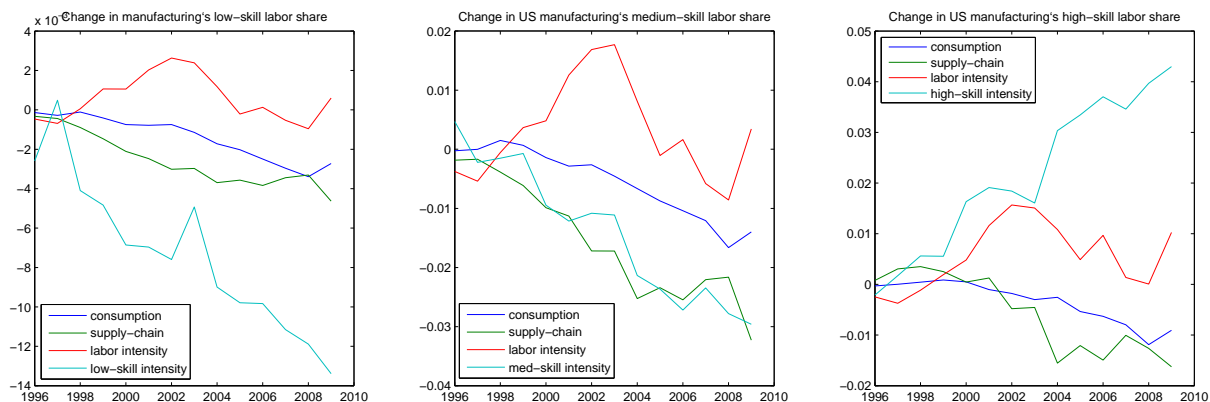
Figure 1.12: *Network-adjusted labor intensity of the manufacturing industries of the US using WIOD data from 1995-2009.*

1.5 Countercyclical Fiscal Policy

Now that we understand the accounting applications of network-adjusted labor intensities, let us consider some of their policy implications. Received wisdom from Keynesian macroeconomics is that governments can use fiscal policy to stimulate employment at the zero-lower bound (for instance, see Christiano *et al.*, 2011; Farhi and Werning, 2012). The question of exactly how they should do this is often left unexplored.

In section 1.3, I showed that the network-adjusted labor intensity tells us how more household demand for an industry's output eventually ends up as more demand for labor. In the neoclassical model of section 1.3, this information did not help us ask any normative policy questions because the equilibrium was efficient. However, we may think that in reality there are times when the government may want to pursue policies to raise employment. One such case, much studied in the literature, is in the context of a New Keynesian model at the zero-lower bound. In such a scenario, labor may be idle and government policy that expands employment can be welfare improving.

In this section, I show, in the context of a model with a production network, the network-adjusted labor intensity determines how much employment expands with increased government spending. When I introduce a nominal friction that causes involuntary

unemployment, optimal government policy will be to target those industries with the highest network-adjusted labor intensities. I will begin by deriving theoretical results linking the employment multiplier to the network-adjusted labor intensity, and then calculate how different types of government spending are expected to effect aggregate employment.

1.5.1   Neoclassical Benchmark

Before introducing any frictions, let us first see how the benchmark neoclassical model of section 1.3 responds to changes in fiscal policy. Since I focus on perturbations to the steady state of this model, changes in government policy are permanent changes to the steady state of the model. This implies that government spending has very strong crowding-out effects, since household's permanent income adjusts one-for-one with government expenditures.

Definition 1.5.1. The relative employment multiplier of government spending in industry i is defined as $\mathrm{d}l/\mathrm{d}\delta_i$, where l is equilibrium employment and δ_i is the share of government expenditures in industry i, holding fixed the total size of the governments' budget.

Proposition 1.5.1. *Government employment multipliers satisfy*

$$\frac{\mathrm{d}l/\mathrm{d}\delta_i}{\mathrm{d}l/\mathrm{d}\delta_j} = \frac{\tilde{\alpha}_i - \delta'\tilde{\alpha}}{\tilde{\alpha}_j - \delta'\tilde{\alpha}}. \tag{1.7}$$

Proof. See Appendix I. ∎

Proposition 1.5.1 shows that the relative multipliers from government spending are pinned down by the network-adjusted labor intensities. So, the government can boost employment by redirecting spending towards sectors with higher network-adjusted labor intensities. In other words, network-adjusted labor intensities also allow us to map how changes in final demand by the government translate into changes in equilibrium employment. Similar results hold for consumption taxes:

Proposition 1.5.2. *The employment multiplier from consumption taxes satisfy*

$$\frac{\mathrm{d}l/\mathrm{d}\tau_i}{\mathrm{d}l/\mathrm{d}\tau_j} = \frac{\beta_i\tilde{\alpha}_i}{\beta_j\tilde{\alpha}_j}\frac{(1+\tau_j)^2}{(1+\tau_i)^2}.$$

Proof. See Appendix I.   ∎

   The employment response of a consumption tax is determined by how intensively an industry uses labor $\tilde{\alpha}$ and how intensively households consume that good β. This is because if households consume very little of a good, a consumption tax on that good will have small effects on employment even if that good uses labor intensively. This proposition demonstrates how targeted consumption taxes or subsidies, like "Cash for Clunkers," affect equilibrium employment. Proposition 1.5.2 shows that the efficacy of such government programs depends on the size of the subsidies, the household's tastes for the good being subsidized, and the network-adjusted labor intensity of the final industry producing the good.

   Of course, in the benchmark model, there is no reason for the government to manipulate employment. The first welfare theorem holds and the optimal level of government taxation and expenditures is zero. However, once we allow for involuntary unemployment, these positive predictions become prescriptive. In the next section, I consider a second-best world where the zero-lower bound constrains the central bank, and neither the fiscal or monetary authority can commit to taking actions in the future. Furthermore, the only tools available to the fiscal authority are direct purchases by the government.

1.5.2   Keynesian Model Setup

Following the distinction made by Werning (2011), countercyclical fiscal policy can be optimal for opportunistic or stimulus reasons. Intuitively, opportunistic fiscal policy occurs when the government can provide useful goods and services to the household, and a recession is a particularly cheap time for the government to provide these goods and services. On the other hand, stimulus fiscal policy occurs when government expenditures are not directly valuable but still raise utility through a Keynesian multiplier effect. This is because government expenditures stimulate households to spend more and this raises private consumption. I add a few ingredients to the benchmark model in section 1.3 to study the model's normative properties for both types of fiscal policy.

To allow for opportunistic stimulus, I allow government purchases to enter the household's utility function directly. This gives fiscal policy a motive to increase expenditures during recessions. Second, I allow for heterogeneity in household types: specifically, there are some credit-constrained households that violate Ricardian equivalence, and consume a constant fraction of their contemporaneous income. The presence of these households allows fiscal policy to have pure "stimulus" effects. Last, I make wages downwardly rigid so that the equilibrium after a shock is not necessarily efficient.

Households

As in Eggertsson and Krugman (2012), suppose there are two representative households with differing discount factors. The more patient household is called the saver and the less patient one the borrower. The fraction of savers in the population is $1 - \chi$ and the fraction of borrowers is χ.

Let the saver maximize

$$\sum_t \rho^t \left[(1 - \lambda) \log(c_t^s) + \lambda \log(G_t) \right], \quad \lambda \in (0, 1)$$

where

$$c_t^s = \prod_k \left(c_{t,k}^s \right)^{\beta_k},$$

is private consumption by the saver, and

$$G_t = \prod_i g_{it}^{\phi_i}.$$

is government consumption services. We maintain the assumption that $\sum_k \beta_k = 1$. Since the government consumption good is additively separable from the household's private consumption, the government's consumption behavior does not directly distort the household's consumption choices through the utility function. The saver has budget constraint

$$\sum_k p_{t,k} c_{t,k}^s + B_t + D_t = (w_t l_t + r_t K_t)(1 - \chi) + (1 + i_{t-1})[D_{t-1} + B_{t-1}] - \tau_t^s,$$

where $p_{t,k}$ is the price of good k in time t. Nominal government bonds are B_t and debts of

other households are D_t. The nominal net interest rate on debt is i_t. The household receives labor income $w_t l_t$ and capital income $r_t K_t$ in proportion to its share of the population. Households are endowed with an exogenous amount of labor and capital. Finally, savers face lump sum taxes τ_t^s.

The borrower, who has a smaller discount factor, faces the same problem as the saver but is subject to a borrowing limit on its debt:

$$D_t \leq \frac{D^h}{1+i_t} \frac{p_{t+1}}{p_t},$$

where p_t is the ideal price index for the households in period t.

Firms

The firms behave exactly as before. That is, the firms are competitive and rent capital and labor on spot markets from the household and reoptimize every period. Therefore, their problems are static.

$$\max_{y_{it}, l_{it}, x_{ijt}} \quad p_{it} y_{it} - \sum_j p_{jt} x_{ijt} - w_t l_{it} - r_t k_{it},$$

such that

$$y_{it} = (l_{it})^{\alpha_i} k_{it}^{\eta_i} \prod x_{ijt}^{(1-\alpha_i-\eta_i)\omega_{ij}}.$$

The government

The government faces the budget constraint

$$B_t = (1+i_{t-1})B_{t-1} + \sum_k p_{t,k} G_{t,k} - \tau_t,$$

where τ_t is income from lump sum taxation. The government cannot target its tax base, so that taxes levied on borrowers and savers are proportional to their share of the population. Furthermore, the government cannot use consumption taxes, since, as shown by Correia *et al.* (2013), a government with access to a rich-enough set of taxes could replicate negative interest rates and achieve the first-best outcome. Unlike the household, the government is not subject to an exogenous borrowing limit (or at least, this limit does not bind for the

44

purposes of our policy exercise).

Market clearing

Prices are flexible and the market for the goods and services clears:

$$p_{t,k} y_{t,k} = p_{t,k} (c_{t,k}^s + c_{t,k}^b + g_{t,k}) + \sum_j p_{t,k} x_{t,j,k}.$$

The rental rate of capital is also flexible and so capital is always fully employed. The bond market also clears; however, the price of bonds are set by the central bank according to a Taylor rule:

$$1 + i_t = \max\{1, (1 + R_t^n) \left(\frac{p_{t+1}}{p_t}\right)^\phi\},$$

where R_t^n is the net Wicksellian natural rate of interest and $\phi > 1$. The model can feature multiple equilibria, and we will discuss equilibrium selection later.

The labor market is subject to a nominal friction. Specifically, wages are downwardly rigid in the spirit of Patinkin (1965), Malinvaud (1977), and more recently Schmitt-Grohé and Uribe (2011). That is

$$l_t \leq \bar{l}, \quad w_{t-1} \leq w_t, \quad (w_t - w_{t-1})(l_t - \bar{l}) = 0.$$

Here, $\bar{l}$ is an exogenous endowment of labor, and w_t is the nominal wage in period t. This is a transparent and tractable way of adding nominal frictions into the model. The key assumption here is that in the event of a shortfall in nominal demand, it is the labor market that fails to clear, and not the capital market. Partial equilibrium in the labor market is shown in log-log terms in figure 1.13, where the sales of firms are held constant. There are two admissible regions: (1) $w_t \geq w_{t-1}$ and the labor market clears, and (2) $w_t = w_{t-1}$ and the labor market fails to clear with $l_t \in (0, \bar{l})$.

There is considerable empirical evidence for downward stickiness in wages, see for instance Barattieri *et al.* (2010), Baqaee (2014a), Dickens *et al.* (2007), and Bewley (1999). Downward wage rigidity is a particularly convenient modelling device in this paper since it makes the intuition for government intervention very transparent – there is idle labor and a
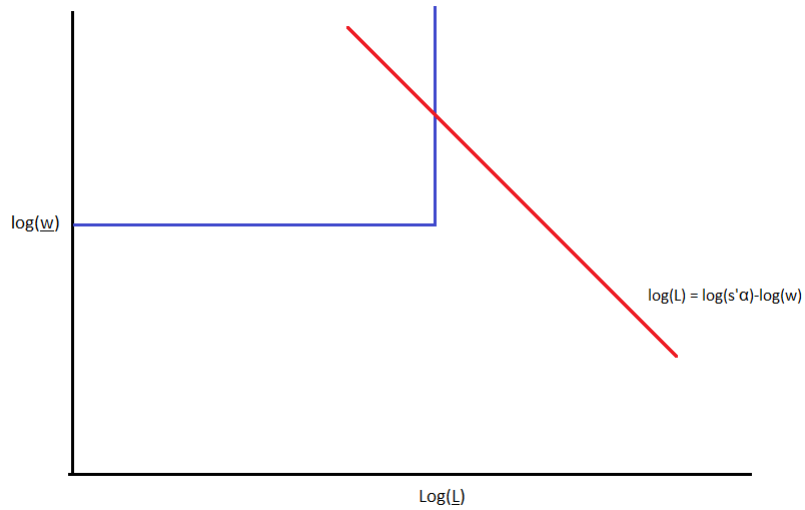
Figure 1.13: *The labor market with downward wage rigidity. The blue line is labor supply and the red line is labor demand. The vector s is the sales of each industry, while the vector α is the gross labor share of each industry.*

unique efficient level of full employment. Such a stark assumption is not strictly necessary however, since similar forces operate as long as output is inefficiently low and labor (rather than capital) is the factor that adjusts. This can be accomplished with, for example, an elastic labor supply curve and sticky prices, and I sketch a version of this model in Appendix V.

1.5.3   Scenario I: Opportunistic Spending

In this section, I analyze optimal fiscal policy during a one-period liquidity trap with only the Ricardian households. In other words, I assume that the fraction of the population corresponding to impatient borrowers is zero. Under this assumption, there is no neoclassical multiplier effect of government spending since labor supply is inelastic (so there is no wealth effect of taxation). There is also no Keynesian multiplier effect for private consumption since we have a one-period shock. Therefore, the results of this subsection pertain to pure opportunistic fiscal policy.

The shock that pushes the economy into the zero lower bound, as is common in the representative agent zero-lower bound literature, following Krugman (1998), is a one-period

unexpected discount factor shock. Suppose that there is an unexpected discount factor shock so that for the next period $\rho^* > 1$. I analyze the government's fiscal policy without commitment – that is, the government reoptimizes its expenditure plans each period.

Lemma 1.5.3. *Aggregate labor demand l_t when the zero lower bound binds is upward sloping in wage inflation, and is given by*

$$l_t = \frac{1}{\rho^*} \frac{w_{t+1}}{w_t} \bar{l} - \frac{1}{\rho^*} \delta'_{t+1} \tilde{\alpha} \tau_{t+1} - (\beta - \delta_t)' \tilde{\alpha} \tau_t. \tag{1.8}$$

Proof. See Appendix I. ∎

Lemma 1.5.3 shows how government spending today τ_t, by deviating from private spending $(\beta - \delta_t)$, can increase employment.

Equilibrium employment is given by combining aggregate demand for labor with the aggregate supply curve for labor. Aggregate supply for labor is defined by $w_{t+1}/w_t \geq 1$ and $l_t \leq \bar{l}$. This situation is graphically depicted in figure 1.14. There are two equilibria. One is the neoclassical equilibrium where the wage rises by exactly enough tomorrow to ensure we maintain full employment. In this case, the government need not intervene to boost employment and government expenditure shares are equal to the Cobb-Douglas parameters. The second equilibrium, which is the equilibrium of interest, features no wage inflation, $w_t = w_{t+1}$, and positive unemployment.

In the non-neoclassical equilibrium, government policy can affect employment. We focus on this second equilibrium since it is the one in which fiscal policy is relevant, and the one that features an employment problem.[4]

Proposition 1.5.4. *The optimal share of expenditures by the government in industry i relative to industry j satisfies*

$$\frac{\delta_i}{\delta_j} = \frac{\phi_i}{\phi_j} \left(\frac{const + (\mu_1 - \mu_2)\tilde{\alpha}_j}{const + (\mu_1 - \mu_2)\tilde{\alpha}_i} \right), \tag{1.9}$$

---

[4]It turns out that the full employment equilibrium is also locally unstable since the aggregate supply relation is steeper than aggregate demand. See Eggertsson and Mehrotra (2014) for more details on the stability of these equilibria.
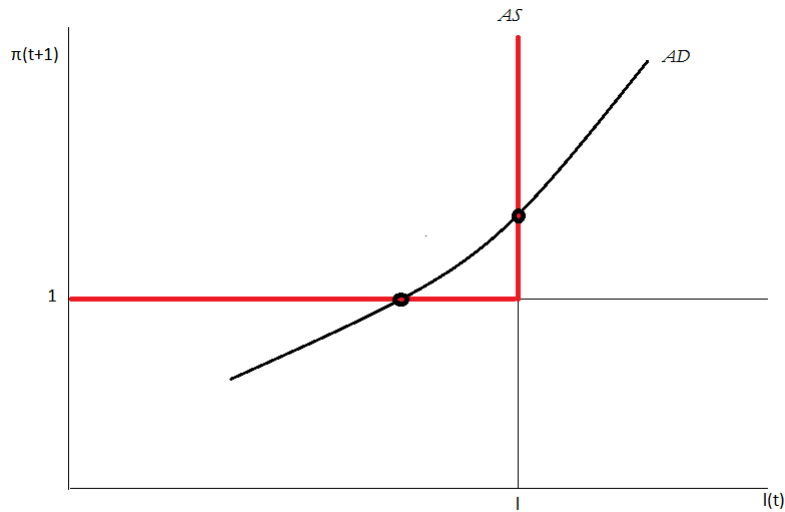
Figure 1.14: *Zero-lower bound is binding and aggregate demand is upward sloping. The y-axis is wage inflation and the x-axis is employment.*

where μ_1 and μ_2 are lagrange multipliers corresponding to the labor and capital markets. When we are at the full employment steady state $\mu_1 = \mu_2$, so that spending shares are equal to the Cobb-Douglas parameters. When there is unemployment $\mu_1 < \mu_2$, so government tilts in favor of firms with higher network-adjusted labor intensities.

Proof. See Appendix I.   ■

The key intuition of this section is that when the zero lower bound binds, the government has an opportunity to provide goods and services to the household more cheaply than usual. Furthermore, the higher the network-adjusted labor intensity of an industry, the more cheaply the government can supply that good to the household. Therefore, the government tilts its expenditures in favor of industries with high network-adjusted labor intensity.

Equation (1.9) is intuitive to interpret. The production of each good uses a certain combination of labor and capital, directly and indirectly through inputs. When there is unemployment, there is idle labor that is essentially free to use for the government. However, capital is not free. Therefore, the government tilts its consumption of goods towards those that use labor more heavily than capital, since any capital used by the government crowds

48

out the private sector, and bids up rents rather than expand production. The existence of a wedge between private and public spending decisions follows from the logic set out by Farhi and Werning (2013).

The intuition for (1.9) can be illustrated by considering an extreme example with only two goods: one good only uses labor and the other only uses capital with no intermediaries. When there is unemployment, the government can use the labor-intensive good without reducing the household's consumption of labor. However any capital used by the government crowds out the household. Therefore, the government will use all the unemployed labor, but only use enough of the capital intensive good to equalize the marginal utility of government and household consumption.

1.5.4 Scenario II: Stimulus Spending

Now, let us consider the case where government expenditures have zero direct value, but since there are credit-constrained borrowers, multiplier effects of government spending give the government a motive to spend during a liquidity trap. Following Eggertsson and Krugman (2012), an exogenous debt limit on borrowers makes them non-Ricardian and forces them to behave as what Galí *et al.* (2007) call "hand-to-mouth consumers." The assumption that a fraction of households' consumption tracks their current income rather than permanent income accords with the empirical findings of Campbell and Mankiw (1990).

Set the fraction of borrowers to χ to be nonzero. To shut down the opportunistic channel, set the utility-value of government consumption $\lambda = 0$, so that government expenditures have no intrinsic value to the household. The only reason why government expenditures may be beneficial in this context is then the stimulus effect of spending due to the presence of non-Ricardian households.

Deleveraging Shock

Since the borrower has a smaller discount factor, the steady-state equilibrium of this model sees the borrower borrow up to his borrowing constraint from the household. We focus on the steady-state equilibrium with no inflation and no government spending or taxation.

Now suppose that a borrowing limit falls unexpectedly in period t so that

$$D_{t-1} = \frac{D^h}{1+i_{t-1}} \frac{p_t}{p_{t-1}}, \quad D_t = \frac{D^l}{1+i_t} \frac{p_{t+1}}{p_t},$$

where $D^h > D^l$. Assume that the borrower has to delever immediately to the new borrowing limit.

I analyze the equilibrium where the steady-state equilibrium features zero inflation, full employment, and no government spending and constant government taxes.

Lemma 1.5.5. *Aggregate demand for labor l_t when the zero lower bound binds is upward sloping in wage inflation and given by*

$$w_t l_t = \frac{\beta'\tilde{\alpha}}{\rho}\left(w_{t+1}\bar{l} + r_{t+1}\bar{k}\right) + \beta'\tilde{\alpha}\frac{D^l - \rho D^h}{\rho(1-\chi)} + \left(\beta'\tilde{\alpha}\left[\frac{1-(1-\rho)(1-\chi)}{\rho(1-\chi)} - 1\right] + \delta'\tilde{\alpha}\right)p_t^g g_t.$$

(1.10)

Proof. See Appendix I.   ∎

Equation (1.10) shows that aggregate demand for labor can be stimulated in two ways by the government. The first channel is the same as the one in the Ricardian model: if the government buys labor intensive goods ($\delta_t'\tilde{\alpha} > 0$), then the government increases nominal demand for labor directly.[5] However, there is now a new, non-Ricardian channel. If $\chi > 0$, then there is a secondary effect because increased government spending increases the contemporaneous income of borrowers and therefore increases private expenditures.

The term

$$\left[\frac{1-(1-\rho)(1-\chi)}{\rho(1-\chi)} - 1\right],$$

---

[5]There is no crowding-out of the private sector because we are at the zero-lower bound, and current private nominal consumption is pinned down by the Euler equation.

is the government multiplier on nominal private GDP. As long as the fraction of borrowers χ is nonzero, this is greater than zero, and so there is a pure stimulus effect to government spending. If we set the fraction of borrowers to be zero, then the model is Ricardian and, as shown in the proof of proposition 1.5.6 in appendix I, the government multiplier on private nominal GDP is exactly zero.

Combining the aggregate demand equation (1.10) with the the aggregate supply relation gives the situation in figure 1.14. Once again, as we see from the graph, there are two equilibria. In the first, there is no inflation, positive unemployment, and increases in government spending increase output and weakly increase inflation. The other equilibrium is the neoclassical equilibrium where we have positive inflation equal to exactly the reciprocal of the gross natural interest rate, full employment, and increased government spending reduces inflation. In the neoclassical equilibrium, the government need not intervene and government spending should be zero. This is the uninteresting equilibrium for our purposes. Therefore, we focus instead on the non-neoclassical equilibrium with positive unemployment.

By inspection of (1.10), we can see that the biggest bang for buck in terms of the government boost to employment comes from maximizing the size of the government's network-adjusted labor intensity $\delta'\tilde{\alpha}$. Maximizing employment is not the government's objective however. Optimal government policy seeks instead to maximize real GDP net of government consumption.

Proposition 1.5.6. *Optimal government spending in the period of the deleveraging shock has the government spend entirely on the industry with the highest network-adjusted labor intensity.*

The intuition here is that the industry with the highest network-adjusted labor intensity not only has the largest employment multiplier, since it employs the most amount of idle labor, but it is also the cheapest resource to waste (since all government spending is wasteful). Therefore, the industry with the highest network-adjusted labor intensity is targeted.

In this setup, borrowers and savers derive the same income from labor and capital. If we modify the model so that borrowers derive more income from labor than capital, as is

51

empirically more relevant, these results would become even stronger. Then not only will high network-adjusted labor intensity imply that those goods are cheaper to waste, but it also means that their owners have higher marginal propensities to consume, and therefore, will have even higher multipliers.

Furthermore, labor is homogenous in this model. A simple extension of the model with different labor types would have the government target sectors that more intensively use low-skilled labor, both because it is cheaper to waste and because it gives larger multipliers. The relevant criteria for the target would be the network-adjusted labor intensity by type.

1.5.5 Practical Application

All three scenarios point to the government targeting its stimulus towards sectors with higher network-adjusted labor intensities. Table 1.4 reports the network-adjusted labor intensity for broad categories of final-use spending, including various types of government spending, assuming Cobb-Douglas functional forms. The network-adjusted labor intensities can be interpreted as determining the relative employment multipliers of an extra dollar of spending from the different final use sectors. Crucially, these numbers are also the fraction of each dollar of spending that is eventually spent on labor — or in other words, labor's share of income from that final-use sector.

If the government has the ability to finely tune stimulus spending, then the industry-level rankings discussed in section 1.3 are the relevant statistics for designing stimulus. If the government used value-added rankings instead, it would, to pick an extreme example, erroneously think that the multiplier for "Soybean and other oilseed processing" is 13 times larger than that of "other residential structures," since the former has a value-added labor share of 0.903, while the latter's value-added labor share is only 0.069. However, using the network-adjustment, we find that Soybean and other oilseed processing has a labor share of 0.247, while other residential structures has a labor share of 0.350 – a reversal in rankings. Generally, we think that governments are unable to perfectly fine-tune stimulus spending, for reasons outside of the model. Therefore in table 1.4, we look at broad classes

of government spending.

Table 1.4: *Network-adjusted and gross labor intensities of various final goods consuming sectors using the BEA's detailed purchase price benchmark input-output table for 2007.*

Final-Use	Network-adjusted	Value-Added
Personal consumption expenditures	0.503	0.537
Private fixed investment	0.572	0.594
Federal Government defense: Consumption expenditures	0.484	0.432
Federal Government defense: Gross investment	0.638	0.726
Federal Government nondefense: Consumption expenditures	0.747	0.796
Federal Government nondefense: Gross investment	0.685	0.800
State and local government consumption expenditures	0.747	0.880
State and local government gross investment	0.584	0.644

Since the data in table 1.4 is very heavily aggregated, we should expect the network structure to matter less for the rankings. Nonetheless, even with this level of aggregation, we see some interesting patterns. Namely, the network-adjusted labor intensities are much closer to one another than the value-added measures. We also see that nondefense investment is less labor intensive than nondefense consumption, despite the value-added measure being larger. The implied relative ranking of multipliers for defense and non-defense spending may help to explain why the literature estimating government multipliers tends to find smaller multipliers for defense spending than other types of government spending. For a summary of the contrasting estimates of defense and non-defense multipliers see Yang *et al.* (2012). They find that non-defense multipliers are 1.5-2.0 times larger than defense multipliers.

1.5.6 Who Gets Paid?

The fact that unemployment rates vary significantly with education levels suggests that this exercise will be more informative if we focus on the multipliers associated with the lower skill types since those labor markets are more likely to experience high cyclical unemployment. In table 1.5, I report each labor type's share of income by final-use normalized by that labor type's share of aggregate income. A number less than one implies that the final-use sector uses that labor type less intensively than average, whereas a number greater than

one implies the opposite. Table 1.5 effectively answers the question of where the money for different kinds of government spending goes – a question that cannot be satisfactorily answered without the use of network-adjusted labor intensities.

Table 1.5: *Fraction of final use expenditures going to each type of labor normalized by that labor type's share of total GDP. A number greater than (less than) 1 indicates the final sector in that row spends more (less) heavily on that labor type compared to total GDP. This table combines data from the 2007 American Community Survey from IPUMS-USA with the detailed Benchmark Input-Output table using purchaser prices for 2007 published by the BEA.*

Final-Use	Grades 0-9	Grade 10-12	1-2 years of college	4 years of college	5+ years of college
Personal consumption expenditures	0.902	0.908	0.938	0.911	0.935
Private fixed investment	1.338	1.243	1.016	1.016	0.715
Federal defense consumption	0.513	0.655	0.793	0.928	1.382
Federal defense investment	0.906	1.041	1.090	1.402	1.145
Federal nondefense consumption	0.569	0.855	1.185	1.523	2.347
Federal nondefense investment	0.851	1.002	1.164	1.608	1.304
State and local government consumption	1.090	1.227	1.304	1.271	1.405
State and local government investment	1.948	1.484	1.043	0.839	0.544

Table 1.5 shows that private fixed investment and state and local government investment use low-skill labor much more and high-skill labor much less than average. On the other hand, federal consumption (defense and nondefense) are overwhelmingly tilted towards high-skill types. Table 1.5 implies that a uniform reallocation of funds from private consumption towards federal government spending would increase high-skilled labor's share of income at the expense of low-skilled labor. To the extent that low-skilled labor markets are experiencing greater slack, this table helps to explain why estimates of fiscal multipliers for state and local government expenditures, like those of Shoag (2010) and Chodorow-Reich *et al.* (2012) tend to find larger effects than estimates from federal expenditures.

1.6   Conclusion

This paper introduces the network-adjusted labor intensity as the relevant notion of labor intensity in an interconnected production economy. This captures how intensively a good or service uses labor in production by taking into account how heavily its entire supply chain relies on labor. Doing this adjusts for the artificial drop in the gross labor intensity resulting

from fragmentation of the production process across industries.

The network-adjusted labor intensity plays a key role in determining how sectoral disturbances translate to aggregate employment, and this has both short-run and long-run implications. For instance, labor's share of income, a central object of interest in the study of growth and inequality, is a weighted average of network-adjusted labor intensities. This allows us to decompose labor's share of income into disaggregated components representing changing consumption patterns, changing supply chains (including trade), and changing capital/labor shares. In a sample of 34 countries over the past 15 years, this decomposition shows that the overwhelming driver of the secular decline in the labor share is a decline in the gross labor share of all industries. This contrasts with the usual value-added decomposition, which over the sample, attributes the drop to changing industrial composition.

The network adjustments also allow us to study individual industries, like manufacturing, without discarding information about changes in their supply chains. For the US manufacturing sector, increasing globalization and changing consumption patterns do explain a sizeable fraction of the decline in manufacturing's network-adjusted labor share, but these effects are not sizeable when aggregated up to the whole economy.

Over the short run, the network-adjusted labor intensities pin down the relative boost to employment from a marginal increase in spending in one industry versus another. This makes network-adjusted labor intensities important to policy makers interested in boosting employment through fiscal policy. I show that if in a recession labor is the factor that adjusts in production, then when the zero lower-bound binds, optimal fiscal stimulus should tilt in favor of stimulating demand in industries with higher network-adjusted labor intensities. The intuition is that, in a recession, the government should aim to expand production rather than to simply bid up rents. The way to expand production is to stimulate industries that are most reliant on unemployed resources. In my model, the unemployed resource is labor and the relevant notion of "reliant" is the network-adjusted labor intensity.

We can also compute measures of how intensively different types of government expenditures use different types of labor. I find that state and local government expenditures are

much more reliant on low-skilled labor than average. On the other hand, federal government expenditures, defense and nondefense, are far more heavily reliant on very high-skilled workers with more than 4 years of college education than average. Furthermore, I find that on the whole defense expenditures are less reliant on labor than other types of government expenditures. These findings go some way towards explaining the heterogeneity in estimates of the effect of government expenditures on employment found in the literature.

Chapter 2

Cascading Failures in Production Networks[1]

2.1 Introduction

In this paper I show how the extensive margin of firm entry and exit can dramatically alter the properties of macroeconomic models with production networks. I model cascades of failures among firms linked through a production network, and show how the network propagates and amplifies shocks through supply and demand chains. This paper contributes to the literature on the microeconomic sources of aggregate business cycle fluctuations.

In a recent paper Acemoglu *et al.* (2012) relate the following anecdote, which illustrates the basic mechanism that I model and demonstrates its real-world relevance:

> In the fall of 2008, rather than asking for government assistance for Ford, Alan R. Mulally, the chief executive of Ford Motor Co., requested that the government supports General Motors and Chrysler. His reasoning for asking government support for his company's traditional rivals was that the failure of either GM or Chrysler would lead to the potential failure of their suppliers, and because Ford depended on many of the same suppliers as the other two automakers, it would also find itself in perilous territory.

---

[1]Co-authored with my other advisor

A government bailout of GM and Chrysler in 2009 prevented the failure of GM and Chrysler (see Goolsbee and Krueger, 2015). However, the scenario Mulally described, which was averted through government intervention in the United States, did come to pass in Australia. In May 2013, Ford Australia announced that they would stop manufacturing cars in 2017. Seven months later, GM Australia announced they would also stop manufacturing cars in 2017. Three months after that, in February 2014, Toyota Australia also announced that it would close its manufacturing plants at the same time. This effectively ended automobile manufacturing in Australia. The Australian government predicts that this will result in the loss of over 30,000 jobs, a figure they arrived at by adding the number of people directly employed by the three automakers and the Australian car parts industry.

While these examples demonstrate that firm exit (and entry) can have important spillovers on other firms, standard macroeconomic models do not allow for this possibility. In this paper, I explicitly incorporate the extensive margin of firm entry into an input-output model of production. I show how the extensive margin alters the quantitative and qualitative properties of the model. First, I show that the standard input-output macroeconomic models that follow Long and Plosser (1983), like Acemoglu *et al.* (2012), Atalay (2013), or Baqaee (2014c) have the property that their responses to productivity shocks can be summarized in terms of a few exogenous sufficient statistics. Once we compute the relevant sufficient statistics, which are closely related to the equilibrium size of firms, we can discard the network structure. In other words, I show that there are disconnected economies, with different structural parameters (and sometimes exogenous wedges), that behave precisely like the network models. This means that, without the extensive margin, it is not the interconnections per se, but how those interconnections affect a firm's size that determines a firm's systemic importance. If we can arrive at the same sufficient statistics using a different (perhaps degenerate) network-structure, the equilibrium responses will be the same. This fact explains why the theoretical implications of the granular hypothesis of Gabaix (2011), where business cycles are driven by large firms, are observationally equivalent to the theoretical implications of the network hypothesis of Acemoglu *et al.* (2012), where business

cycles are driven by well-connected firms. Furthermore, in canonical input-output models, systemic importance depends only on a firm's role as a supplier. As long as firms i and j have the same strength connections to the same customers, then their systemic influence will be the same, regardless of what i and j's supply chains look like.

When we allow for entry and exit, the sufficient-statistic approach breaks down. This is because the extensive margin makes "systemic importance" an endogenously determined value that is not well-approximated by equilibrium size or prices. A firm that may seem like a small player, when measured by sales, can have potentially large impacts on aggregate outcomes. On the other hand, a firm that may seem like a key player, as measured by sales, can have relatively minor effects on the equilibrium. Furthermore, the endogenously-determined measures of systemic influence depend on a firm's role as both a supplier *and* as a consumer, as well as on how many close substitutes there are for the firm.

This allows us to combine the granular hypothesis of Gabaix (2011) and the network hypothesis of Acemoglu *et al.* (2012) in a new and interesting way. In particular, I show that once firms have positive mass, extensive margin shocks can be *locally* amplified via interconnections. In other words, when a large and well-connected firm exits, it can set off an avalanche of firm failures that actually gets larger as it gathers steam. This type of amplification is not possible unless we have both granularity and network connections. With positive mass, the model satisfies the criteria of Scheinkman and Woodford (1994) for self-organized criticality: it exhibits strong local interactions that are significantly nonlinear.

This paper also contributes to the wider literature on diffusion on social networks, by bridging the gap between two alternative modelling traditions. Loosely speaking, there are two popular approaches to modelling diffusion on social networks. First, there are continuous input-output type models like Acemoglu *et al.* (2012). Here, nodes influence each other in continuous ways – shocks travel away from their source like waves and slowly die out. The strength of the connections between the nodes controls the rate of decay. Such shocks, sometimes called pulse processes, are characterized by geometric sums. I show that these models are incapable of *local* amplification: a productivity shock to an industry

59

will always have its largest effect at its source, and the shock decays as it travels through the connections. These structures were first studied by Leontief (1936) in his input-output model of the economy.

The other camp consists of models that behave discontinuously, as typified by Morris (2000) or Elliott *et al.* (2012). In this class of models, sometimes called threshold models, each node has a threshold and is either active or inactive. When a node crosses its threshold it changes states and, by changing states, pushes its neighbors closer to their thresholds. Such models are frequently used to study the spread of epidemics, products, or even ideas. One of the earliest and most influential threshold models is the Schelling (1971) model of segregation. Threshold models do not have wave-like properties since the rate at which shocks decay are not geometric. Crucially, these models are capable of generating local amplification – that is, shocks can be amplified as they travel through the network; however these models are notoriously difficult to analyze.

In this paper, I consider a model that bridges the gap between the continuous and discrete models. Specifically, I explicitly account for the mass of firms in a given industry. Industries with a continuum of firms behave continuously – the mass of firms responds continuously to shocks. On the other hand, lumpy industries, with only a few firms, behave discontinuously. For instance, a negative shock to an unconcentrated industry, say hairdressing, will result in some fraction of hairdressers exiting. The fraction exiting will be a continuous function of the size of the shock. The effect on a neighboring industry will be attenuated by the strength of its connections to hairdressing. However, a negative shock to a highly concentrated industry, like automobile manufacturing, will have no effect on the number of firms unless it is large enough to force an exit. But once a large firm exits, it imparts an additional impulse to the size of the shock which can trigger a cascade. Because the model is flexible enough to express both behaviors, I can provide conditions under which we can expect a continuous approximation to a discontinuous model to perform badly.

The idea of cascading – domino-like – chain reactions also appears outside of economics.

In particular, models of contagion and diffusion like the threshold models considered by Kempe *et al.* (2003) have these features. In these models, notions of connectedness play a key role since the only way contagion can spread is via connections between nodes. An interesting implication of embedding a contagion model into a general equilibrium economy is the role prices and aggregate demand play – a role that does not have analogues in other threshold models. Typically, in a threshold model, shocks can only travel along edges. Contagion can only spread to nodes who are connected to an infected node. This important intuition breaks down in general equilibrium models since all firms are linked together via aggregate demand. This means that general equilibrium forces can act like long-distance carriers of disease. Shocks in one fragile industry, like the financial industry, can jump via aggregate demand, to a different fragile industry like automobile manufacturing even if these two industries are not connected.

The structure of paper is as follows. In section 2.2, I set up the model and define its equilibrium. In section 2.3, I characterize the equilibrium conditional on the mass of entrants in each industry and define some key centrality measures. In section 2.4, I study how the model behaves when the extensive margin of firm entry and exit is shut down. I prove results showing that the network structure can be summarized by sufficient statistics related to size. I also show that we can think of these models as non-interconnected models with different parameters. Finally, I show that this class of models is incapable of local amplification of shocks. In section 2.5, I allow firm entry and exit. First, I characterize the model's responses to shocks in the limit where all firms are massless. I show that sufficient statistics are no longer available and that systemic importance is endogenous. Then, I consider the case when firms can have positive mass, and show that with atomistic firms, shocks can be locally amplified. I prove an inapproximability result showing conditions under which we should expect a continuous approximation to a discontinuous model to perform badly. In section 2.6, motivated by the Ford example I discuss above, I consider conditions under which firms' incentives align with those of society. Specifically, when can we trust one firm's testimony about whether or not another firm should be bailed out. I

conclude in section 2.7.

2.2 Model

In this section, I spell out the structure of the model and define the equilibrium. There are three types of agents: households, firms, and a government. Each firm belongs to an industry, and there are N industries.

The households in the model are homogenous with a unit mass. The representative household maximizes utility

$$U(c_1, \ldots, c_N) = \left(\sum_{k=1}^{N} \beta_k^{\frac{1}{\sigma}} c_k^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

where c_k represents composite consumption of varieties from industry k and $\sigma > 0$ is the elasticity of substitution across industries. The composite consumption good produced by industry k is given by

$$c_k = \left(\sum_{i=1}^{N_k} \Delta_k c(k, i)^{\frac{\varepsilon_k-1}{\varepsilon_k}} \right)^{\frac{\varepsilon_k}{\varepsilon_k-1}},$$

where $c(k, i)$ is household consumption from firm i in industry k and $\varepsilon_k > 1$ is the elasticity of substitution across firms within industry k. Here, N_k is the number of firms active in industry k and Δ_k is the mass of each firm. The assumption that Δ_k is constant for all firms in industry k means firms in each industry are homogenous. The total mass of firms in industry k is given by $M_k = N_k \Delta_k$. The household's budget is given by

$$\sum_{k,i} p(k, i) c(k, i) = wl + \sum_{k,i} \pi(k, i) - \tau,$$

where $p(k, i)$ is the price of firm i in industry k and $\pi(k, i)$ is firm i in industry k's profits. The wage is w and labor is inelastically supplied at l. For the rest of the paper, and without loss of generality, we take labor to be the numeraire so that $w = 1$, and fix the supply of labor $l = 1$. Lump sum taxes by the government are denoted τ.

Let firm i in industry k maximize profits

$$\pi(k,i) = p(k,i)y(k,i) - \sum_{l=1}^{N}\sum_{j}^{N_l} p(l,j)x(k,i,l,j) - wl(k,i) - wf_k + \tau_k,$$

where $p(k,i)$ is the price and $y(k,i)$ is the output of the firm. Inputs from firm j in industry l are $x(k,i,l,j)$ and labor inputs are $l(k,i)$. Finally, in order to operate, each firm pays a fixed cost f_k in units of labor and the firm potentially receives a lump sum subsidy τ_k. The mass parameter Δ_k controls how finely the fixed costs of industry k can be split up – in other words, it captures increasing returns to scale at the industry level. The firm's production function (once the fixed cost has been paid) is constant returns to scale

$$y(k,i) = \left(\alpha_k^{\frac{1}{\sigma}}(z_k l(k,i))^{\frac{\sigma-1}{\sigma}} + \sum_{l=1}^{N} \omega_{k,l}^{\frac{1}{\sigma}} x(k,i,l)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}.$$

Here, $\sigma > 0$ is again the elasticity of substitution among inputs, and ω_{kl} is the CES share parameter for how intensively firms in industry k use composite inputs from industry l. The $N \times N$ matrix of ω_{kl} determine the network-structure of this economy. One can think of this matrix as the adjacency matrix of a weighted directed graph. The parameter $\alpha_k > 0$ gives the intensity with which firms in industry k use labor.

Labor productivity shocks, like the ones considered by Acemoglu *et al.* (2012) and Atalay (2013) are denoted by z_k. Note that when $\sigma \neq 1$, a productivity shock z_k to industry k is equivalent to changing that industry's labor intensity from α_k to $\alpha_k z_k^{\sigma-1}$. Therefore, as long as $\sigma \neq 1$, we can think of α_k as including both the productivity shock and the labor intensity. This way we do not need to directly make reference to the shocks z since they are just equivalent to changing α. This equivalence breaks down when $\sigma = 1$, and in those cases, we shall have to work directly with z. For the majority of this paper, I focus on the propagation of productivity shocks. When the elasticity of substitution is equal to one, these are precisely the shocks considered by Acemoglu *et al.* (2012). The same methods can easily be used to study other shocks. I defer the discussion of how other shocks, like fixed-cost shocks or demand shocks, would affect the results to the end of the paper.

The composite intermediate input from industry l used by firm i in industry k is

$$x(k,i,l) = \left(\sum_{j=1}^{N_l} \Delta_l x(k,i,l,j)^{\frac{\varepsilon_l - 1}{\varepsilon_l}} \right)^{\frac{\varepsilon_l}{\varepsilon_l - 1}},$$

where ε_l is the elasticity of substitution across different firms within industry l. Note that the elasticities of substitution are the same for all users of an industry's output.

The government runs a balanced budget so that

$$\sum_k \tau_k = \tau.$$

We study the subgame perfect Nash equilibrium. In period 1, entry decisions are made simultaneously. In period 2, firms play monopolistic competition conditional on period 1's entry decisions.

Definition 2.2.1. A monopolistically competitive equilibrium is a collection of prices $p(i,k)$, wage w, and input demands $x(i,k,l,j)$, outputs $y(i,k)$, consumptions $c(i,k)$ and labor demands $l(i,k)$ such that for mass of entrants $\{M_k\}_{k=1}^N$ and vector of productivity shocks z_k,

  (i) Each firm maximizes its profits taking as given the industrial price level and industrial demand,

 (ii) the representative household chooses consumption to maximize utility,

(iii) the government runs a balanced budget,

(iv) markets for each good and labor clear.

Note that the productivity shock is known at the start of the game. Changing the information structure to study the effects of uncertainty on the actions of agents is an interesting extension that I leave for future work.

Let $\Pi : \mathbb{R}_+^N \times \mathbb{R}_+^N \to \mathbb{R}^N$ be the function mapping the masses of entrants M and vector of productivity shocks z to industrial profits assuming monopolistic competition in period 2. In theorem 2.3.3, I analytically characterize this function.

Definition 2.2.2. A vector of integers $\{N_k\}_{k=1}^{N}$ is an equilibrium number of entrants if

$$\Pi_i(M_i, M_{-i}, z) \geq 0 > \Pi_i(M_i + \Delta_i, M_{-i}, z) \quad (i \in \{1, 2, \ldots, N\}),$$

where $M_j = N_j \Delta_j$, for all j.

Intuitively, a vector of integers is an equilibrium number of entrants if all firms make non-negative profits, and the entry of an additional firm in any industry results in firms in that industry making negative profits.

Notation

Let e_i denote the ith standard basis vector. Let Ω be the $N \times N$ matrix whose ijth element is equal to ω_{ij}. Let α and β be the $N \times 1$ vectors consisting of α_is and β_is. Let $\tilde{M}$ be the $N \times N$ diagonal matrix whose ith diagonal element is equal to $M_i^{\frac{1}{\varepsilon_i - 1}}$, and let M be the $N \times N$ diagonal matrix whose ith element is M_i. Finally, let μ be the $N \times N$ diagonal matrix whose ith diagonal element is the mark-up $\varepsilon_i / (\varepsilon_i - 1)$ charged by firms in industry i. Let $\circ$ denote the element-wise or Hadamard product, and $\text{diag} : \mathbb{R}^N \to \mathbb{R}^{N^2}$ be the operator that maps a vector to a diagonal matrix.

Definition 2.2.3. An economy E is defined by the tuple $E = (\beta, \Omega, \alpha, \varepsilon, \sigma, f, \Delta)$. The vector β contains household taste parameters, Ω captures the input-output share parameters, α contains the industrial labor share parameters, $\varepsilon > 1$ is the vector of industrial elasticities of substitution, $\sigma > 0$ is the cross-industry elasticity of substitution, f is the vector of fixed costs, and Δ is the vector of masses of firms in each industry.

Before analyzing the model, it helps to define some key statistics. These are standard definitions from the literature on monopolistic competition. See, for example, Bettendorf and Heijdra (2003).

Definition 2.2.4. The price index for industry k is given by

$$p_k = \left(\sum_i^{N_k} p(k, i)^{1 - \varepsilon_k} \right)^{\frac{1}{1 - \varepsilon_k}},$$

and the total composite output of industry k is given by

$$y_k = \left(\sum_i^{N_k} y(k,i)^{\frac{\varepsilon_k-1}{\varepsilon_k}} \right)^{\frac{\varepsilon_k}{\varepsilon_k-1}}.$$

The consumer price index, which represents the price level for the household, is given by

$$P_c = \left(\sum_k^N \beta_k p_k^{1-\sigma} \right)^{\frac{1}{1-\sigma}},$$

and total consumption by the household is given by

$$C = \left(\sum_{k=1}^N \beta_k^{\frac{1}{\sigma}} c_k^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

which is just the utility of the household.

These are the "ideal" price and quantity averages for each industry. The reason we do not simply average prices to get a price index or add outputs to get total output is because even within each industry, each firm is producing a slightly different product. When the elasticity of substitution $\varepsilon_k = 0$, then the price index for that industry is simply the sum of all the prices since there is no substitution and a consumer of this industry must buy all the varieties. When the elasticity of substitution $\varepsilon_k \to \infty$, the price index for an industry is just the minimum price, since households will only purchase from the cheapest firm in each industry. An important special case is when $\varepsilon_k \to 1$, where the price index is a geometric average of the industrial prices.

2.3   Monopolistic Competition Subgame

In this section, I characterize the equilibrium in the monopolistic competition subgame of period 2, conditional on the number of entrants in each industry. Before stating any results, it helps to define some key industrial statistics.

Let us define supply-side and demand-side centrality measures.

Definition 2.3.1. The supplier centrality is

$$\tilde{\beta}' = \beta' \Psi_s,$$

where

$$\Psi_s = (I - \tilde{M}^{\sigma-1} \mu^{-\sigma} \Omega)^{-1} = \sum_{n=0}^{\infty} \left(\tilde{M}^{\sigma-1} \mu^{-\sigma} \Omega \right)^n.$$

This captures the frequency with which each industry appears in supply chains. The kth element of $\tilde{\beta}$ captures demand from the household that reaches industry k, whether directly or indirectly through other industries who use k's products. The following helps establish why we might care about $\tilde{\beta}$

Lemma 2.3.1. *In equilibrium,*

$$\tilde{\beta}_i = \left(\frac{p_i^{\sigma} y_i}{P_c^{\sigma} C} \right),$$

where P_c is the ideal price index for the household, C is total consumption by the household, p_i is industry i's ideal price index and y_i is industry i's composite output.

Note that in the Cobb-Douglas limit, where the elasticity of substitution across industries σ is equal to one, $\tilde{\beta}_i$ is precisely an industry's share of sales. In the Cobb-Douglas case, $\tilde{\beta}$ coincides with the *influence measure* defined by Acemoglu *et al.* (2012).

There are two reasons to think of $\tilde{\beta}$ as a supplier centrality. First, since it measures how frequently an industry appears in other agents' supply chains, it means that star suppliers have high $\tilde{\beta}$. Secondly, as we shall see, $\tilde{\beta}$ captures the response of output to productivity shocks. For instance, Acemoglu *et al.* (2012) show that in a Cobb-Douglas model without an extensive margin, $\tilde{\beta}$ captures the extent to which industry-specific labor productivity shocks affect output. Furthermore, since $\tilde{\beta}$ is share of sales in a Cobb-Douglas economy, the work of Hulten (1978) implies that $\tilde{\beta}$ maps marginal TFP shocks to aggregate output. The intuition for these results is clear: if a firm is supplying a large fraction of the economy, then its productivity shocks have a large impact on output.

Note two important facts about $\tilde{\beta}$. First, if there is no entry, so that $\tilde{M}$ is constant, then $\tilde{\beta}$ is exogenous with respect to productivity shocks. Second, note that $\tilde{\beta}_k$ depends *only* on

industry k's consumers, and not on industry k's suppliers. That is, two industries with the same demand-chain will have the same $\tilde{\beta}$ regardless of their own supply chains (I formally show this in the next section).

An analogous demand-side centrality measure can also be defined.

Definition 2.3.2. The consumer centrality is

$$\tilde{\alpha} = \Psi_d \alpha,$$

where

$$\Psi_d = (I - \mu^{1-\sigma} \tilde{M}^{\sigma-1} \Omega)^{-1} \mu^{1-\sigma} \tilde{M}^{\sigma-1} = \sum_{n=0}^{\infty} \left(\mu^{1-\sigma} \tilde{M}^{\sigma-1} \Omega \right)^n \mu^{1-\sigma} \tilde{M}^{\sigma-1}.$$

This is the flip-side to the supply-side centrality measure. It captures how frequently an industry appears in demand-chains. Whereas $\tilde{\beta}_k$ depended only on who bought from k, the consumer centrality $\tilde{\alpha}_k$ depends only on who k buys from. Baqaee (2014c) shows that, in a model without an extensive margin, $\tilde{\alpha}$ captures the response of output to demand shocks. As the following lemma shows, consumer centrality $\tilde{\alpha}$ is a transformation of an industry's price:

Lemma 2.3.2. *In equilibrium,*

$$\left(\frac{p_i}{w} \right)^{1-\sigma} = \tilde{\alpha}_i,$$

where p_i is industry i's price index and w is the nominal wage.

Since prices are collinear with marginal costs, this means that $\tilde{\alpha}$ is a measure of marginal costs. This makes clear why $\tilde{\alpha}_k$ depends on industry k's supply-chain, since the it is suppliers and not consumers, who contribute to marginal costs.

In defining the consumer centrality $\tilde{\alpha}$ and proving lemma 2.3.2, I have not made any reference to the productivity shocks z. As alluded to earlier, this is an abuse of notation, because I treat α_k as already incorporating the productivity shock. This is because when $\sigma \neq 1$, a productivity shock z_k to industry k is equivalent to changing that industry's labor intensity from α_k to $\alpha_k z_k^{\sigma-1}$. Therefore, as long as $\sigma \neq 1$, we can think of α_k as including

both the productivity shock and the labor intensity. When $\sigma = 1$, the consumer centrality is trivially equal to a vector of ones regardless of the productivity shocks.

A key result, that delivers much of the intuition of the results in the paper, is the following characterization of active firms' profit functions in terms of supplier and consumer centrality measures:

Theorem 2.3.3. *The payoffs of firm i in industry k are equal to*

$$\pi(k,i) = \underbrace{\frac{1}{\varepsilon_k M_k}}_{\text{industrial competition}} \times \underbrace{P_c^\sigma C w^{1-\sigma}}_{\text{GE terms}} \times \underbrace{\tilde{\beta}_k}_{\text{supplier centrality}} \times \underbrace{\tilde{\alpha}_k}_{\text{consumer centrality}} - \underbrace{w f_k}_{\text{fixed cost}}.$$

Note that $M_k \pi(k,i)$ gives industry k's profits Π_k.

Without loss of generality, we can set the nominal wage $w = 1$. The expression in theorem 2.3.3 tells us that the profits of a firm are determined by a few intuitive key statistics. The product of $\tilde{\beta}_k$ and $\tilde{\alpha}_k$, which are the supply-side and demand-side centrality of the industry give us an industry's share of sales. The term $P_c^\sigma C$ is an economy-wide shifter of all industry's profits, akin to aggregate demand. The division by ε_k converts an industry's sales into profits since the within-industry elasticity of substitution determines mark-ups. Dividing gross industrial profits by the mass firms M_k in that industry turns gross industrial profits into gross firm-level profits. Finally, we arrive at a firm's profits by subtracting the fixed costs of entry from its gross profits.

Path Example

Before moving on to an analysis of the equilibrium, first let us demonstrate the intuition so far using a simple example of a production chain, shown in figure 2.1.

Begin by computing the supplier-centrality for node k in this chain:

$$\tilde{\beta}_k = \beta'(I - \tilde{M}^{\sigma-1}\mu^{-\sigma}\Omega)^{-1}e_k,$$

$$= \prod_{i=0}^{k-1}(1-\alpha_i)M_i^{\frac{\sigma-1}{\varepsilon_i-1}}\left(\frac{\varepsilon_i}{\varepsilon_i-1}\right)^{-\sigma}.$$

First, note that $\tilde{\beta}_k$ is a product. Therefore, if any industry $i < k$ disappears $M_i = 0$, then the
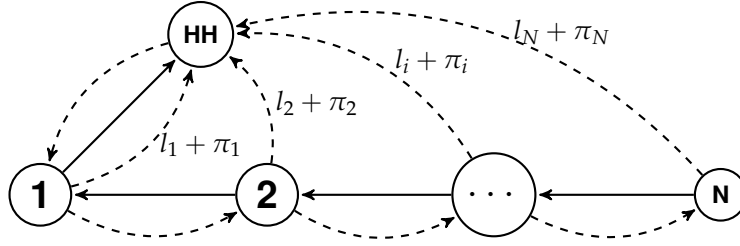
Figure 2.1: *A production path example. The solid arrows represent the flow of goods and services, and the dashed arrows indicate the flow of money. The household HH buys from industry 1 who buys from industry 2 and so on. Each industry in turn pays labor income and rebates profits to the household.*

supplier centrality of industry k drops to zero. This intuitive, since if a downstream industry collapses, that cuts all upstream industries off from any demand. The centrality of k as a supplier is increasing in the strength of its downstream connections $(1 - \alpha)$ and decreasing in the size of downstream markups $\varepsilon_i / (\varepsilon_i - 1)$. The latter represents double-marginalization in this economy. Note that as long as the elasticity of substitution is greater than 1, an industry's supplier centrality is increasing in the mass of downstream industries. Intuitively, when the elasticity of substitution is greater than one, more industries downstream attract more demand from the household. Lastly, observe that kth industry's supplier centrality depends purely on its customers, customers' customers, and so on. It does not depend on its suppliers. This is a general property of the supplier centrality, and we shall prove it in section 2.4.

Now, compute the consumer-centrality for node k in this chain:

$$
\tilde{\alpha}_k = e_k'(I - \tilde{M}^{\sigma-1}\mu^{1-\sigma}\Omega)^{-1}\tilde{M}^{\sigma-1}\mu^{1-\sigma}\alpha,
$$

$$
= \sum_{j=k+1}^{N} \left(\prod_{i=k}^{j-1}(1 - \alpha_i)M_i^{\frac{\sigma-1}{\varepsilon_i-1}} \left(\frac{\varepsilon_i}{\varepsilon_i - 1} \right)^{1-\sigma} \right) \alpha_j + \alpha_k M_k^{\frac{\sigma-1}{\varepsilon_k-1}} \left(\frac{\varepsilon_k}{\varepsilon_k - 1} \right)^{1-\sigma}.
$$

The consumer centrality is slightly more complex. It is an arithmetico-geometric. The intuition is that even if an upstream supplier i for industry k disappears, industry k still has access to all suppliers j where $j < i$. But any supplier $j > i$ also drops out. This gives rise to the arithmetic series, where each term is a product. Once again, the consumer centrality is increasing in the strength of the connection. And as long as the elasticity of substitution is

greater than one, consumer centrality is decreasing in markups and increasing in the mass of upstream firms.

When the elasticity of substitution is less than one, consumer centrality is increasing in upstream markups. This sounds perverse until we recall that proposition 2.3.1 implies that

$$\tilde{\alpha}_k = \left(\frac{p_k}{w}\right)^{1-\sigma}.$$

Therefore, when the elasticity of substitution is less than one, higher consumer centralities indicate higher prices, not lower prices. Therefore, it is intuitive that in this case, higher upstream markups correspond to higher consumer centrality, since this corresponds to higher prices. Finally, note that, consumer centrality of a firm depends only on who it buys from, and not on who it sells to. Once again, this is a general property of consumer centrality that we discuss in more detail later.

Analysis of the Full Equilibrium

Our analysis of the equilibrium of this model, and its responses to the shocks, will proceed in parts. First, we fix the mass of firms in each industry to isolate the intensive margin responses. Firms may make nonzero profits in equilibrium, and the model is a generalization of the models in Acemoglu *et al.* (2012) or Long and Plosser (1983). With the entry margin shut down, I show that the relevant notion of a firm's systemic importance is exogenous and is approximated by its share of sales. Furthermore, a firm's importance depends only on that firm's role as a supplier of goods. A firm's role as a consumer of inputs is irrelevant. Lastly, I show productivity shocks can never be amplified in this class of models. This is because a shock is always largest at its source – as it travels through the network, the shock is attenuated and the aggregate impact of the shock is a convergent geometric sum. The more influential the firm, the slower the decay. However, there are no cases where the shock actually gets bigger as it travels through the network.

After analyzing the model without entry, I allow free-entry but consider the limit of the model when firms have no mass, $\|\Delta\| \to 0$. In this case, systemic importance is endogenous

and can change depending on economic conditions. Furthermore, importance is no longer well-approximated by size. Finally, a firm's systemic importance depends not just on its importance as a supplier of inputs but also as a consumer of inputs. In this sense, both the in-degrees and out-degrees matter. However, in this limit, the model still cannot amplify shocks. As in the model without an extensive margin, shocks decay geometrically as they travel from the source.

Finally, I consider the case when $\|\Delta\| > 0$. This model features amplification mechanisms, and cascading failures, that are not present in the continuous model. This model can be thought of as an interpolation between discontinuous threshold models and continuous pulse models. To solve for the discontinuous model's equilibrium, we need to solve a computationally intractable integer programming problem. Most of the time, this model's behavior can be approximated by its continuous limit. However, in certain cases, its behavior is very different. This is because in some cases, a small shock can cause a systemically important firm to discontinuously exit. This can snowball into further failures and build on itself. I prove an inapproximability result that gives conditions for when the continuous model behaves very differently to the discrete model, along with some informative examples.

2.4   No Extensive Margin

In this section, we consider the behavior of this model when the extensive margin is shut down. This assumption means that this model is a generalization of the canonical input-output model of Long and Plosser (1983). I show that in this class of models, an industry's influence on other industries and on the aggregate economy is exogenously determined, and that it is only dependent on how important the industry is as a supplier of inputs. Furthermore, I show that a firm's size is an important determinant of a firm's influence. Last, I show that this class of models is incapable of local amplification. A shock to a firm or industry will always decay as it travels from the source to its neighbors. This makes clear why systemically important industries must be large ones, since the decaying propagation of shocks means that a shock's aggregate impact is necessarily limited by its immediate

impact.

2.4.1 Perfect Competition

Standard models that lack an entry margin are typically perfectly competitive, with no fixed costs, so that the representative firm in each industry makes zero profits. To get to this benchmark, let $f_i = 0$ for every industry so that there are no fixed costs; and let $\varepsilon_i \to \infty$ so that all industries are perfectly competitive and firms have no market power. Then, due to constant returns to scale, the size of firms in each industry is indeterminate. Without loss of generality, we can fix $M = I$ so that there is a unit mass of firms in each industry. Technically, there may be entry or exit, but since firms in each industry are completely homogenous, entry and exit is observationally equivalent to firms getting larger or smaller. That is, the extensive and intensive margins are indistinguishable.

Note that in this special case $\Psi_s = \Psi_d = (I - \Omega)^{-1}$. To simplify notation, let $\Psi = (I - \Omega)^{-1}$. In subsection 2.4.2, I show how allowing for fixed costs and monopolistic competition, but not allowing entry, would affect the results of this section.

First, let us consider real GDP C as a function of productivity shocks z for a Cobb-Douglas economy E.

Proposition 2.4.1 (Productivity shock). *Let the elasticity of substitution across industries be equal to one, then*

$$\log(C(z|E)) = \tilde{\beta}'(\alpha \circ \log(z)) = \sum_{k=1}^{N} \frac{wl_k}{GDP} \log(z_k).$$

That is, the network-structure Ω which has N^2 parameters is summarized by N sufficient statistics. These sufficient statistics are each industry's expenditures on labor as a share of GDP.

This proposition is a slight generalization of results in Acemoglu *et al.* (2012). If we assume that each industry's labor share is constant, so that $\alpha_i = \alpha$ for all i, and household expenditure shares are uniform so that $\beta_i = 1/N$, then we exactly recover the result in Acemoglu *et al.* (2012), which tells us that the effect of a productivity shock on real GDP depends on share of sales.

Corollary (Productivity shock). *Let the elasticity of substitution across industries be equal to one, and all industries have the same labor share α, then*

$$\log(C(z|E)) = \alpha \sum_{k=1}^{N} \frac{p_k y_k}{GDP} \log(z_k).$$

That is, the network-structure Ω which has N^2 parameters is summarized by N sufficient statistics. These sufficient statistics are each industry's share of sales.

These propositions imply that when $\sigma = 1$, the supplier centrality $\tilde{\beta}$ is a sufficient statistic for translating productivity shocks into real GDP. In particular, the aggregate impact of a vector of shocks depends on the supplier centrality weighted average of the productivity shocks. As lemma 2.3.1 shows, $\tilde{\beta}_k$ is equal to industry k's share of sales. Therefore, the bigger the industry in equilibrium, the larger the impact of its productivity shocks on GDP. The exact nature of the network-structure is irrelevant since an industry i could be big because it sells a lot to the household (large β_i) or because it supplies many other firms (the ith column of Ψ is large).

These intuitions also carry over to the case where $\sigma \neq 1$. When the elasticity of substitution across industries $\sigma \neq 1$, productivity shocks to labor are isomorphic to changes in the labor share parameters α. In particular, a productivity shock z_k to industry k is equivalent to changing industry k's labor share parameter from α_k to $\alpha_k z_k^{\sigma-1}$. Therefore, we simply investigate changes to α.

Proposition 2.4.2 (Productivity shock). *When the elasticity of substitution across industries is not equal to one, then*

$$C(\alpha|E) = \left(\tilde{\beta}'\alpha\right)^{\frac{1}{\sigma-1}}.$$

That is, the network-structure Ω, which has N^2 parameters, is summarized by N sufficient statistics: $\tilde{\beta}$.

Outside of the Cobb-Douglas case, the supply-side centrality $\tilde{\beta}' = \beta'\Psi_s$ is no longer an industry's share of sales. Lemma 2.3.1 shows that it is still related to an industry's share of

sales however. In fact, even though $\tilde{\beta}'$ is not equal to share of sales globally, it is still closely related to share of sales.

Proposition 2.4.3. *Around the steady-state, where $z_k = 1$ for all k, we have that*

$$\frac{\tilde{\beta}_k}{\tilde{\beta}_l} = \frac{p_k y_k}{p_l y_l},$$

so although $\tilde{\beta}$ is not equal to share of sales everywhere, at the steady-state, it does reflect an industry's size.

This is a consequence of that fact that in the steady-state with no markups or productivity shocks, all firms have the same price. And, as long as all firms have roughly the same price, we can interpret $\tilde{\beta}$ as being roughly equal to share of sales. Away from this steady state, $\tilde{\beta}$ is still the relevant sufficient statistic, although it is no longer equal to the share of sales.

Even though $\tilde{\beta}$ no longer corresponds to an industry's size everywhere, it is still exogenous, and depends solely on the amount of household demand that reaches the industry, whether directly through retail sales, or indirectly through other industries. As before, the exact nature of the network-structure is irrelevant since an industry i could be big because it sells a lot to the household (large β_i) or because it supplies many other firms (the ith column of Ψ_s is large).

Furthermore, not only is this statistic exogenous with respect to productivity shocks, but it also depends only on the industry's role as a supplier to other agents. If two industries sell the same amount to the same group of industries, they will have the same $\tilde{\beta}$ regardless of who they themselves are buying from. We can formally state this intuition as follows.

Proposition 2.4.4. *Consider two industries k and l such that*

$$\omega_{jk} = \omega_{jl}, \quad (j = 1, \dots, N), \quad and \quad \beta_k = \beta_l,$$

then $\tilde{\beta}_k = \tilde{\beta}_l$. In other words, if two industries have the same immediate customer base, their supplier-centralities are the same.

In light of these propositions, we can state the following result.

Theorem 2.4.5. *For every household share parameter vector β and input-output share parameter matrix Ω, there exists a different economy with household share parameter vector $\tilde{\beta}$ and degenerate input-output share parameter matrix $\mathbf{0}$, such that*

$$C(\alpha|\beta,\Omega) = C(\alpha|\tilde{\beta},\mathbf{0}).$$

This result means that the existence of input-output connections is isomorphic to a change in household share parameters. This underlies the earlier claim that, it is not the interconnections themselves, but how intensively the household ultimately consumes goods from each industry that determines the model's equilibrium responses to shocks.

Local Amplification

One reason why a firm's size matters for its impact on aggregate outcomes is the continuous nature of shock propagation in this class of models. In input-output models, a shock to industry k has the largest impact on industry k. The shock influences k's neighbors, but the effect of the shock is attenuated by the weakness of its connections. The effect of the shock on k's neighbors' neighbors is further attenuated by the weakness of its connections's connections. The shock decays geometrically with a decay rate controlled by the input-output matrix Ω. This implies that the shock has its greatest impact at its source.

To show this, first we need the following technical lemma, which is a novel result in linear algebra.

Lemma 2.4.6. *Let A be a non-negative $N \times N$ matrix whose rows sum to one. Let a, b, c be $N \times 1$ vectors in the unit cube with $c = a + b$. Let $B = (I - \mathrm{diag}(1 - c)A)^{-1}$. Then*

$$\frac{e_i' B e_i}{e_i' B a} \geq \frac{e_j' B e_i}{e_j' B a} \quad (i \neq j),$$

where e_i is the ith standard basis vector. When a is strictly positive, then the inequality is strict.

The following proposition shows that the equilibrium effect of a productivity shock z_i to industry i's sales, in percentage change terms, is greatest for industry i. Furthermore, the effect on other industries decays geometrically according to the sum of geometric matrix

series Ψ. Since sales are proportional to gross profits, the proposition also applies to gross profits.

Proposition 2.4.7. *Consider the semi-elasticity of firm i's sales relative to firm j's sales with respect to a shock to firm i's productivity around the steady state $z = 1$:*

$$\frac{d \log(\text{sales}_i)}{dz_i^{\sigma-1}} - \frac{d \log(\text{sales}_j)}{dz_i^{\sigma-1}} = \alpha_i \left(\Psi_{ii} - \Psi_{ji} \right) > 0.$$

This proposition formalizes the intuition that a shock has to decay as it travels, and the input-output connections Ω control the rate of decay. The rate of decay can be slow enough to overturn the law of large numbers, as shown by Acemoglu *et al.* (2012). However, even though the network can slow the decay, proposition 2.4.7 shows that shocks cannot be amplified as they travel out from their source. In order for shocks to industry k to have a big impact, industry k must have strong connections to the household, or strong connections to other industries that have strong connections to the household, and so on. But having strong connections to the household implies that the industry has to be large in equilibrium. The lack of local amplification means that size and influence are closely linked in this class of models. This, and the earlier results in this section, show that as long as firms are small, and they do not supply large fractions of the economy, then shocks to them cannot have significant aggregate effects.

2.4.2 No Extensive-Margin and Monopolistic Competition

So far, we have assumed that there are no fixed costs of entry and all industries are perfectly competitive. In this subsection, I consider the case where there is no entry but firms have some market power and positive fixed costs. This makes it easier to understand how adding the extensive margin will affect the results, since the extensive margin will only matter once we have fixed costs and product differentiation. I state a few key propositions in this subsection to show how the intuition of the previous section will carry over to this case.

First, let us consider the special case where the elasticity of substitution across industries $\sigma = 1$, which is the Cobb-Douglas case.

Proposition 2.4.8 (Productivity shock). *Let the elasticity of substitution across industries be equal to one, then*

$$\log(C(z|E)) = \text{const} + \beta'(I - \Omega)^{-1}(\alpha \circ \log(z)).$$

That is, the network-structure Ω, which has N^2 parameters, is summarized by N sufficient statistics. These sufficient statistics are $\beta'(I - \Omega)^{-1} = \beta'\Psi_d$.

The sufficient statistic $\beta'\Psi_d$ is closely related to the sales shares $\tilde{\beta} = \beta'\Psi_s$. To see this, note that

$$\tilde{\beta} = \beta' \sum_{t=0}^{\infty} \left(\mu^{-1}\Omega\right)^t,$$

where μ is the diagonal matrix of mark-ups. Whereas,

$$\beta'\Psi_d = \beta' \sum_{t=0}^{\infty} (\Omega)^t.$$

Therefore, the relevant statistics $\beta'\Psi_s$ are the sales shares that would prevail if there were no mark-ups. Intuitively, they are still an exogenous supplier-centrality measure.

Proposition 2.4.9 (Productivity shock). *When the elasticity of substitution across industries is not equal to one and there is no entry or exit, then*

$$C(\alpha|E) = \frac{1 - M'f}{1 - \text{diag}(\varepsilon)^{-1}\tilde{\beta}'\tilde{\alpha}/\beta'\tilde{\alpha}} \left(\beta'\tilde{\alpha}\right)^{\frac{1}{\sigma-1}},$$

$$= f(\beta'\Psi_s\Psi_d\alpha, \beta'\Psi_d\alpha) = f(\tilde{\beta}'\Psi_d\alpha, \beta'\Psi_d\alpha).$$

That is, the network-structure Ω, which has N^2 parameters, is summarized by $2N$ sufficient statistics: $\beta'\Psi_d$ and $\tilde{\beta}'\Psi_d$.

$\tilde{\beta}$ is just the supplier centrality we have dealt with previously. It is an exogenous supplier centrality, and depends solely on the amount of household demand that reaches the industry, whether directly through retail sales, or indirectly through other industries. As before, the exact nature of the network-structure is irrelevant since an industry could be big because it sells a lot to the household (large β) or because it supplies many other firms (large $\Psi_d e_i$).

The Cobb-Douglas case constitutes a very special knife-edge scenario where, because expenditure shares are exogenous, the length of supply chains do not matter. Once we

allow for market power $\varepsilon_k < \infty$, and endogenous expenditure shares $\sigma \neq 1$, each time funds flow from one industry to another, they are attenuated by monopoly profits. Longer chains accrue larger mark-ups, and this changes the expenditure shares. Due to this effect, we need a measure that not only takes the intensity of supply chains into account, but also their length. This is the role that the second sufficient statistic $\beta'\Psi_s\Psi_d = \tilde{\beta}'\Psi_d$ plays.

To see this, suppose that we have a unit mass of firms in each industry $M = I$ and that all mark-ups are zero. In this extreme case, $\Psi_d = \Psi_s$. So, we can interpret

$$\Psi_d\Psi_s = (I - \Omega)^{-2} = I + 2\Omega + 3\Omega^2 + 4\Omega^3 + \dots.$$

This calculation gives some intuition for why $\beta'\Psi_s\Psi_d$ is a supply-side centrality measure that controls for the *length* of the supplier relationship as well as its intensity.

Local Amplification

The following proposition shows that the equilibrium effect of a productivity shock z_i to industry i's sales, in percentage change terms, is greatest for industry i. Furthermore, the effect on other industries decays geometrically according to the sum of geometric matrix series Ψ^s. Since sales are proportional to gross profits, the proposition also applies to gross profits.

Proposition 2.4.10. *For any economy E, if we fix the mass of firms in each industry, we have*

$$\frac{d\log(\text{sales}_i)}{dz_i} - \frac{d\log(\text{sales}_j)}{dz_i} = \frac{\Psi^s_{ii}\alpha_i}{\tilde{\alpha}_i} - \frac{\Psi^s_{ji}\alpha_i}{\tilde{\alpha}_j} \geq 0.$$

2.5 Extensive Margin

Now that we understand the properties of the model without the extensive margin, let us consider how entry and exit changes the model's properties. In this case, the network structure is endogenous since the number of firms in each industry is endogenous. This means that our notions of centrality $\tilde{\beta}$ and $\tilde{\alpha}$ are determined in equilibrium, and they change in response to shocks. It is no longer the case that an industry's influence is exogenous,

nor is it the case that only its role as a supplier matters. In other words, an industry's influence depends both on its in-degrees and its out-degrees, and its degrees are determined in equilibrium depending on the mass of firms in other industries. Two firms with identical roles as suppliers can have markedly different effects on the equilibrium depending on their roles as consumers.

Our analysis will proceed in steps. First, we consider the limiting case of the model where all firms have zero mass. This means that the mass of firms in each industry will continuously adjust to ensure that all active firms make zero profits in equilibrium. This will allow for sharp analytical results about the equilibrium response to marginal shocks.

Once we have characterized the properties of the continuous limit, we consider the case where firms can have positive mass $\Delta > 0$. In this case, the model inherits some of the cascading and amplification properties of linear threshold models like Schelling (1971), but it also retains the linear geometric structure of Leontief (1936). Most of the time, the model can be expected to behave like its continuous limit; however, occasionally, this approximation can dramatically break down. The intuition is that each industry can support an integer number of firms. If firms are sufficiently small and there are many of them in an industry, then a shock to the industry will continuously perturb the mass of firms in that industry. However, if there are few firms in an industry, then shocks below a certain threshold are attenuated and will not result in any change to the number of active firms. However if shocks are bigger than that threshold, they will cause a firm to discontinuously exit or enter. This adds an additional impulse to the original shock, which will now travel to the neighbors of the affected industry with more force. This process can feed on itself as a small firm's failure can trigger a chain reaction that results in a large of mass of firms exiting.

A full characterization of the equilibrium of the model with lumpy firms is not possible. A natural solution is to use the continuous model to approximate the behavior of the model with lumpy firms. However, I show conditions under which the continuous limit provides a bad approximation to the discontinuous model. I show that the network can make firm's payoffs non-monotonic in each other's entry decisions. So, there are regions where entry

decisions are strategic complements and and regions where they are substitutes. I show that when a shock puts us close to this intermediate region, the approximation error will explode.

To foreshadow the formal model, consider the extreme toy example in figure 2.2. The household HH is served by three industries: GM, F, and L. Two of these industries share a common supplier D. To make the intuition transparent, suppose that each industry consists of one firm. Recall that theorem 2.3.3 implies that in order for a firm to remain profitable, its sales $\tilde{\beta}_i \tilde{\alpha}_i P_c^\sigma C$ have to be greater than an exogenous threshold $\varepsilon_i f_i$. Each term of $\tilde{\beta}$, $\tilde{\alpha}$, and $P_c^\sigma C$ can respond to a shock. Therefore, the shock to a firm can travel via network connections $\tilde{\beta}$ and $\tilde{\alpha}$ or it can travel through household demand $P_c^\sigma C$.

The example in figure 2.2 demonstrates. In panel 2.2a, I show a negative shock to L's fixed cost so that L is forced to exit. Then, through the effect of L on $P_c^\sigma C$, it can be the case that GM is forced to exit, and this is shown in panel 2.2b. GM exiting will cause D's supplier centrality $\tilde{\beta}_D$ to fall, and so D can be forced out, shown in panel 2.2c. Finally, once D is gone, this can cause F's consumer centrality $\tilde{\alpha}$ to drop, which can cause them to also exit. Of course, this is not a realistic calibration of the model, and the assumptions of monopolistic competition are hard to justify with single-firm industries. However, this stark example does illustrate the forces operating in the model. In a canonical input-output model, like the one in section 2.4.2, a change to the fixed costs of L would have no network effects and simply lower real GDP C by the amount of the fixed cost.

This example shows that the model is capable of formalizing the intuition for why a bail-out of GM may have been crucial. As pointed out by Goolsbee and Krueger (2015), in this scenario, it might be "essential to rescue GM to prevent an uncontrolled bankruptcy and the failure of countless suppliers, with potentially systemic effects that could sink the entire auto industry."

(a) *L exists due to increased fixed costs.*

(b) *Effect on GM through household.*

(c) *Effect on D through GM.*

(d) *Effect on F through D.*
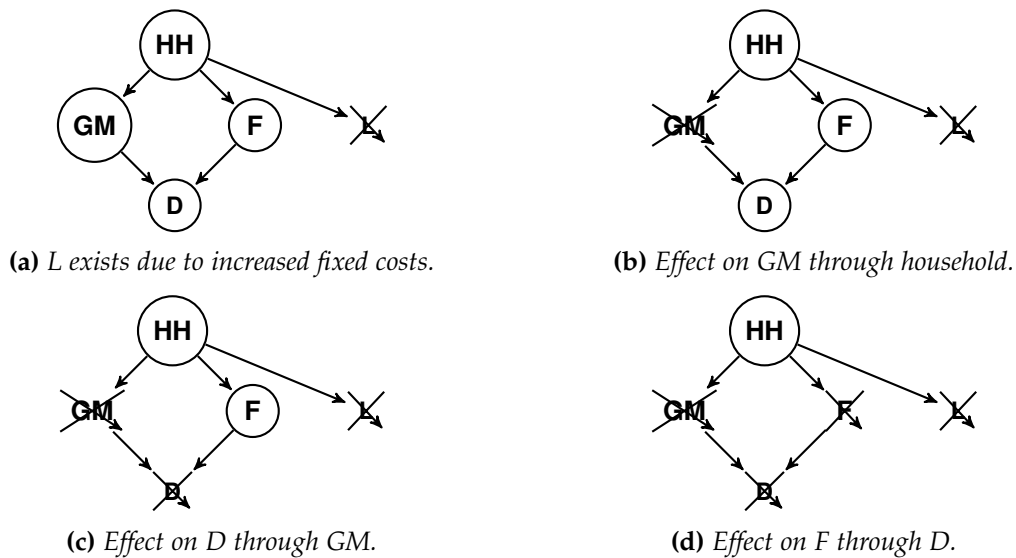
Figure 2.2: *A toy representation of an economy where the direction of arrows denote the flow of funds. The node HH denotes the household. This shows how the shock can travel in both directions on the extensive margin.*

2.5.1 Massless Limit

To begin with, let us first consider the equilibrium of the model in the limit where $\Delta \to 0$. Computationally, this corresponds to the case where the mass of firms M_k in industry k will adjust so that firms in industry k make exactly zero profits. This can be seen by taking the limit of the expression in definition 2.2.2 as $\Delta \to 0$. The primary motivation for looking at this limit is analytical tractability. By considering massless firms, we can glean useful intuition about the marginal effects of shocks on the equilibrium.

In this subsection, I show that once we allow for firm entry and exit, the model's equilibrium responses can no longer be characterized in terms of sufficient statistics. The intuition is simple: the centrality measures $\tilde{\beta}$ and $\tilde{\alpha}$ are functions of the masses of firms in each industry. Since the mass of firms responds to shocks, the centrality measures also respond to shocks. Furthermore, I show that an industry's impact on supplier centralities depends on the industry's own supplier centrality *and* its role as a consumer of inputs. On the other hand, an industry's impact on consumer centralities depends on the industry's own consumer centrality *and* its role as a supplier of inputs. So, although the supplier (consumer)

centrality only depends on out-degrees (in-degrees), the way supplier (consumer) centrality changes in response to more entry depends on both the out-degrees and the in-degrees.

First to get intuition for the general model, let's consider a very special case – the case where the elasticity of substitution is equal to one. In this case, all expenditure shares are exogenous. In this Cobb-Douglas case, we have the following result.

Proposition 2.5.1. *When the elasticity of substitution is equal to one*

$$C(z, f) = \tilde{\beta}' \left(\alpha \circ \log(z) - \frac{1}{\varepsilon - 1} \circ \log(f) \right) + \text{const.}$$

Therefore, the network-structure is summarized by N sufficient statistics $\tilde{\beta}$. Furthermore, in equilibrium

$$\tilde{\beta}_i = \frac{p_i y_i}{GDP}.$$

This shows that with Cobb-Douglas, the share of sales, which is exogenous, is again a sufficient statistic. Once again, the details and the complexity of the network are irrelevant, once we know each industry's share of sales. This is a knife-edge case. Once we deviate from Cobb-Douglas, expenditure shares respond to relative prices, and centrality measures become endogenous. This special case shows that the mechanism for our upcoming results depends crucially on the fact that expenditure shares respond to relative prices.

Out-of-Equilibrium Effect

To get intuition for the model's properties, let's consider the following out-of-equilibrium comparative statics: how does industry k's supplier centrality change when there is entry in industry $i \neq k$? This comparative static holds fixed the mass of firms in all industries except industry i. In equilibrium, of course, the masses in other industries would respond, but it helps our understanding if we first isolate the partial equilibrium effect. If we had lags in entry and exit, these partial equilibrium results would be relevant for understanding short-run effects. Once we have characterized the out-of-equilibrium effects, we turn our attention to the equilibrium responses of the model.

Lemma 2.5.2. *The derivative of $\tilde{\beta}_k$ with respect to a percentage change in the mass of firms in industry i, holding fixed the mass of firms in all other industries is given by*

$$\frac{1}{M_i}\frac{\partial \tilde{\beta}_k}{\partial M_i} = \left(\frac{\sigma - 1}{\varepsilon_i - 1}\right)\tilde{\beta}_i(\Psi_{ik}^s - 1(i = k)).$$

This expression is very intuitive. The impact to industry k's centrality as a supplier depends on i's importance as a supplier, and on how much i buys from k (whether directly or indirectly). So big effects are felt if i is a key supplier *and* i buys a lot from k. Note that if the elasticity of substitution $\sigma = 1$, then these derivatives are identically zero, which explains the neutrality result in proposition 2.5.1.

For ease of notation, let

$$\frac{\partial \tilde{\beta}}{\partial \log(M\mathbf{1})} = \Psi_1,$$

then the result in lemma 2.5.2 can be written in matrix notation as

$$\Psi_1' = \text{diag}(\tilde{\beta})\text{diag}\left(\frac{\sigma - 1}{\varepsilon - 1}\right)(\Psi_s - I).$$

Similar results apply to the consumer centrality measure.

Lemma 2.5.3. *The derivative of $\tilde{\alpha}_k$ with respect to a percentage change in the mass of firms in industry i, holding fixed the mass of firms in all other industries is given by*

$$\frac{\partial \tilde{\alpha}_k}{\partial M_i} = M_i^{1-\sigma}\left(\frac{\sigma - 1}{\varepsilon_i - 1}\right)\left(\frac{\varepsilon_i}{\varepsilon_i - 1}\right)^{\sigma - 1}\tilde{\alpha}_i\Psi_{ki}^d,$$

This expression is the demand-side analogue to lemma 2.5.2. The impact to industry k's centrality as a consumer depends on on i's importance as a consumer, and on how much i sells to k (whether directly or indirectly). So big effects are felt if i is a key consumer *and* i sells a lot of k. Once again, the impact to industry k's consumer centrality depends on industry i's suppliers and i's customers.

For ease of notation, let

$$\frac{\partial \tilde{\alpha}}{\partial \log(M\mathbf{1})} = \Psi_2,$$

84

then from lemma 2.5.3 we can write Ψ_2 in matrix notation,

$$\Psi_2 = \Psi_d \text{diag}(\tilde{\alpha}) \text{diag}\left(\frac{\sigma - 1}{\varepsilon - 1}\right) \mu^{\sigma - 1} \text{diag}(M)^{1 - \sigma}.$$

Path Example

Before moving on to an analysis of the equilibrium, first let us demonstrate the intuition of our partial equilibrium results using a simple example of a production chain depicted in figure 2.3.
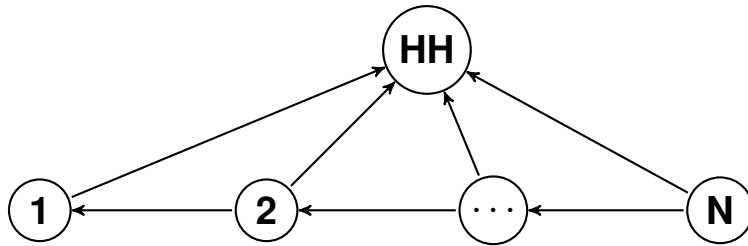
Figure 2.3: *The arrows represent the flow of goods and services. Let $\omega_{k-1,k}$ be constant for all k, and β_k, α_k and ε_k be constant for all k.*

In this example, each industry k sells some goods directly to the household, and some goods to the industry below it $k - 1$. For ease of exposition, consider $\tilde{\beta}$ and $\tilde{\alpha}$ at the point where all industries have a unit mass of firms. In this example, $\tilde{\beta}_k$ is increasing in k and $\tilde{\alpha}_k$ is decreasing in k. That is, industry N is the most central supplier and least central consumer, and industry 1 is the most central consumer and least central supplier.

Now consider changing the mass of firms in industry N. Industry N is the most central supplier so $\tilde{\beta}_N > \tilde{\beta}_k$ for $k \neq N$. However,

$$\frac{\partial \tilde{\beta}_k}{\partial M_N} = 0,$$

because industry N buys from no other industries. Therefore, its impact on supplier centralities is zero despite being the most central supplier.

Now consider changing the mass of firms in industry 1. Industry 1 is the most central

consumer so $\tilde{\alpha}_1 > \tilde{\alpha}_k$ for $k \neq 1$. However,

$$\frac{\partial \tilde{\alpha}_k}{\partial M_1} = 0 \quad (k > 1),$$

because industry 1 sells to no other industries. Therefore, its impact on consumer centralities is zero despite being the most central consumer.

This simple example illustrates why both the in-degrees and the out-degrees will matter for how the centrality measures will respond to a change in the mass of firms in each industry. Being a central supplier does not mean that entry or exit in your industry will have any effects on the supplier centralities of other industries. Similarly, being a central consumer does not imply that entry or exist in your industry will have any effects on the consumer centralities of other industries. To be influential, a firm must be central both as a consumer and as a supplier.

Equilibrium Impact of Shocks

So far, we have been focusing on out-of-equilibrium results. However, using lemmas 2.5.2 and 2.5.3, we can also analyse the general equilibrium impact of a productivity shock to an industry.

Proposition 2.5.4. *The derivative of the equilibrium mass of firms in each industry M with respect to a labor productivity shock in industry k is given by*

$$\left(\frac{\mathrm{d} \log(M)}{\mathrm{d} z_k} \right) = (I - \mathrm{diag}(\tilde{\beta})^{-1} \Psi_1 - \mathrm{diag}(\tilde{\alpha})^{-1} \Psi_2)^{-1} \left(\frac{\mathbf{1} \mathrm{d} P_c^\sigma C / \mathrm{d} z_k}{P_c^\sigma C} + \mathrm{diag}(\tilde{\alpha})^{-1} \Psi_d e_k \right).$$

So we see that in equilibrium, a productivity shock to industry k will first affect the masses in all other industries through its effect on the aggregate objects $P_c^\sigma C$ and through its effect on the marginal costs of anyone who buys from k. However, the initial change in masses results in the supplier and consumer centralities of all industries to change (captured by Ψ_1 and Ψ_2). This change in centrality measures, in turn, causes the masses to adjust again, and this changes the centralities again, and so on ad infinitum. This gives rise to a geometric series and the equilibrium effect is the sum of this geometric series.

For ease of notation, let

$$\Lambda = \frac{d\log(\tilde{\beta}) + \log(\tilde{\alpha})}{d\log(M)} = \text{diag}(\tilde{\beta})^{-1}\Psi_1 + \text{diag}(\tilde{\alpha})^{-1}\Psi_2.$$

Note that $\log(\tilde{\beta}) + \log(\tilde{\alpha})$ is proportional to each industry's share of sales. Therefore, Λ is the elasticity of industry sizes relative to the mass of firms in each industry. Now we can see the intuition of 2.5.4 most clearly by expressing the derivative as

$$\left(\frac{d\log(M)}{dz_k} \right) = \sum_{t=0}^{\infty} \Lambda^t \left(\frac{1 d P_c^g C/dz_k}{P_c^g C} + \text{diag}(\tilde{\alpha})^{-1}\Psi_s e_k \right)$$

$$= \underbrace{\sum_{t=0}^{\infty} \Lambda^t \left(\frac{1 d P_c^g C/dz_k}{P_c^g C} \right)}_{\text{GE effect}} + \underbrace{\sum_{t=0}^{\infty} \Lambda^t \left(\text{diag}(\tilde{\alpha})^{-1}\Psi_s e_k \right)}_{\text{network-structure effect}}.$$

The term

$$\text{diag}(\tilde{\alpha})^{-1}\Psi_s e_k$$

is the intensive-margin effect of a productivity shock to industry k. If industry k is more productive, that increases the productivity of any industry that buys inputs from k. The degree to which an industry downstream from k is affected depends on how intensively it uses inputs from k (directly or indirectly). This is precisely the effect of a shock when the extensive margin is shut down as shown in proposition 2.4.10. Therefore, we can think of this as the traditional input-output effect.

With the extensive margin, the initial change in $\tilde{\alpha}$ also causes the mass of firms to change. This change in the mass of firms causes further changes in the masses of firms. The cumulative effect on the equilibrium of these changes is captured by the "network-structure effect" term. The shock also has an effect on the general price level and real GDP, and the second "GE effect" captures this general equilibrium effect. It is a simple matter to show that in equilibrium, $dP_c^g C/dz_k$ is just a weighted average of the network-structure effects.

As alluded to earlier, this complexity depends on the endogeneity of expenditure shares. For the Cobb-Douglas case, these formulas lose their interesting properties, and we only have the aggregate general equilibrium effects.

Proposition 2.5.5. *If Cobb-Douglas, then the matrix is diagonal, and we have*

$$\frac{\mathrm{d}M_i/\mathrm{d}z_k}{\mathrm{d}M_j/\mathrm{d}z_k} = \frac{\frac{1}{\tilde{\beta}_i} + \frac{1}{\tilde{\alpha}_i} + \frac{1}{M_i}}{\frac{1}{\tilde{\beta}_j} + \frac{1}{\tilde{\alpha}_j} + \frac{1}{M_j}},$$

because entry in industry i in response to shock to industry k is controlled just by exposure to aggregate objects and not via network interactions.

Now, let us analyze how real GDP responds to productivity shocks in this massless limit. This result should be compared to proposition 2.4.9, which gave the response when there was no extensive margin of entry and exit.

Proposition 2.5.6. *With free entry,*

$$\left(\frac{\mathrm{d}\log(C)}{\mathrm{d}z_k}\right) = \text{scalar } \beta' \Psi_2 (I - \Lambda)^{-1} \left(\mathrm{diag}(\tilde{\alpha})^{-1} + I\right) \Psi_d e_k \alpha_k z_k^{\sigma-2},$$

where

$$\text{scalar} = \left[\beta'\tilde{\alpha} + \beta'\Psi_2(I-\Lambda)^{-1}\mathbf{1}\right]^{-1}.$$

It helps to compare this result to our earlier results. The presence of Λ, which depends on Ψ_1 and Ψ_2, and therefore depends on in-degrees and out-degrees, shows that the most influential industry is not simply the industry who supplies the most or consumes the most. Furthermore, all these matrices are now endogenous objects and depending on how many firms are in each industry, they will take different values. This means that the sufficient statistics approach of the earlier sections will no longer work. Furthermore, the influence measure

$$\beta'\Psi_2(I-\Lambda)^{-1}\left((\mathrm{diag}(\tilde{\alpha}))^{-1} + I\right)\Psi_d,$$

which maps industrial shocks to movements in real GDP is not tied to sales.

To isolate the effect of the shocks to just the extensive margin, we could look at shocks to the fixed costs f_k rather than to labor productivity. Such shocks only propagate through the network due to the change in masses, and therefore give rise to cleaner analytical expressions. However, to keep the results comparable to standard models, I have restricted my attention to productivity shocks (which have both intensive and extensive margin effects).

The results in this section show show that the extensive margin is an important channel through which the network structure affects the equilibrium. However, despite the influence measures being endogenous and related in more realistic ways to the network-structure, this model is still incapable of generating local amplification.

Local Amplification

Now we can turn our attention to the behavior of a shock as it travels through the network. The equilibrium impact of a change in an industry's productivity at the firm level is very stark.

Proposition 2.5.7. *In equilibrium, the effect of a productivity shock z_k on the sales of firm i in industry k is the same as its effect on firm j in industry l.*

$$\frac{\mathrm{d}\log(\text{sales}_{ki})}{\mathrm{d}z_k} - \frac{\mathrm{d}\log(\text{sales}_{lj})}{\mathrm{d}z_k} = 0.$$

At the firm level, the mass of firms in each industry adjusts to ensure that all firms are equally exposed to productivity shocks. At the industry-level, the conclusion is less stark.

Proposition 2.5.8. *In equilibrium, the difference in the response of the profits of industry k relative to industry j to a productivity shock to industry k is given by*

$$\frac{\mathrm{d}\log(s_k)}{\mathrm{d}z_k} - \frac{\mathrm{d}\log(s_j)}{\mathrm{d}z_k} = (e_k - e_j)'(I - \Lambda)^{-1}\left(\mathrm{diag}(\tilde{\alpha})^{-1}\Psi_s e_k\right). \tag{2.1}$$

Numerical simulations suggest that, in equilibrium, this expression is always negative. In other words, the model is still incapable of amplifying shocks locally, and shocks decay as they move away from their source.[2] The intuition for equation (2.1) is the following: The initial impact of a productivity shock to k, holding fixed the mass of firms in each industry, is given by $\mathrm{diag}(\tilde{\alpha})^{-1}\Psi_s e_k$ – this is the traditional input-output effect which captures the total intensity with which each industry uses inputs from industry k. However, with entry, the traditional input-output effect must be multiplied by a new term $(I - \Lambda)^{-1}$. This captures

---

[2]The proof of this result is work in progress.

how the industry's sizes change in response to the intensive margin shock. Recall that

$$\Lambda = \frac{\partial \log(\tilde{\beta}) + \log(\tilde{\alpha})}{\partial \log(M)} \propto \frac{\partial \log(\text{share of sales})}{\partial \log(M)},$$

where the proportionality sign follows from lemmas 2.3.1 and 2.3.2. Therefore, $(I - \Lambda)^{-1}$ captures the cumulative effect of the shock on the size of the industries.

2.5.2   Lumpy Firms

In this section, we consider the equilibrium where $\|\Delta\| > 0$. The equilibrium we focus on is the subgame perfect Nash equilibrium where firms make a simultaneous entry decision in period 1, and in period 2 they play monopolistic competition general equilibrium. Note that for tractability, I restrict the firms' pricing strategies to take the aggregate industry price level and output to be exogenous.

An equilibrium will now feature as many firms in each industry as possible, in the sense that adding a firm to any industry would drive that industry's profits to be negative.

Proposition 2.5.9. *A mixed strategy equilibrium always exists.*

Pure-strategy equilibria need not exist. Technically, the non-existence of a pure-strategy equilibrium means that this model of the economy has non-fundamental or non-exogenous randomness. Intuitively, pure-strategy equilibria can fail to exist due to cycling, where the entry of a firm causes the profits of another firm to go negative. Once that firm exists, another firm can enter that causes the profits of the first entrant to be negative, and so on.

Proposition 2.5.10. *If $N > 2$, then a pure-strategy equilibrium does not always exist.*

We do not focus on non-existence or multiplicity of equilibria in this paper, although these issues are present. Instead, we focus on equilibria that, if they exist, converge to an equilibrium of the non-atomistic economy when we take the limit $\|\Delta\| \to 0$.

An equilibrium when $\|\Delta\| > 0$ is a solution to a nonlinear integer programming problem. Results from computational complexity theory show that we cannot hope to fully characterize the set of equilibria. A naive brute-force computation would become infeasible

very rapidly, since with even a few hundred firms, we may need to consider more cases than there are atoms in the observable universe! Instead, our approach here will be to compare the discontinuous model's equilibria to the equilibria of the continuous model. To demonstrate some of the subtleties, consider the following example illustrated in figure 2.4.
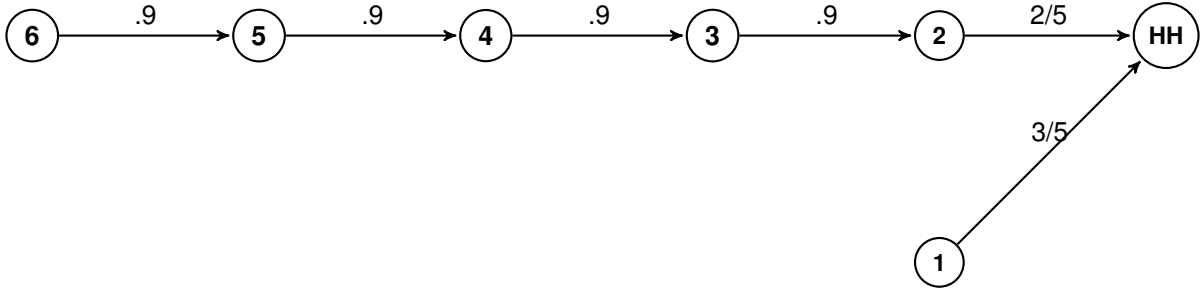
Figure 2.4: *Two production lines selling to the household HH. Industries $2 - 6$ form on production line and industry 1 forms a second (degenerate) production line. The direction of arrows represent the flow of goods and services.*

$$\beta = (2/5, 2/5, 0, 3/5)',$$

$$f = (0.01, 0.01, 0.01, 0.01, 0.01, f_6),$$

$$\sigma = 1.1,$$

$$\varepsilon = (1.2, 3, 3, 3, 3, 3). \Delta = (0.5, 1, 1, 1, 1, 1).$$

Let us consider the equilibrium mass of entrants $M(\Delta, f_6)$ for different values of f_6 and mass vector Δ. First, let us consider the massless limit,

$$M(\mathbf{0}, 0.04) = \begin{pmatrix} 15.7 \\ 17.6 \\ 10.1 \\ 5.8 \\ 3.3 \\ 1.2 \end{pmatrix}, \; M(\mathbf{0}, 0.05) = \begin{pmatrix} 16.9 \\ 16.2 \\ 9.4 \\ 5.3 \\ 3.0 \\ 0.8 \end{pmatrix}, \; M(\mathbf{0}, 0.06) = \begin{pmatrix} 17.8 \\ 15.4 \\ 8.7 \\ 4.9 \\ 2.7 \\ 0.6 \end{pmatrix}, \; M(\mathbf{0}, 0.07) = \begin{pmatrix} 18.7 \\ 14.6 \\ 8.2 \\ 4.6 \\ 2.5 \\ 0.5 \end{pmatrix}.$$

We see how increasing the fixed costs of the final supplier of the chain reduces the mass of firms active in the long chain and increases the mass of firms in the competing short chain in a continuous and intuitive way.

Now, let us consider this same equilibrium away from the massless limit.

$$M(\Delta, 0.04) = \begin{pmatrix} 16 \\ 17 \\ 10 \\ 5 \\ 3 \\ 1 \end{pmatrix}, \quad M(\Delta, 0.05) = \begin{pmatrix} 19 \\ 15 \\ 8 \\ 4 \\ 2 \\ 0.5 \end{pmatrix}, \quad M(\Delta, 0.06) = \begin{pmatrix} 19 \\ 14 \\ 8 \\ 4 \\ 2 \\ 0.5 \end{pmatrix}, \quad M(\Delta, 0.07) = \begin{pmatrix} 33 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

In going from $f_6 = 0.04$ to $f_6 = 0.05$, the discontinuous model amplifies the impact of the shock because the shock is large enough to force a discontinuous change. From $f_6 = 0.05$ to $f_6 = 0.06$ however, the discontinuous model attenuates the impact of the shock since the firms are large enough that they can absorb the losses without exiting. This represents a phase transition since shocks below a threshold are attenuated, and above that threshold, they are amplified. In going from $f_6 = 0.06$ to $f_6 = 0.07$, the interior equilibrium disappears, and the long supply chain collapses. Furthermore, this is the unique equilibrium, so this is not a coordination failure resulting from equilibrium switching. The situation is illustrated in figure 2.5.

This example illustrates that most of the time, we can expect the model to behave like its continuous limit. However, some shocks are amplified when a discontinuous change occurs and other shocks are attenuated. Furthermore, occasionally, when whole industries exit, the discontinuous equilibrium can be very different. These differences are not only large, but they can also be very complex. Figure 2.5 illustrates that when dealing with firm entry and exit, the size or mass of firms in an industry is not a good guide to how important those industries are to the equilibrium. It also illustrates the fact that bailout policy is very important in this model.

2.5.3 Approximation Error

To formalize the approximation error, we need a few definitions. These definitions have some resemblance to the concepts of *natural friends* and *natural enemies* in international trade theory. Motivated by the Stolper and Samuelson (1941) result, Deardorff (2006) defines
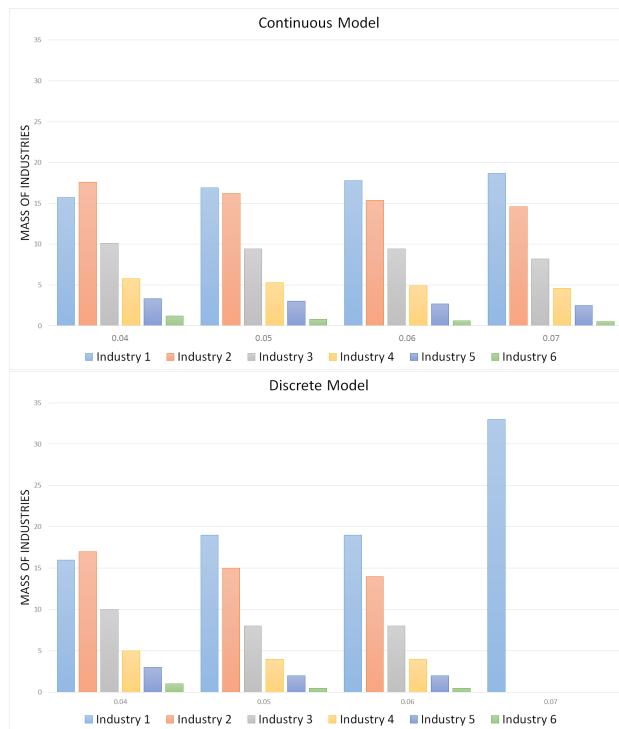
Figure 2.5: *Masses of industries with different levels of f_6. Note that the discrete model amplifies the shock when going from the first panel to the second panel, attenuates it in going from the second to the third, and experiences a chain reaction in going from the third to the fourth.*

an industry to be a natural friend (enemy) of a factor when an increase in its prices will increase (decrease) the returns to that factor. Inspired by these terms, I define the following.

Definition 2.5.1. Let $\pi_i(M_1, \ldots, M_N)$ be the profit function of industry i. Industry j is an *enemy* of industry i when

$$\frac{\partial \pi_i}{\partial M_j} < 0,$$

and industry j is a *friend* of industry i when

$$\frac{\partial \pi_i}{\partial M_j} > 0.$$

Industry j is a *frenemy* of industry i when

$$\frac{\partial \pi_i}{\partial M_j}$$

takes both positive and negative values.

Although the definition is reminiscent of the one in Deardorff (2006), these definitions are *out-of-equilibrium*. In other words, we change the mass of firms in one industry without changing the masses in other industries.

When industries are frenemies, the entry decisions of firms in one industry are strategic substitutes for some regions and strategic complements for other regions. I show that frenemies can only occur with non-degenerate (non-diagonal) input-output connections.

Proposition 2.5.11. *Let $\sigma > 1$. When the network structure Ω is degenerate (diagonal), all industries are enemies. When the network is non-degenerate, industries can be frenemies.*

We can say more than this.

Proposition 2.5.12. *If industry j is a frenemy of industry i, then initially industry i's profits are increasing in M_j and eventually industry i's profits are decreasing in M_j.*

It turns out that we can guarantee that approximating the discontinuous model with a continuous limit will be bad when firms are frenemies.

Theorem 2.5.13. *Let $M(\Delta)$ correspond to the mass of firms in each industry in an equilibrium where firm-level masses are given by Δ. Let $M(0)$ denote equilibrium masses in each industry when $\Delta \to 0$. Then*

$$\|M(\Delta) - M(0)\| \geq \|D\pi(\tilde{M})\|^{-1}\|\pi(M(\Delta))\|,$$

where $\tilde{M} \in co\{M(0), M(\Delta)\}$. Here, co refers to the convex hull and $D\pi$ to the derivative of $\pi(M)$ as a function of M.

In particular, the error gets larger when the profit functions have flat slopes, which can only occur if industries are frenemies. When $\sigma > 1$, this can only happen with network connections, because a model with no network has monotone profit functions. The network makes the profit functions non-monotonic even when industry outputs are substitutes $\sigma > 1$.

2.5.4 Endogenous Markups

So far we have made very strong assumptions about the firms' pricing decisions. Firms in industry k face the following demand function

$$
\begin{aligned}
y(i,k) &= \left(\frac{p(i,k)}{p_k}\right)^{-\varepsilon_k} y_k, \\
&= \left(\frac{p(i,k)}{p_k}\right)^{-\varepsilon_k} \left(\beta_k \left(\frac{p_k}{P_c}\right)^{-\sigma} C + \sum_l \omega_{lk} M_l \left(\frac{p_k}{\lambda_l}\right)^{-\sigma} y(l,j)\right), \\
&= \left(\frac{p(i,k)}{p_k}\right)^{-\varepsilon_k} p_k^{-\sigma} \left(\beta_k P_c^{\sigma} C + \sum_l \omega_{lk} M_l^{\frac{\sigma-1}{\varepsilon_l-1}} \lambda_l^{\sigma} y_l\right).
\end{aligned}
$$

I have so far assumed that a firm only internalizes their effect on demand via its direct effect on $p(i,k)$. That is, the firm assumes that demand is isoelastic with elasticity ε_k and therefore charges constant markups. This assumption is fully rational when firms have zero mass, since they cannot affect the industry-level aggregates by changing their prices. The assumption is also defensible when there at least a handful of firms in the industry, since the firm's effect on the industry-level statistics is negligible.

However, in some of the examples I discuss, I consider cases where there is only one or two firms in an industry, and these firms *do* have appreciable effects on the industry-level

aggregates. One might fear that if these firms were more sophisticated, they could avoid exit by endogenously increasing their markups as they gain more market power. In other words, perhaps, Ford can take greater advantage of its market power once it is the only firm left in an industry, and by charging higher markups, it can withstand greater shocks.

A relatively simple way to address this concern is to have the firm internalize its effects on the industry-level aggregates p_k and y_k (but not on economy-wide aggregates like P_c and C) when it sets its prices. Then we can write the profit function of firm i in industry k as being proportional to

$$p(i,k)^{1-\varepsilon_k} p_k^{\varepsilon-\sigma} - \lambda_k p(i,k)^{-\varepsilon} p_k^{\varepsilon-\sigma}.$$

Maximizing this gives the optimal price as

$$p(i,k) = \frac{\varepsilon_k N_k - (\varepsilon_k - \sigma)}{(\varepsilon_k - 1) N_k - (\varepsilon_k - \sigma)} \lambda_k.$$

In the limit, as $N_k \to \infty$, this gives an exogenous markup of $\varepsilon_k/(\varepsilon_k - 1)$. In the case where the firm is a monopolist, and $N_k = 1$, this gives the markup $\sigma/(\sigma - 1)$. To make this sensible, we need that $\varepsilon_k > \sigma > 1$.

2.6 Bail-outs and the Nature of Externalities

The examples in the previous section suggest that bail-outs and government interventions may be desirable, since the failure of an important firm or industry can take down entire parts of a network. However, figuring out which failures are efficient and which ones are not is impracticable. To implement the optimum, a social planner would not only need access to an infeasible amount of information, but it would also have to solve an intractable computational problem.

In this section, we imagine a scenario where the policy-maker can, in response to a shock, ask firms if other firms should be rescued. The industries will truthfully tell the policy-maker whether a rescue would increase their profits or decrease their profits. Under this assumption, we can investigate conditions under which the profits of an industry align

with those of society.

This exercise is inspired by the example in the introduction, where the president of Ford Motor Company testified in favor of General Motors. The results in this section give sufficient conditions under which the policy-maker can trust a firm's recommendation. The results of this section should not be taken literally as a guide to policy. Rather, this thought experiment tells us about the nature of externalities in this model. The aim of this section will be to prove a series of "innocent by-stander" results, which imply that firms' requests for the bail-outs of other firms can only be trusted if the firms are not reliant on one another.

To start with, we need to define the following notion of connectedness.

Definition 2.6.1. Two firms u and v are connected if there exists a directed path from u to v or a directed path from v to u, on a directed graph defined by Ω.

This lemma will be the workhorse result for the rest of the section.

Lemma 2.6.1. *If firm u and firm v are not connected, then in the event that u fails, and we hold fixed the number of firms in all other industries, $\tilde{\beta}_v$ and $\tilde{\alpha}_v$ remain constant.*

Now we are in a position to state our first result.

Theorem 2.6.2. *Let $\mathcal{B}$ be the set of firms not connected to v. If all firms in $\mathcal{B}$ prefer for v to be rescued, then it is Pareto-efficient for v to be rescued. If the firms in $\mathcal{B}$ disagree with each other, it must be because the firms who want v to be rescued are badly affected through a cascade.*

Consider the example in figure 2.6. Theorem 2.6.2 implies that in the event that GM, and only GM, is about to fail, Ford, Toyota, and Mitsubishi will agree with each other about whether or not GM should be rescued. Furthermore, if they agree that GM should be rescued, then the rescue is Pareto-efficient. However, if Ford, Toyota, and Mitsubishi disagree, then it must be that Ford is adversely affected by GM's failure through the failure of their common supplier Dunlop Tires. Furthermore, the efficiency of the bailout is ambiguous. The welfare implication of theorem 2.6.2 is that Ford's plea that GM be bailed out can only be trusted if Ford is worried about aggregate demand, not cascades. In
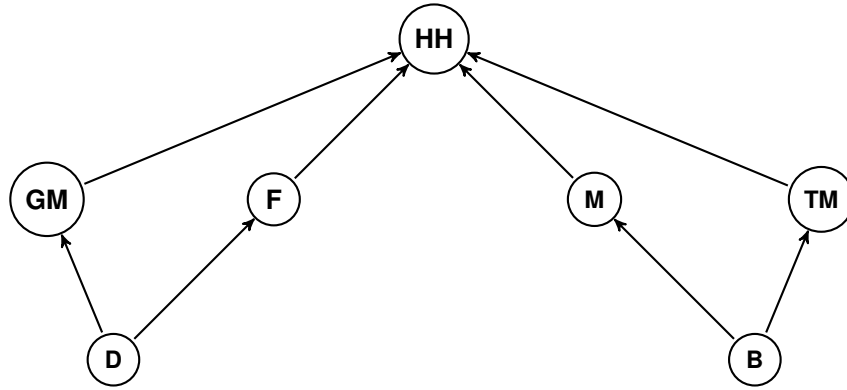
Figure 2.6: *A stylized example of a supply chain. HH is household; GM is General Motors; F is Ford Motor Company; D is Dunlop Tires; M is Mitsubishi Motors; T is Toyota; and, B is Bridgestone Tires, a Japanese tire manufacturer. The arrows represent the flow of goods and services.*

his December 2008 testimony to Congress, Mulally referred to both the general economic downturn, as well as their overlapping supply base, as the reason why GM and Chrysler should be bailed out.

We can sharpen theorem 2.6.2 into the following "innocent by-stander" principle.

Theorem 2.6.3. *Let u and v be two firms. Suppose that the only undirected path from u to v goes through the household. Then if u prefers for v to be rescued, rescuing v increases the utility of the household.*

Note that theorem 2.6.3 considers *undirected* paths, meaning that we disregard the direction of the arrows. So, while *F* and *GM* are unconnected according to definition 2.6.1, there does exist an undirected path from *F* to *GM* that does not go through the household. Theorem 2.6.3 implies that, in the example of figure 2.6, Ford's profits may not be aligned with societal preferences. However, the profits of Ford may be aligned with society when considering the bail-out of an unconnected firm like Lehman Brothers. That is, in a network like the one in figure 2.7, if a shock to L causes F's profits to go down, we can safely infer that we should have bailed L out.

Unfortunately, the cases where we can trust a firm are precisely the cases where we would expect that firm to not be well-informed about the consequences of a failure. These results cast doubt on the idea that we reliably implement bailout policy by surveying a
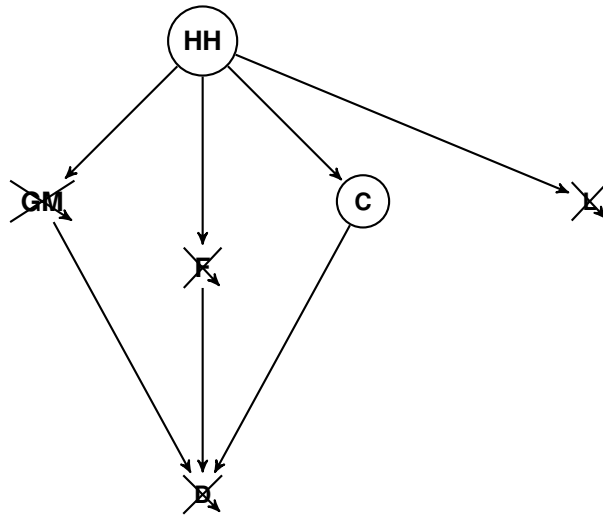
Figure 2.7: *A toy representation of Detroit's economy with GM, Ford, Chrysler, and a Bank selling to a representative household.*

firm's direct rivals.

Bounding Inefficiency

Although the asking mechanism I investigate is simple to understand, it does not necessarily achieve first-best. In particular, asking is sufficient but not necessary for an efficient bail-out. We can bound the losses from using this simplistic mechanism relative to first best, or other, more elaborate mechanisms that can implement first-best.

Theorem 2.6.4. *The fraction of utility lost by asking mechanism relative to utility maximizing*

$$\text{fraction of utility lost} \leq \left(\frac{w^{fb}/P_c^{fb}}{w_c^A/P_c^A} \right)^{\sigma} - 1,$$

where P_c^a is the consumer price index under the asking equilibrium and P_c^{fb} is the consumer price index under the first best equilibrium.

This theorem give us a sense of how large the "mistakes" will be in terms of the real wage.

2.7 Conclusion

This paper highlights the importance of firm entry and exit in propagating and amplifying shocks in a production network. I show that without the extensive margin, canonical input-output models have several crucial limitations. First, a firm's influence depends solely on the firm's role as a supplier to other firms, and its role as a consumer is irrelevant. Second, each firm's influence measure is exogenous, and the exogenous influence measures are sufficient statistics for the input-output matrix. In this sense, for every input-output model, there exists a non-interconnected model with different parameters that has the same equilibrium responses. Third, a firm's influence is well-approximated by the firm's size, so the nature of interconnections does not matter as long as it results in the same distribution of firm sizes. Finally, I show that canonical input-output models lack the ability to locally amplify shocks. That is, a shock to one industry always has the largest impact at its source, and its effect decays geometrically as it travels away from its source. This shows why a firm has to be large in order to have a meaningful aggregate impact.

However, when the mass of firms in each industry can adjust endogenously, all of these properties disappear. The influence of an industry is endogenous, depends on its role as a supplier and as a consumer of inputs, and is not well approximated by its size. Two industries with the same demand-chains can have very different effects on the equilibrium if their supply-chains are different, and vice versa. Furthermore, there are no sufficient statistics that summarize the network. In this sense, the model is not isomorphic to a non-network model with different parameters.

Despite these features, I show that in the limit where all firms have zero mass, the model with entry still lacks the ability to locally amplify shocks. The size of each industry responds smoothly to shocks and the effects still decay geometrically as a shock travels from its source. However, when firms in some industries are "granular," then a failure of one of these firms can snowball into an avalanche of failures. I show that we can expect these cascades to occur when firms' entry decisions switch from being strategic substitutes to being strategic complements.

Finally, motivated by the possibility of catastrophic failure in this model, I show conditions under which the objectives of a firm are aligned with those of society. In principle, this would allow a policy-maker to formulate bail-out policy by surveying other firms. Unfortunately, the results show that we should not be very optimistic about this strategy, since the only circumstances when firms are trustworthy are precisely those circumstances where we would expect them to be as uninformed as the policy-maker. This is because a firm's incentives are only aligned with society when a firm's exposure to the failure of the other firm is through general equilibrium effects. In other words, in cases where the firm is not directly linked to the troubled industries. Despite this, these results do show that the welfare consequences of shocks that travel from one firm to another are different depending on whether they arrive via network connections or general equilibrium effects.

Chapter 3

Asymmetric Inflation Expectations, Downward Rigidity of Wages, and Asymmetric Business Cycles

3.1 Introduction

Keynesian macroeconomic theory posits that sticky wages are a crucial feature of labor markets. Rigid wages can cause involuntary unemployment, amplify fluctuations in employment at business cycle frequencies, and break monetary neutrality. In recent years, a large number of empirical papers have shown not only that wages are very sticky, but that there is a clear asymmetry in the way that wages are sticky. In particular, wages appear more flexible when they are rising than when they are falling. Examples include Barattieri *et al.* (2010), Dickens *et al.* (2007), Daly and Hobijn (2013), as well as the seminal contribution by Bewley (1999). Furthermore, recent work by Kaur (2012) shows that downward nominal rigidities distort labor market outcomes in rural India.

In this paper, I argue that informational frictions for households can help to explain the asymmetric adjustment of wages during the business cycle. This paper makes two main contributions: (1) it documents the existence of a statistically robust asymmetry in

how households form their expectations of inflation; in particular, I show that households are much better at anticipating accelerations in the inflation rate than decelerations. In a typical model, this asymmetry in household beliefs feeds into asymmetry in wage-setting where wages respond more vigorously to inflationary forces than disinflationary forces. This makes demand-driven business cycles asymmetric. Positive monetary policy shocks (or more generally positive demand shocks) are highly inflationary but do not increase output by very much, whereas negative monetary policy shocks (or negative demand shocks) are not very disinflationary but cause large unemployment. I show that this asymmetry in beliefs are unique to households and are not present for professional forecasters. (2) I micro-found the source of asymmetric belief formation in an equilibrium model where households are ambiguity-averse and are trying to make robust decisions. I show that with this microfoundation, optimal monetary policy is still subject to the Lucas critique, and the central bank does not have a systematic inflationary bias despite the existence of downward rigidities.

The intuition for my microfoundation is simple. When negotiating their wages, workers observe their nominal wages perfectly, but foresee the real cost of the goods and services that they will consume imperfectly. As in Lucas (1973), workers face a signal extraction problem when trying to determine their purchasing power. However, households' signals of the general price level are subject to Knightian uncertainty, since households do not know precisely how informative their signals are. This means that they will be more sensitive to inflationary news than disinflationary news because, for a fixed nominally denominated employment contract, inflationary news lowers their purchasing power whereas disinflationary news raises it (relative to their prior expectations). This asymmetry of beliefs can then show up in wage-setting since distrustful workers will, to the extent that they can, refuse wage cuts in the presence of deflation, but demand wage increases in the presence of inflation.

This distrustful attitude of workers towards the inflation rate is attested to in many surveys. For example, according to Shiller (1997), the "biggest gripe about inflation" expressed by 77% of the general public is that "inflation hurts my real buying power. It

makes me poorer." Interestingly, only 12% of economists chose this answer.[1] The households' answer makes sense in partial equilibrium, since over short horizons, households can treat their wages as known and exogenous to the inflation rate.

I embed ambiguity into a general equilibrium model and find that monetary shocks have very asymmetric effects on wage inflation and output. In particular, positive monetary shocks result in high wage inflation and small booms, since households react strongly to the inflationary signals by demanding wage increases. On the other hand, negative monetary shocks cause large unemployment and relatively small disinflation, since households distrust the disinflationary signals and refuse wage cuts. I verify the model's predictions about the asymmetric impact of monetary policy on output and wage inflation using time series data for the US.

The asymmetry implied by the model substantially alters the welfare costs of business cycles when compared to Lucas (1987). Whereas in the Lucas model, positive shocks cancel out with negative shocks so that the welfare cost of fluctuations is second order, in this model, positive demand shocks do not cancel with negative demand shocks, so stabilization policy reaps first order gains. This harks back to the point made by De Long and Summers (1988) that demand stabilization may fill in the troughs without shaving the peaks. Schmitt-Grohé and Uribe (2011) have also recently drawn attention to this point. I also investigate optimal monetary policy in my model. The received wisdom in the literature, following Akerlof *et al.* (1996), is that if wages are downwardly rigid, then central banks should have an inflationary bias to "grease the wheels" of the labor market. This way real wage cuts can be masked by a positive inflation rate. Unfortunately, in this model, inflationary biases from the central bank are not helpful since household expectations adjust to take them into account. Any inflationary bias built into central bank policies are undone by endogenously-formed household expectations. In other words, the model predicts an asymmetric equilibrium but the central bank is powerless to do anything about the asymmetry.

---

[1]Instead, the most popular reason given by economists was "inflation makes it hard to compare prices, forces me to hold too much cash, and is inconvenient." Only 7% of households chose this answer.

Other theoretical treatments of downwardly rigid wages often take the rigidity as given and investigate its consequences (e.g. Daly and Hobijn (2013); Schmitt-Grohé and Uribe (2011); Kaur (2012); Akerlof *et al.* (1996); Hall (2005)). These models are usually motivated by an exogenous fairness norm, and assumptions about the function relating wages to worker effort. Akerlof (1982) is a seminal paper in this strand of the literature. Other attempts to microfound downward wage rigidity are based on implicit contracts, where firms insure their workers against fluctuations by uncoupling the real wage from marginal product of labor; A leading example is Holmstrom (1983), but this literature is focused on real wages and does not bring inflation into the analysis. The theory I propose delivers downward rigidity that is broadly consistent with the data at the micro level, but it also has novel macro implications. The model predicts that monetary policy has asymmetric effects on output and wage inflation. There is already evidence that monetary policy has asymmetric effects on output, for example in De Long and Summers (1988), Cover (1992), and Angrist *et al.* (2013). I present new evidence that monetary policy also has asymmetric effects on wage inflation consistent with the model's prediction.

The model of Knightian uncertainty I use in this paper is the one axiomatized by Gilboa and Schmeidler (1989). In this framework, workers have multiple priors about the information content of their signals, and they act according to their worst-case prior when making decisions. A similar modelling device is used by Epstein and Schneider (2008) in the context of asset pricing to model skewness in asset returns. Kuhnen (2012) finds empirical evidence in support of asymmetric learning in the context of financial markets, where agents are overly pessimistic in the loss region. Another recent paper that incorporates ambiguity aversion into a macroeconomic model is by Ilut and Schneider (2012). However my results differ markedly from theirs both in terms of the research question and the set up of the model. [2]

The outline of this paper is as follows. In section 3.2, I set out a basic partial equilibrium

---

[2]Ilut and Schneider (2012) are interested in how ambiguity about the level of productivity affects an economy, whereas I am interested in how ambiguity about the informativeness of price signals affect wage setting.

model with ambiguity aversion that demonstrates my mechanism. In section 3.3, I endogenize prices and output, and study the effects of monetary shocks on real and nominal variables. In section 3.4, I present empirical evidence in favor of asymmetric adjustment of beliefs, and time series evidence that wage inflation responds asymmetrically to positive and negative monetary shocks as predicted by the model. I also discuss the extent to which the model can explain the cross-sectional distribution of wage-changes. In section 3.5, I investigate a simple optimal policy problem and compute the welfare cost of business cycles. In section 3.6, I embed my mechanism into a standard New Keynesian model with sticky wages, draw out some of its implications. I find that time-varying Knightian uncertainty can act like a cost-push shock in the economy, creating a tradeoff between inflation and output stabilization for the central bank even though there are no supply shocks in the model. I summarize and conclude in section 3.7.

3.2   Partial Equilibrium Model

Consider the following partial equilibrium model that establishes the intuition for the rest of the paper. Suppose that there is an employer and a worker. The worker has log utility in his real wage, is endowed with a unit of labor, and an exogenous outside option d (I take this to be the utility of leisure). In other words, his preferences are given by

$$u(w_t/p_t, x_t) = \log(w_t/p_t)\mathbf{1}(x_t = 1) + d\mathbf{1}(x_t = 0),$$

where x_t is a binary variable for whether or not he works, w_t is the nominal wage, p_t is the price level in period t and d is an exogenous outside option. The employer makes a nominal wage offer w_t to the worker, who then chooses whether or not to work. If the worker does not work, he receives the exogenous outside option d.

   The worker chooses to work if

$$\mathrm{E}_t\left(\log\left(\frac{w_t}{p_t}\right)\right) \geq d,$$

where the expectation is taken with respect to the worker's information set. Since the

nominal wage is always known with certainty, we can rearrange this expression to get that the lowest wage for which the worker will work is

$$w_t = \exp\left(d + E_t(\log(p_t))\right). \tag{3.1}$$

Suppose that workers receive a public signal s_t about the inflation rate. Then we can rewrite (3.1) as

$$\log(w_t) = d + \log(p_{t-1}) + E(\log(\pi_t)|s_t),$$

where π_t is the inflation rate from period $t-1$ to t and workers are assumed to know the price level in the previous period. This equation makes clear that the wage inherits the properties of the conditional expectation function when viewed as a function of the signal s_t. If households' expectations of inflation are, for some reason, asymmetric (they rise more quickly than they fall), then the wage will also behave asymmetrically with respect to inflationary pressures.

The basic mechanism of the model in this paper is that households place greater weight on inflationary news than disinflationary news. To motivate this assumption, we can look for evidence of this asymmetry by using inflation expectation surveys of households. I use the Michigan survey of inflation expectations. Denote inflation in period t by π_t and median household inflation expectations of inflation 12 months ahead by $\hat{\pi}_{t+12|t}$. In figure 3.1, expected revisions to the inflation rate $\hat{\pi}_{t+12|t} - \pi_t$ are plotted against actual changes to the inflation rate $\pi_{t+12} - \pi_t$. As expected, we see a steep convexity, indicating that the median household's expectations of inflation are more responsive to positive rather than negative changes to inflation. This finding is a direct confirmation of the model's underlying mechanism.

In figure 3.2, we see that the median forecasts made by professional forecasters do not exhibit this convexity. This suggests that source of the asymmetry, at least in the United States, is in how households process information, rather than in the information itself. Furthermore, the fact that there is an asymmetry in the beliefs of households is unique to this model and would not be found in preference-based theories that rely on loss-aversion

or fairness norms. This empirical finding is in the same spirit as the experimental results of Fehr and Tyran (2001) who emphasize that subjects respond weakly to deflationary shocks and strongly to inflationary shocks, although my results are about accelerations and decelerations in the inflation rate.

To demonstrate how ambiguity-aversion can deliver convex conditional expectations, suppose that the price level p_t is given by

$$\log(p_t) = \mu + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2),$$

and both the worker and the employer know μ. Let s_t be a noisy public signal of the price shock ε_t,

$$s_t = \varepsilon_t + \varepsilon_s, \quad \varepsilon_s \sim \mathcal{N}(0, \sigma_s^2).$$

Note that

$$\varepsilon_t | s_t, \mu, \sigma^2, \sigma_s^2 \sim \mathcal{N}\left(\frac{\sigma^2}{\sigma^2 + \sigma_s^2} s_t, \frac{\sigma_s^2 \sigma^2}{\sigma^2 + \sigma_s^2}\right).$$

This means we can rewrite the work condition (3.1) as

$$w_t \geq \exp\left(d + \mu + \frac{\sigma^2}{\sigma^2 + \sigma_s^2} s_t\right).$$

Now Suppose that there is ambiguity about how informative the signal s_t is. That is, the signal-to-total variance ratio $\frac{\sigma^2}{\sigma^2 + \sigma_s^2}$ is unknown. For example, suppose that the worker knows only that $\sigma_s \in [\underline{\sigma}_s, \overline{\sigma}_s]$. Gilboa and Schmeidler (1989) have axiomatized and provided a representation theorem for the preferences of such agents. In particular, such agents follow a minmax procedure, where they make decisions that maximize their worst-case expected utility. In other words, when information quality is ambiguous, expression (3.1) becomes

$$
\begin{aligned}
w_t &= \max_{\sigma_s \in [\underline{\sigma}_s, \overline{\sigma}_s]} \exp\left(d + \mu + E(\log(p_t)|s_t)\right) \\
&= \exp\left(d + \mu + \max_{\sigma_s} E(\log(p_t)|s_t)\right) \\
&= \exp\left(d + \mu + \tilde{E}(\log(p_t)|s_t)\right).
\end{aligned}
\tag{3.2}
$$

For notational convenience, I denote the expectations taken with respect to the worst-case
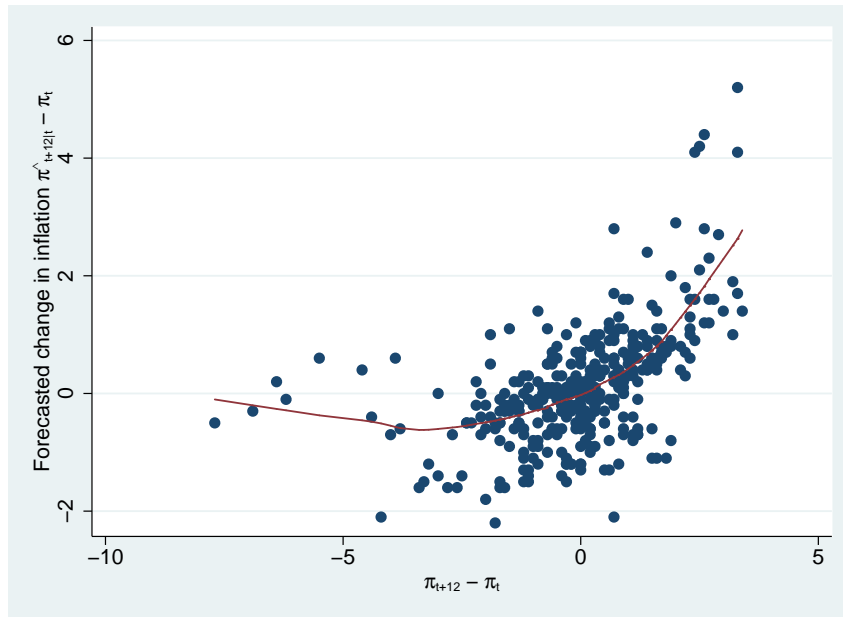
Figure 3.1: *Forecast revisions of the annual inflation rate by the median household in the Michigan Survey of Inflation Expectations from 1983-2012, plotted against realized changes in the annual inflation rate as measured by the CPI.*
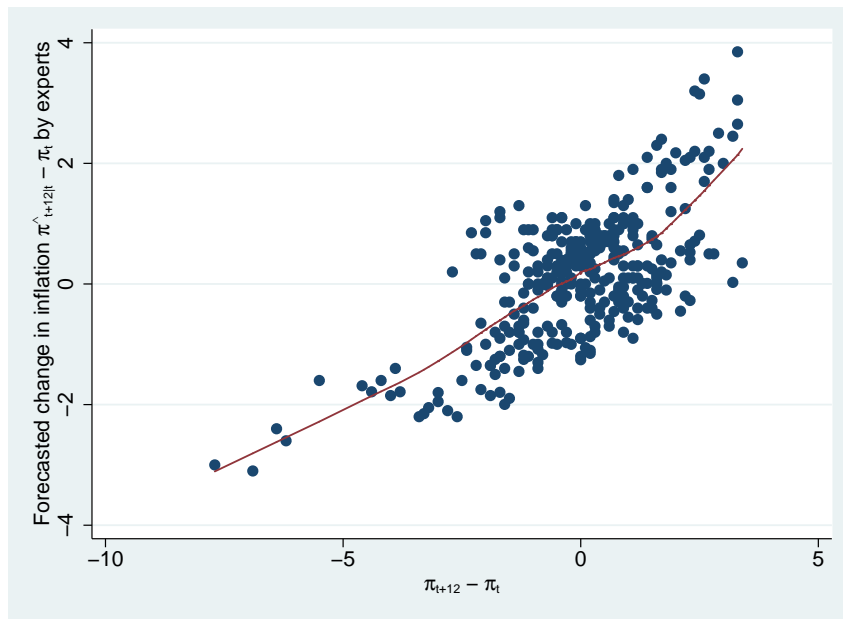
Figure 3.2: *Forecast revisions by the median professional forecaster in the Michigan Survey of Inflation Expectations from 1983-2012, plotted against realized changes in the annual inflation rate as measured by the CPI.*

prior by $\tilde{E}$. Maximizing worst-case expected outcomes in this way is very similar to the "Robustness" framework proposed by Hansen and Sargent (2011). The specific information structure I use is similar to the one posited by Epstein and Schneider (2008), who use it to study skewness in asset prices.

In this section, the source and nature of the ambiguity is not important. It could be due to Knightian uncertainty about official statistics, or a reduced form representation of the fact that consumers have idiosyncratic consumption baskets and there is ambiguity about the extent to which official statistics are relevant to one's individual consumption basket. Alternatively, we could assume that there is ambiguity about σ^2, to capture Knightian uncertainty about demand shocks. A final interpretation is that households are playing a game against the statistical agency in the country, and political pressures on the statistical agency result in the public signals being less informative in the presence of inflation than disinflation. We will return to these issues later, for now, let us take (3.2) as given.

Expression (3.2) implies that $\sigma_s = \underline{\sigma}_s$ when $s_t \geq 0$, and $\sigma_s = \overline{\sigma}_s$ when $s_t < 0$. In figure 3.3, we see an asymmetry in the adjustment of wages to signals of the price level. In particular, wages increase much more rapidly in response to inflationary signals than they fall in response to disinflationary signals. Similar results obtain for the more general constant-relative-risk-aversion utility case, and in the case with ambiguity in the variance of the monetary shock σ instead of ambiguity in the variance of the noise term σ_s.
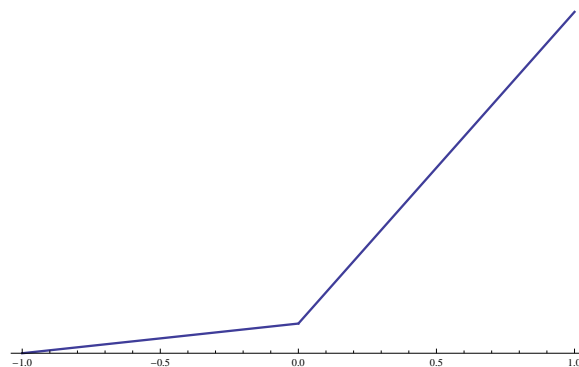
Figure 3.3: *Critical wage as a function of s.*

The intuition for this result is that households take inflationary news very seriously,

since the worst case scenario is that bad news is very informative. On the other hand, households distrust or ignore disinflationary news, since the worst case scenario is that good news is uninformative.

3.3 A simple general equilibrium model

So far, I have used a partial equilibrium model to draw out the implications of asymmetric expectations and ambiguity for the individual workers and employers. In this section, I show that the insights of the previous section survive in general equilibrium with endogenous prices and output.

Consider a two period model with a representative firm and a continuum of identical households. In period 1, nature sets money supply M, and there is a noisy public signal s of M. The firm posts a nominal wage W conditional on the signal, and households decide whether or not to apply for a job. In period 2, M becomes common knowledge, the firm chooses the fraction of the population it wishes to employ, and workers spend the money supply on consuming the output. Intuitively, in period 2, nominal wages are fixed, but the price level changes – this proxies a world where nominal wages are fixed over the length of a contract while prices continue to change.

In period 1, following our earlier discussion, the firm sets the wage to equate the utility of working with the outside option

$$\tilde{E}(u(C)|s) = d, \tag{3.3}$$

where C is consumption of workers when employed and d is the exogenous outside option in utility terms. Let households have log utility so that the wage, in period 1, is given by

$$\tilde{E}(u(C)|s) = \tilde{E}(\log(C)|s) = \tilde{E}(\log(W/P)|s) = d. \tag{3.4}$$

This makes the households indifferent between working and consuming their outside option.

Rearrange this for the wage to get

$$\log(W) = d + \tilde{E}(\log(P)|s). \tag{3.5}$$

In period 2, the stock of money is revealed, the firm sets marginal product of labor equal to the real wage

$$f'(L) = \frac{W}{P}. \tag{3.6}$$

To give a role to money, suppose that households have a cash-in-advance constraint, so that their total expenditures have to equal the money supply

$$PC = M. \tag{3.7}$$

Market clearing for the consumption good implies that

$$C = f(L).$$

Let the firm's production technology be given by

$$f(L) = L^{\alpha}, \quad \alpha \in (0,1),$$

then the firm's first order condition (3.6) implies that

$$f(L) = L^{\alpha} = \left(\frac{W}{\alpha P}\right)^{\frac{\alpha}{\alpha-1}}.$$

Combine this with market clearing, and (3.7) to get

$$\frac{M}{P} = C = L^{\alpha} = \left(\frac{W}{\alpha P}\right)^{\frac{\alpha}{\alpha-1}}. \tag{3.8}$$

Rearrange this to get

$$P = \frac{M^{1-\alpha}W^{\alpha}}{\alpha^{\alpha}}, \tag{3.9}$$

so the equilibrium price is a geometric average of the money stock and the wage. Substitute this expression for P into the wage setting equation (A.2) to get

$$\log(W) = \frac{d}{1-\alpha} + \tilde{E}\left(\log M|s\right) - \frac{\alpha}{1-\alpha}\log(\alpha) \tag{3.10}$$

as the wage in equilibrium. To get equilibrium output, substitute the equilibrium price (3.9) into (3.7) to get

$$C = \frac{M}{P} = \frac{\alpha^\alpha M}{M^{1-\alpha} W^\alpha} = \left(\frac{\alpha M}{W}\right)^\alpha.$$

Finally, equilibrium labor is given by using the production function $L^\alpha = C$ to get

$$L = \alpha \frac{M}{W}. \tag{3.11}$$

Equations (3.10) and (3.11) show that if, for whatever reason, conditional expectations of the money shock as a function of the signal are more convex than the signal is as a function of the money shock, we should observe asymmetries in wage-setting and in employment fluctuations. As before, ambiguity aversion towards the underlying shocks to the money supply can deliver the asymmetric conditional expectation function seen in figure 3.1. Suppose that s is a normal noisy signal of the shock to $\log(M)$. Denoting logs in lower case letters,

$$m|s \sim \mathcal{N}\left(\mu + \frac{\sigma^2}{\sigma^2 + \sigma_s^2} s, \frac{\sigma^2 \sigma_s^2}{\sigma^2 + \sigma_s^2}\right).$$

Denote the signal-to-total variance ratio $\frac{\sigma^2}{\sigma^2 + \sigma_s^2}$ by ψ, and note that $\psi \in [\underline{\psi}, \overline{\psi}]$. Then the equilibrium wage (3.10) is

$$w = \frac{d}{1-\alpha} + \mu + \overline{\psi} s \mathbf{1}(s \geq 0) + \underline{\psi} s \mathbf{1}(s < 0) - \frac{\alpha}{1-\alpha} \log(\alpha),$$

and equilibrium employment (3.11) is

$$l = \frac{1}{1-\alpha} \log(\alpha) + m - \frac{d}{1-\alpha} - \mu - \overline{\psi} s \mathbf{1}(s \geq 0) - \underline{\psi} s \mathbf{1}(s < 0). \tag{3.12}$$

In the benchmark case of full information, $\underline{\psi} = \overline{\psi} = 1$, employment is independent of monetary shocks and the wage is a linear function of the size of the monetary shock. This corresponds to the neoclassical case without frictions. In the case with no ambiguity, $0 \leq \underline{\psi} = \overline{\psi} < 1$, the nominal wage and the level of employment are linear in monetary shocks. The intuition here is the same as for the Lucas (1973) islands model. In the case with ambiguity, $0 \leq \underline{\psi} < \overline{\psi} \leq 1$, shown in figure 3.4, we have asymmetric nominal wage adjustment and employment fluctuations in response to monetary shocks. So we recover the

intuition from the partial equilibrium model in section 3.2, but with additional predictions about the level of employment and the effects of monetary policy (both of which were absent in the partial equilibrium model).
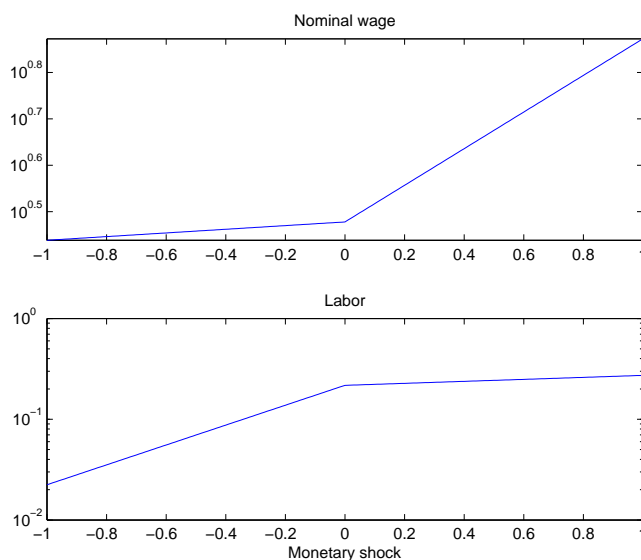
Figure 3.4: *The nominal wage and employment as a function of shocks to money supply.*

3.4   Empirical Evidence

In this section, we look at the extent to which this stylized model is consistent with patterns found in the data. I present three types of evidence: (1) direct evidence of asymmetry in household beliefs towards inflation; (2) time-series evidence from the United States relating wage and price inflation to monetary shocks; (3) evidence from the cross-sectional distribution of wage changes in different countries and different time periods;

3.4.1   Evidence on Asymmetric Expectations

The basic mechanism of the model is that households place greater weight on inflationary news than disinflationary news. We can try to test for this mechanism directly by using inflation expectation surveys of households. Technically, the expectations of the agents in

the model are not unique, since they have multiple priors. So, I assume that individuals report their "worst-case" or "effective" beliefs in surveys – these are the beliefs that would rationalize their behavior if they were Bayesians.

I use the Michigan survey of inflation expectations. Denote inflation in period t by π_t and median household inflation expectations of inflation 12 months ahead by $\hat{\pi}_{t+12|t}$. As mentioned previously, figure 3.1 plots expected revisions to the inflation rate $\hat{\pi}_{t+12|t} - \pi_t$ against actual changes to the inflation rate $\pi_{t+12} - \pi_t$. As expected, we see a kink at zero, indicating that the median household's expectations of inflation are more responsive to positive rather than negative changes to inflation. The asymmetry is a direct confirmation of the model's underlying mechanism. A placebo test, plotted in figure 3.2, shows that the median forecasts made by professional forecasters do not exhibit a kink or convexity.

A regression version of these graphs will allow us to control for covariates and do a formal hypothesis test of the piecewise linear conditional expectation function generated by ambiguity-aversion. To that end, consider the following reduced-model for household inflation expectations

$$\hat{\pi}_{t+12|t} = c_0 + c_1\hat{\pi}_{t+11|t-1} + c_2 s_t \mathbf{1}(s_t \geq 0) + c_3 s_t \mathbf{1}(s_t \leq 0) + c_4 \pi_{t-1} + \varepsilon_t.$$

So, the median household's expectations of future inflation are a linear function of the median household's expectations last month, the value of inflation last month, which reflects the publicly available information in period t, and a piecewise linear function of the signal. In table 3.1, I proxy for the signal received by the households by using the realized change in the inflation rate. In other words, I set

$$s_t = \pi_{t+12} - \pi_t.$$

The results of this regression are reported in table 3.1. In both specifications, we can reject the null hypothesis that $c_2 = c_3$ at the 1% significance level. All of the results on beliefs are entirely robust to controlling for the demographic characteristics of the respondents, namely, their age group (18-34, 35-54, 55+), their region (West, North Central, North East,

115

South), their gender, their income group (bottom tertile, middle tertile, top tertile), and their education level (high School or less, some college, college).

In table 3.1, I proxy for the signal received by households using the realized change in the inflation rate. Therefore, my estimates suffer from attenuation bias, since the realized change is an imperfect measure of the signal received by the household. An alternative approach is to suppose that the signal received by households and professional forecasters is the same. After all, the information used by professional forecasters is mostly publicly available. Denote the forecasts of headline inflation 12 months ahead made in period t by $\pi^e_{t+12|t}$. I use the change in inflation predicted by professional forecasters

$$expert^+_t = (\pi^e_{t+12|t} - \pi_t)\mathbf{1}(\pi^e_{t+12|t} - \pi_t \geq 0), \quad expert^-_t = (\pi^e_{t+12|t} - \pi_t)\mathbf{1}(\pi^e_{t+12|t} - \pi_t \geq 0)$$

as my measure for the signal received by households. To the extent that this is a better measure of the signal received by households (say the signal received by households is literally the median forecast), then this regression should suffer from less attenuation bias. The results of this test are reported in table 3.2. As before, we can reject the hypothesis that positive and negative news are treated symmetrically at the 5% or 1% significance level depending on the specification. The general lesson we learn from these results is that in the US, households are much better at anticipating accelerations in the inflation rate than decelerations.

Since the supporting data is from after the Great moderation, one may may question the extent to which such asymmetries can persist in countries with high (but stable) inflation rates. Using household expectations data from Argentina, I verify that that higher average inflation does not appear to affect the existence of the asymmetry. To this end, I run the same regression with data from Argentina and present the results in table 3.3. As predicted by the theory, the point estimate for c_2 is much larger than for c_3. Since we have many fewer observations, the parameters are imprecisely estimated, and we cannot reject the hypothesis that the coefficients are the same. The inflation data used here are from a private consulting firm, used in the study by Perez-Truglia et al, and are not official figures from

116

Table 3.1: *Responsiveness of Household Inflation Forecasts to Positive and Negative Shocks*

	(1) $\hat{\pi}_{t+12\|t}$	(2) $\hat{\pi}_{t+12\|t}$
$\hat{\pi}_{t+11\|t-1}$	0.957***	0.716***
	(0.02)	(0.06)
$(\pi_{t+12} - \pi_t)\mathbf{1}(\pi_{t+12} - \pi_t \geq 0)$	0.062**	0.131***
	(0.03)	(0.03)
$(\pi_{t+12} - \pi_t)\mathbf{1}(\pi_{t+12} - \pi_t < 0)$	−0.019	−0.001
	(0.02)	(0.03)
π_{t-1}		0.164***
		(0.04)
Constant	0.100*	0.317***
	(0.06)	(0.08)
Observations	407	407

Newey-West *t* statistics in parentheses with lag parameter 4.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Columns regress median 12-months ahead inflation expectations of households on realized positive and negative changes to the actual headline CPI inflation rate and other covariates. The inflation expectation data comes from the Michigan Survey of Consumers and the inflation data comes from the BLS. The question households are responding to in the Michigan survey is "During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now? By what percent do you expect prices to go up, on the average, during the next 12 months?" Column (1) and Column (2) are the same except that column (2) controls for the lagged inflation rate. In both specifications, the coefficient on positive changes to the inflation rate have larger magnitude than the one for negative changes in the inflation rate at the 1% significance level. Hats indicate forecasts, subscripts indicate time periods. The sample period is monthly data from January 1978 to December 2012. Observations are at the month level.

Table 3.2: *Responsiveness of Household Inflation Forecasts to Professional Forecasts*

	(1) $\hat{\pi}_{t+12|t}$	(2) $\hat{\pi}_{t+12|t}$	(3) $\hat{\pi}_{t+12|t}$
π_t	0.430*** (0.03)	0.261*** (0.05)	0.113** (0.04)
$expert_t^+$	0.312*** (0.08)	0.270*** (0.08)	0.098** (0.04)
$expert_t^-$	−0.002*** (0.00)	−0.002*** (0.00)	−0.001*** (0.00)
π_{t-1}		0.164*** (0.05)	0.057 (0.05)
$\hat{\pi}_{t+11|t-1}$			0.565*** (0.06)
Constant	1.686*** (0.11)	1.719*** (0.10)	0.797*** (0.12)
Observations	415	414	414

Newey-West t statistics in parentheses with lag parameter 4.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Columns regress median 12-months ahead inflation expectations of households on positive and negative changes to the inflation rate forecasted by the median professional forecaster and other covariates. The inflation expectation data for households and experts comes from the Michigan Survey of Consumers and the inflation data comes from the BLS. The question households are responding to in the Michigan survey is "During the next 12 months, do you think that prices in general will go up, or go down, or stay where they are now? By what percent do you expect prices to go up, on the average, during the next 12 months?" All columns control for the actual inflation rate. Column (2) also controls for the lagged inflation rate, and Column (3) controls for the lagged inflation rate and the lagged median household inflation forecast. For column (1) and (2) we can reject the hypothesis that the coefficient on $expert^+$ and $expert^-$ are the same at the 1% significance level, and for column (3) we can reject this hypothesis at the 5% significance level. Hats indicate forecasts, subscripts indicate time periods. The sample period is from January 1978 to December 2012. Observations are at the month level.

the government (which are widely known to be unreliable).

Further supportive evidence of asymmetry is found in the working paper by Perez-Truglia et al. They perform a randomized controlled experiment on Argentinean households and find that household expectations respond more (almost five times more strongly) to inflationary news than disinflationary news. This asymmetry disappears if instead of news about the inflation rate, households are given news about the change in the price of a specific set of goods. This indicates that the source of ambiguity might be in how aggregate statistics relate to one's individual consumption basket, rather than the conduct of monetary policy or the source of the information.

3.4.2 Evidence from aggregate time-series

Next, we look at time series evidence of the relationship between wage inflation, price inflation, and monetary policy shocks. The model implies that wage inflation should respond more strongly to positive monetary shocks than negative ones. To test this, I use a measure of structural monetary shocks from Coibion and Gorodnichenko (2012). Following Romer and Romer (2004), I estimate the following reduced form model

$$\pi_t^w = a_0 + \sum_{j=1}^{J} a_j \pi_{t-j}^w + \sum_{k=0}^{K} b_k \varepsilon_{t-k}^+ + \sum_{l=0}^{L} c_l \varepsilon_{t-l}^- + v_t,$$

where π_t^w is annual wage inflation, ε_t^+ and ε_t^- are positive and negative monetary shocks, and v_t is the error term. The measure of wage inflation is the seasonally adjusted annual percent change of average hourly earnings of production and nonsupervisory employees for the total private sector taken from Federal Reserve Economic Database.

We can then test the hypothesis that

$$\sum_{k=0}^{K} b_k + \sum_{l=0}^{L} c_l = 0,$$

or that the cumulative effect of a positive shock on wage inflation is the same as the cumulative effect of a negative shock. I use the BIC to select the autoregressive lag length J, although the results are robust to changing the number of lags to be higher (for example,

Table 3.3: *Responsiveness of Household Inflation Expectations to Positive and Negative Shocks in Argentina*

	(1) $\hat{\pi}_{t\|t}$	(2) $\hat{\pi}_{t\|t}$
$\hat{\pi}_{t-1\|t-1}$	0.778*** (0.07)	0.911*** (0.04)
π_{t-1}	0.195* (0.08)	
$(\pi_t - \pi_{t-1})\mathbf{1}(\pi_t - \pi_t \geq 0)$	0.909* (0.37)	1.015** (0.38)
$(\pi_t - \pi_{t-1})\mathbf{1}(\pi_t - \pi_t < 0)$	0.365 (0.62)	0.185 (0.60)
Constant	0.910 (1.01)	1.843 (1.08)
Observations	79	79

t statistics in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Columns regress median contemporaneous inflation expectations of households on positive and negative changes to the inflation rate as measured by a private consulting company. The inflation expectation data and the inflation data were kindly shared by Perez Truglia et al (2014). Both columns control for the lagged expected inflation rate. Column (2) also controls for the lagged inflation rate. The hypothesis that the coefficients for positive and negative changes are equal in magnitude cannot be rejected. Hats indicate forecasts, subscripts indicate time periods. The sample period is from August 2006 to March 2013. Observations are at the month level.

Table 3.4: *Responsiveness of Wage Inflation to Monetary Policy Shocks*

	(1) π_t^w	(2) π_t^w	(3) π_t^w
ε_t^+	0.234***	0.202***	0.201***
	(0.08)	(0.07)	(0.07)
ε_t^-	−0.110*	−0.076	−0.081
	(0.06)	(0.06)	(0.07)
ε_{t1}^+		0.132*	0.126*
		(0.08)	(0.07)
ε_{t-1}^-			0.025
			(0.10)
Constant	Yes	Yes	Yes
Number of autoregressive lags	4	4	4
Observations	471	471	471

Newey-West *t* statistics in parentheses

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

Notes: Columns regress wage inflation on positive and negative monetary policy shocks and other covariates. Wage inflation is seasonally adjusted annual percent change of average hourly earnings of production and nonsupervisory employees for the total private sector taken from the Federal Reserve Economic Database. The monetary policy shocks are the structural shocks from Coibion and Gorodnichenko (2012). The number of autoregressive lags is chosen by maximizing the BIC. Column (1) has only the contemporaneous monetary policy shock, while column (2) includes the first lag of the positive shock, and column (3) includes lagged values for both the positive and negative shock. We can reject symmetry at either the 10% or 5% significance level for all specifications. The sample is monthly data from March 1969 to December 2008.

the results are virtually unaffected by using 12 autoregressive lags). The results are in table 3.4. On the whole, the positive shocks are much larger in magnitude and more statistically significant. We can reject symmetry at either the 10% or 5% significance level depending on the specification.

The prediction that these shocks should have asymmetric effects on wage inflation is a purely nominal implication of this model that is not generated by alternative theories of asymmetric business cycles like the ones driven by financial frictions that only bind in recessions.

As a further check on these results, I conduct a placebo test by replacing wage inflation

with headline CPI inflation and report the results in table 3.5. The model predicts that price inflation should exhibit smaller asymmetries than wage inflation in response to monetary shocks. In table 3.5, the point estimates in the specifications for contemporaneous positive and negative monetary shocks are virtually identical, including specifications with no autoregressive lags or differing numbers of lags for the monetary shock, and the hypothesis that positive and negative shocks have the same impact cannot be rejected.

3.4.3 Cross-sectional distribution of wage changes

Next, we look at the cross-sectional distribution of wages in different countries and time periods. Dickens *et al.* (2007) demonstrate that the cross-sectional distribution of wage changes is both skewed towards the right, and exhibits bunching at zero.
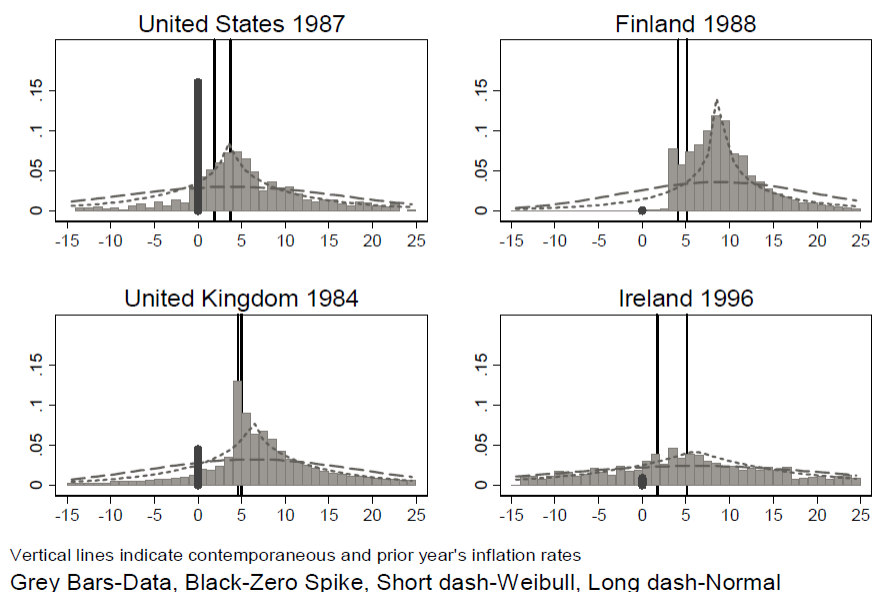
Figure 3.5: *Dickens et al (2007) cross sectional distribution of wages in different countries and time periods.*

Dickens *et al.* (2007) identify two forms of wage rigidity. The first, which they call "nominal rigdity", is a large point mass at zero wage change, the second, which they call "real rigidity", is an asymmetric distribution of wage changes around the average inflation

Table 3.5: *Responsiveness of Price Inflation to Monetary Policy Shocks*

	(1) π_t	(2) π_t	(3) π_t
ε_t^+	0.073	0.065	0.052
	(0.07)	(0.07)	(0.07)
ε_t^-	-0.069	-0.056	-0.077
	(0.08)	(0.09)	(0.09)
ε_{t-1}^+		0.059	0.021
		(0.07)	(0.07)
ε_{t-1}^-			0.134
			(0.10)
Constant	Yes	Yes	Yes
Number of autoregressive lags	4	4	4
Observations	475	475	475

Newey-West t statistics in parentheses

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

Notes: Columns regress headline annual CPI inflation on positive and negative monetary policy shocks and other covariates. CPI inflation data is from the Federal Reserve Economic Database. The monetary policy shocks are the structural shocks from Coibion and Gorodnichenko (2012). The number of lags of the inflation rate are chosen to mimic the ones in table 3.4, though the results are robust to including more lags. Column (1) has only the contemporaneous monetary policy shock, while column (2) includes the first lag of the positive shock, and column (3) includes lagged values for both the positive and negative shock. We cannot reject the hypothesis of symmetry. The sample is monthly data from March 1969 to December 2008.

rate. An illustration of this can be seen in figure 3.5, taken from Dickens *et al.* (2007). We see that the UK in 1984, where both current and last years' inflation rates were around five percent, had an asymmetry at five percent. On the other hand, in countries where the inflation rates were lower and less stable, the asymmetry was at zero.

The model presented in section 3.2 has a degenerate cross-sectional distribution, but it can easily be extended to have cross-sectional heterogeneity. Following Lucas (1973), consider a continuum of islands indexed by elements of the $[0,1]$ interval, with each island inhabited by a worker and an employer. Worker-employer pair i observe a noisy island specific signal $s_i = \varepsilon + \varepsilon_i$ of the log price level $p = \mu + \varepsilon$, and then write a wage contract. Crucially, we assume that the worker considers the signal-to-noise ratio to be ambiguous. As before, when worker utility is log, the prevailing log wage w in island i is

$$w_i = \mu + \frac{\sigma^2}{\sigma^2 + \sigma_s^2(s_i)} s_i + d.$$

where $\sigma_s(s_i) = \overline{\sigma}_s$ whenever $s_i \leq 0$, and $\underline{\sigma}_s$ otherwise. This means that the cross-sectional distribution of wages will be discontinuous around the expected price level μ, with higher variance on the right-hand side and bunching on the left-hand side. If the aggregate shock ε is sufficiently large, the discontinuity and bunching disappear. That is, when the unexpected monetary shock is large (high surprise inflation), no asymmetry is observed in the cross-sectional distribution, see figure 3.6.

Crucially, the model implies that the key point of asymmetry is the ex-ante expected price level (the price level before the signal is observed). This means that, in this model, there is nothing special about zero per se, unless households believe that they are in a very low inflation environment and that absent any signal, prices are not going to change. Conversely, in environments with high and stable inflation, we would observe an asymmetry not at zero, but around expected inflation. Where the asymmetry appears will depend on household expectations in the absence of any new information. It is difficult to infer this from the data available, but the empirical evidence presented in Dickens *et al.* (2007) is consistent with the idea that the location of asymmetry is higher than zero in environments with persistently
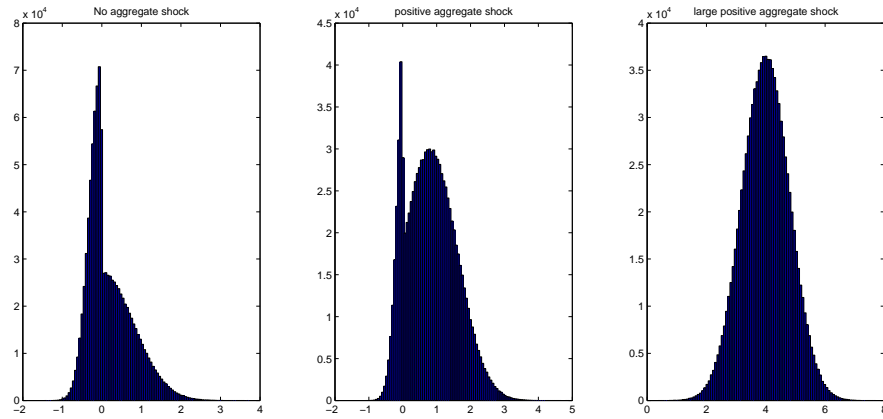
Figure 3.6: *Theoretical cross-sectional wage-change distributions for different aggregate monetary shocks.*

high inflation. Furthermore, if the model presented in this paper is augmented with agents who exhibit asymmetric money illusion (as argued by Fehr and Tyran (2001)) then we can capture both the asymmetry around the expected inflation rate and the bunching at wage-freezes using only the price-expectations of households without invoking loss-aversion or fairness norms.

A further finding by Dickens *et al.* (2007), as well as Holden and Wulfsberg (2009), is that the degree of real wage rigidity is strongly correlated with union density. Again, the basic model in section 3.2 can naturally be extended to account for this finding. In particular, note that the degree of rigidity does not solely depend on the degree of ambiguity, but also on the relative elasticities of labor supply and labor demand – or, loosely speaking, the bargaining power of workers and employers. The kinked beliefs of the workers only affect their wage to the extent that workers can withdraw labor in response to their perceived real wage. In particular, the model in section 3.2 assumed that the workers' outside option is exogenous and constant. This makes labor supply completely elastic, since workers effectively make an ultimatum to the the employers and refuse to supply any labor when wages are lower than what they demand. Consider, instead the following log labor supply curve

$$l = \gamma(w - \tilde{E}(p|s) - d),$$

where l is log labor and the elasticity of labor supply is given by γ. When γ tends to infinity, we recover the previous set up. On the other hand, when $\gamma = 0$, labor supply is completely inelastic at fixed supply. Let log labor demand be given by

$$l = \frac{1}{1-\alpha}(\log(\alpha) + E(p|s) - w),$$

easily derived from profit maximization with a Cobb-Douglas production function. The equilibrium log wage is given by

$$w = \kappa\left(\log(\alpha) + E(p|s)\right) + (1-\kappa)\left(\tilde{E}(p|s) + d\right),$$

where $\kappa = \frac{1}{\gamma(1-\alpha)+1} \in [0,1]$.

The equilibrium nominal wage is a convex combination of the beliefs of the workers and the beliefs of the employers. In the extreme case of infinitely elastic labor supply, $\gamma = \infty$, only the beliefs of the workers matters and we get maximum rigidity. In the other extreme of completely inelastic labor supply, $\gamma = 0$, only the beliefs of the employer matter. In particular, in the completely inelastic labor supply case, ambiguity has no effect on equilibrium wages. This is intuitive: if workers cannot withdraw their labor in response to the wage offer, then wages are determined solely through competition between employers. If employers are not adversely affected by inflation, then the equilibrium wage will not exhibit a discontinuity around the expected inflation rate even if employers are ambiguity-averse.

So, the degree of bargaining power, captured here by the workers' ability to withdraw labor when wages fall, affects the degree of rigidity in wages. This is consistent with the empirical findings of Dickens *et al.* (2007) who find that countries where unions have more power exhibit more wage rigidity of the kind generated by this model (i.e. asymmetry around the expected inflation rate).

3.4.4 Is Ambiguity-Aversion Necessary?

While ambiguity aversion is consistent with the evidence I provided, it is not the only theory that could generate these predictions. In particular, any theory that delivers convex

conditional expectations will generate similar predictions. The simplest alternative theory is a Bayesian expected-utility maximizer. Since professional forecasters exhibit no convexity in their beliefs, there must then be something unusual about the priors of the households. In particular, priors that assume the signal is more accurate when there is accelerations of the inflation rate than decelerations will generate very similar results to mine. However, these priors will be incorrect and hard to justify intuitively. Furthermore, the professional forecasters will be objectively doing a better job.

The virtue of ambiguity-aversion, other than its tractability, is that that there is no sense in which households are acting irrationally or making a mistake. If households do not know the precise mapping between aggregate statistics and the prices relevant for them, it is reasonable for them to rely on robustness heuristics – indeed, this was the original motivation for axiomatic theories of ambiguity-aversion.

3.5   Normative Analysis

One of the advantages of having a microfoundation for downward wage rigidity is that it allows us to deal with normative questions in a more satisfying manner. In this section 3.5.1, I derive optimal monetary policy, and in section 3.5.2, I analyse the welfare costs of business cycles.

3.5.1   Optimal Monetary Policy

It is typically assumed that downward wage rigidity implies that the central bank should have an inflationary bias. Inflation is said to "grease the wheels" of the labor market since it allows wage cuts to take place that would otherwise not have occurred. In this section, I show that this intuition holds in my model if we take the conditional expectation function of the households as exogenous, but fails if we account for the fact that household expectations will react to the change in policy.

Consider a scenario where the central bank has some, but not complete, control over the distribution of demand shocks that hit the economy. Crucially, suppose that although the

central bank can affect the distribution of shocks, it has no control over the distribution of the public signal. In other words, the central bank chooses a distribution of demand shocks to minimize expected losses, taking as given the information content (and ambiguity) of a noisy signal.

Most central banks are tasked with maintaining price stability and full employment. In a model like the one sketched above, with only aggregate demand shocks, price stability and deviations from first-best employment are both log-linear functions of the level of the monetary surprise. In particular, if we denote first-best employment by l^{fb}, then

$$l - l^{fb} = \frac{1}{1-\alpha}(m - \tilde{E}(m|s)),$$

where first-best employment is employment in the perfect information world. On the other hand, $m - \tilde{E}(m|s)$ also captures price instability. So, I assume that the central bank's loss function is given by

$$L(g) := E_g\left((m - \tilde{E}(m|s))^2\right),$$

where g is the marginal distribution of demand shocks m, and the expectation is taken with respect to g. I assume that if the central bank takes no action, demand shocks will have a reference distribution q. The central bank chooses g to minimize its losses subject to the requirement that g is not too different from q. I formalize "not too different" using the Kullback-Leibler divergence, an analytically tractable measure of difference between probability distributions.

Naive Policy

In this section, I consider the problem of a central bank who takes the function mapping signals to beliefs of the household as given, and does not internalize the fact that changing the distribution of demand shocks will affect how signals are mapped to conditional expectations. In other words, the central bank solves the following problem

$$\min_{g(m)} \iint (m - \phi(s))^2 f(s|m)g(m)\mathrm{d}s\mathrm{d}m$$

such that

$$\int q(m) \log(g(m)) \mathrm{d}m - \int q(m) \log(q(m)) \mathrm{d}m \leq K$$

$$\int g(m) \mathrm{d}m = 1,$$

where q is the distribution of demand shocks when the central bank is passive, $f(s|m)$ is the density of the signal conditional on the monetary policy shock, and $\phi(s)$ is household expectations of the demand shock conditional on the signal, which the bank takes as exogenous. The first constraint requires that the distribution of demand shocks the bank chooses be sufficiently close to the reference distribution q in Kullback-Leibler terms. The second constraint ensures that the chosen density implies a valid probability distribution. The slack non-negativity constraints have been suppressed since they are implied by the first constraint.

The Lagrangian is given by

$$\min_{g(m)} \iint (m - \phi(s))^2 f(s|m) g(m) \mathrm{d}s \mathrm{d}m - \lambda \int q(m) \log(g(m)) \mathrm{d}m - \mu \int g(m) \mathrm{d}m.$$

The first order condition is given by

$$\frac{\mathrm{d}}{\mathrm{d}t} \iint (m - \phi(s))^2 f(s|m)(g(m) + th(m)) \mathrm{d}s \mathrm{d}m$$

$$- \lambda \int q(m) \log(g(m) + th(m)) \mathrm{d}m - \mu \int g(m) + th(m) \mathrm{d}m \bigg|_{t=0} = 0, \quad \forall h$$

At the optimum, $g(m)$ solves the following equation

$$\int (m - \phi(s))^2 f(s|m) - \lambda \frac{q(m)}{g(m)} - \mu = 0.$$

Rearrange this to get

$$g(m) = \frac{\lambda q(m)}{\int \left(m - \tilde{E}(m|s)\right)^2 f(s|m) \mathrm{d}s - \mu}, \tag{3.13}$$

where $\tilde{E}(m|s)$ is substituted for $\phi(s)$.

This first order condition is very intuitive to interpret. Draws of the monetary shock m with large expected squared error in the household's forecast, $E((m - \tilde{E}(m|s))^2|m)$, are

129

less likely to occur relative to the reference distribution q. In other words, if households are more likely to have incorrect beliefs during deflationary episodes than inflationary episodes, then the central bank will reduce the probability of deflationary shocks. This is despite the fact that the central bank's loss function treats under- and over-employment symmetrically.

Since we found household beliefs to be more likely to be incorrect after disinflationary periods than inflationary periods, equation (3.13) suggests that the central bank should maintain an inflationary bias in policy. This is in keeping with the intuition, and the advice, found in papers like Akerlof *et al.* (1996) or Kim and Ruge-Murcia (2009), that recommend positive steady state inflation in the presence of downwardly sticky wages. However, in the context of this model, this line of reasoning is susceptible to the Lucas critique if the conditional expectations of households respond endogenously to the distribution of demand shocks.

Before proceeding to the case with endogenous expectations, let us get a better sense for how the solution behaves with the following numerical example. This example shows that, the naive optimal policy will feature an inflationary bias. Suppose that the signal s is given by

$$s = m + \varepsilon,$$

where ε is a mean-zero normally distributed noise term with variance $\sigma \in [\underline{\sigma}, \overline{\sigma}]$. Then

$$\tilde{E}(m|s) = \max_{\sigma \in [\underline{\sigma}, \overline{\sigma}]} \int m \frac{f(s|m,\sigma)g(m)}{f(s|\sigma)} \mathrm{d}m, \tag{3.14}$$

where

$$f(s|m,\sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(s-m)^2}{2\sigma^2}\right), \tag{3.15}$$

and

$$f(s|\sigma) = \int f(s|m,\sigma)g(m)\mathrm{d}m. \tag{3.16}$$

Equations (3.13), (3.14), (3.15), and (3.16) determine the equilibrium of this economy.

Let the reference distribution q be a standard normal distribution with mean 0 and variance 1. By calibrating $\underline{\sigma}, \overline{\sigma}$, and λ we can compute the equilibrium distribution of

monetary shocks. We can calibrate $\underline{\sigma}$ and $\overline{\sigma}$ by fitting a piecewise linear regression to the expectations data. The slope of the piecewise linear function ψ gives the signal-to-total variance ratio, which in turn pins down $\underline{\sigma}$ and $\overline{\sigma}$. A good estimate seems to be $\underline{\sigma} = 0.7$ and $\overline{\sigma} = 2$. Calibrating λ is harder, so we can plot solutions for a range of λ to get a sense of what the optimal solution looks like. In figure 3.7, we see that as the constraint on the central bank becomes looser, the distribution of shocks becomes more positive and concentrated. For comparison, figure 3.8 shows that without ambiguity, the distribution simply becomes more concentrated, but there is no inflationary bias. This lines up with the received wisdom that central banks should have an inflationary bias because of downward wage rigidity.
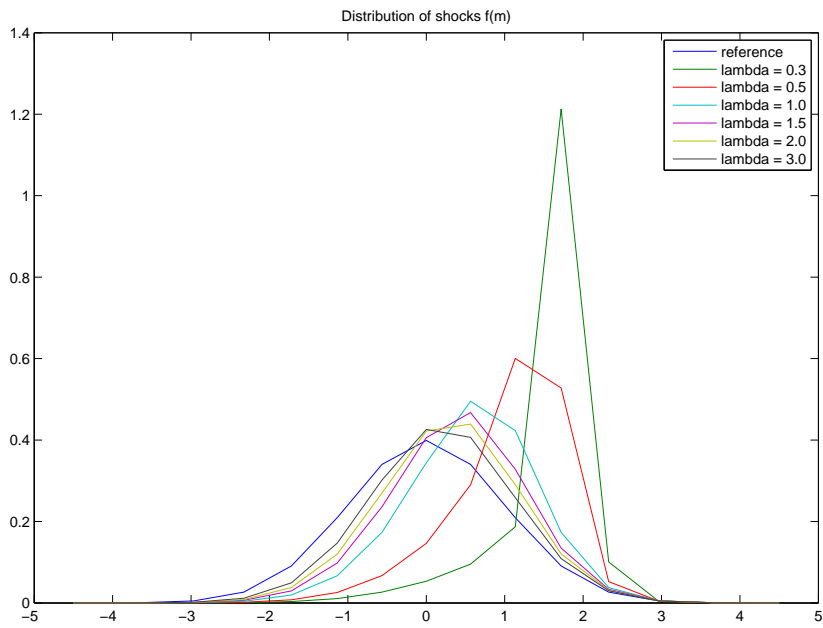
Figure 3.7: *The marginal distribution of monetary shocks*
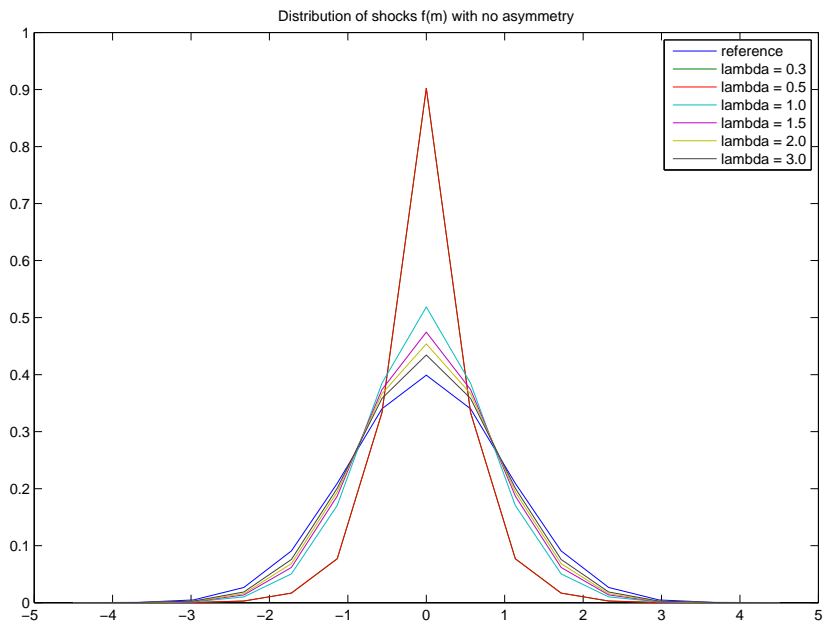
Figure 3.8: *The marginal distribution of monetary shocks with no ambiguity.*

Sophisticated Policy

The intuitive result in the previous section is in line with other work in recommending inflationary bias in the presence of downward wage rigidity. However, this result depends crucially on the assumption that the function mapping signals to conditional expectations for the households is fixed. If the central bank takes into account the fact that changing the distribution of monetary shocks changes the signal-extraction problem faced by households, then the inflationary bias disappears.

To that end, consider a central bank that faces the following problem:

$$\min_{g(m)} \iint (m - \phi(s))^2 f(s|m)g(m)\mathrm{d}s\mathrm{d}m$$

such that

$$\int q(m)\log(g(m))\mathrm{d}m - \int q(m)\log(q(m))\mathrm{d}m \leq K$$

$$\int f(m)\mathrm{d}m = 1,$$

$$\phi(s) = \max_{\sigma \in [\underline{\sigma}, \overline{\sigma}]} \int \frac{mf(s|m,\sigma)g(m)}{f(s|\sigma)}\mathrm{d}m.$$

The first order condition (omitted) for this problem is harder to interpret. Instead I plot example solutions using a normal error term and a normal reference distribution in figure 3.9. Unlike the previous section, we see no inflationary bias in the central bank's optimal response, even though the degree of the asymmetry is very extreme. The reason is that if the central bank attempted to skew the distribution towards more inflationary shocks, conditional expectations of households would take this skew into account when interpreting the signal.

The results of this section do not prove that zero percent inflation is the optimal inflation rate. In fact, in this model, the mean value of the inflation rate, as long as it is known by all agents, has no effect on welfare, since wages and prices are flexible. In practice, there are other reasons why we might want to implement a positive inflation target, ranging from concerns about hitting the zero lower bound to other causes of downward wage rigidity

besides the one studied here (for instance a nominal fairness norm).
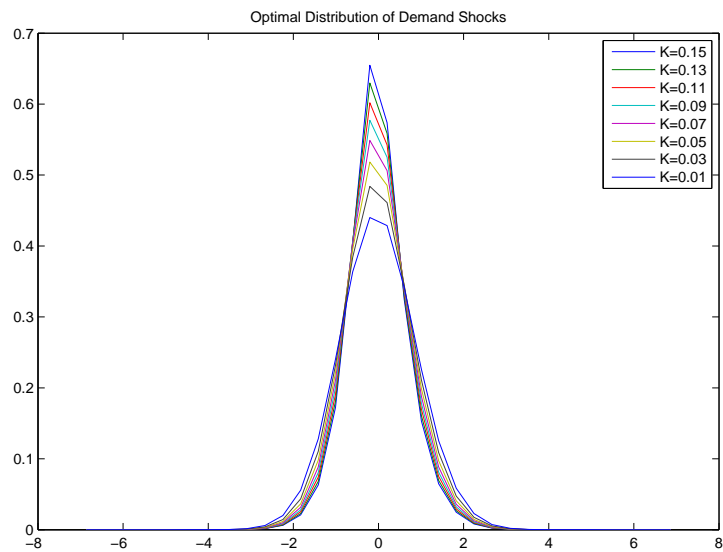
Figure 3.9: *The optimal distribution of monetary shocks.*
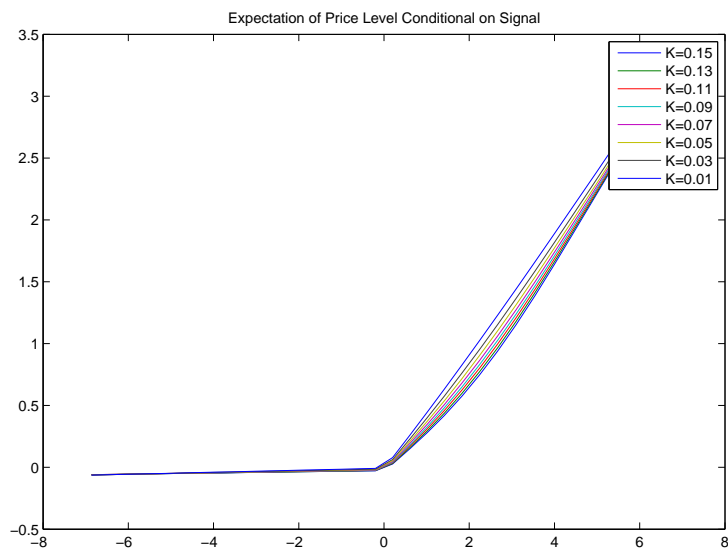
Figure 3.10: *The expectation of the price level conditional on the signal.*

3.5.2 Costs of Business Cycles

In this environment, demand shocks are more costly to welfare than in standard models of business cycles. Lucas (1987), in a highly influential study, performs a back-of-the-envelope calculation that implies that the welfare costs of ordinary business cycles, measured in units of life-time consumption, are extremely small (around one-twentieth of one percent). The basic intuition underlying this result is that negative shocks are cancelled out by positive shocks, resulting in second order gains from demand-management policies. However, as pointed out by Schmitt-Grohé and Uribe (2011), in a world with asymmetric rigidities, such calculations need not be true. In the present environment, as seen in figure 3.4, negative shocks cause far larger drops than positive shocks – therefore, demand-management policy can reap first-order gains.

To formalize this intuition, we can replicate the calculation in Lucas (1987) for the present model. Let

$$c_t^{det} = c_0 e^{gt}$$

represent a deterministic consumption path growing at rate g starting from c_0. Let the stochastic consumption stream be

$$c_t = c_0 e^{gt} \exp\left(\overline{\kappa}\mathbf{1}(\varepsilon_t \leq 0)\varepsilon_t + \underline{\kappa}\mathbf{1}(\varepsilon_t > 0)\varepsilon_t\right),$$

where ε_t is a Gaussian-$(0, \sigma^2)$ demand shock, and $\overline{\kappa} > \underline{\kappa}$ corresponds to the piecewise-linear slopes of (3.12). This is equilibrium consumption in a model like the one presented above that also features deterministic growth. To measure the welfare costs of demand shocks in permanent consumption units, set

$$\sum_{t=0}^{\infty} \beta^t u(\lambda c_t^{det}) = \mathrm{E}\left(\sum_{t=0}^{\infty} \beta^t u(c_t)\right),$$

and solve for λ. Note that the expectation on the right-hand side features no ambiguity-aversion, but is the objective probability distribution of $\{c_t\}_{t=0}^{\infty}$ given normally distributed demand shocks. Following Lucas (1987), assume CRRA utility with risk aversion parameter

γ. Then we can derive the following analytical expression for λ

$$(1+\lambda)^{1-\gamma} = \frac{1}{2}\left\{ \exp\left((\gamma-1)\bar{\kappa}(1+(\gamma-1)\bar{\kappa})\sigma^2\right)\left[1 + \Phi\left(\frac{1+2(\gamma-1)\bar{\kappa}\sigma}{2\sqrt{2}}\right)\right] + \right.$$
$$\left. \exp\left((\gamma-1)\underline{\kappa}(1+(\gamma-1)\underline{\kappa})\sigma^2\right)\left[1 - \Phi\left(\frac{1+2(\gamma-1)\underline{\kappa}\sigma}{2\sqrt{2}}\right)\right]\right\},$$

where Φ is the CDF of a standard normal distribution. Note that if we set $\bar{\kappa} = \underline{\kappa} = 1$, we recover the original calculation done by Lucas. To get a sense for how large this cost is, Lucas calibrates his model by letting $\gamma \in [1,4]$ and $\sigma = 0.032$. The standard deviation of the shocks are taken from the residuals of a linear regression. The present model implies that such a procedure would underestimate the variance of the true underlying shocks. This harks back to the debate between Romer (1986) and De Long and Summers (1988) about whether macroeconomic policy reduces the variance of shocks or fills in the troughs without shaving the peaks. In the table below, I take $\sigma = 0.032$ and $\bar{\kappa} = 1$, to make my results directly comparable with those of Lucas (1987), with the caveat, that re-estimating σ would result in even larger differences. In particular, the variance of the consumption process I specify is $(1+\underline{\kappa}^2)\sigma^2/2$ which is strictly less than the variance used by Lucas as long as $\underline{\kappa} < 1$.

Table 3.6: *The ratio of λ in this model to that in Lucas (1987).*

		γ			
		1	2	3	4
	0	25.28	25.62	25.98	26.34
	0.1	22.86	23.24	23.63	24.03
	0.2	20.45	20.87	21.3	21.74
	0.3	18.03	18.5	18.98	19.47
	0.4	15.6	16.13	16.67	17.21
$\underline{\kappa}$	0.5	13.18	13.77	14.36	14.97
	0.6	10.75	11.41	12.07	12.74
	0.7	8.315	9.048	9.787	10.53
	0.8	5.88	6.695	7.514	8.337
	0.9	3.441	4.345	5.251	6.159

The results of the calibration are in table 3.6. We see that even a modest amount of asymmetry can substantially increase the welfare costs of demand shocks for the US. In

particular, if we calibrate κ using table 3.1, then the welfare cost of demand-driven business cycles are approximately 1% percent of life-time consumption, or about 20 times the cost found by Lucas (1987). There are reasons to believe that these estimates are a lower bound on the welfare costs even in the context of a complete markets, representative consumer economy. First, according to the model, the variance of the underlying shocks in the data is larger than what one would estimate from the residuals of a least squares regression. Second, as we will see in the next section, in a dynamic model, ambiguous information can cause distortions in the deterministic steady-state of a linearized model and make negative shocks more persistent than positive shocks, further increasing the welfare costs of shocks in the model.

3.6   New Keynesian Model with ambiguous sticky wages

In this section, I embed ambiguous information quality into a standard New Keynesian model with sticky wages. Other than showing that our earlier intuitions survive in this context, I show that in a New Keynesian model time-varying ambiguity is observationally equivalent to a supply or cost-push shock. I also show that ambiguity not only causes the amplitude of positive and negative shocks to be asymmetric, but it can also change their persistence.

In a typical New Keynesian model with sticky wages, households are monopolistically-competitive suppliers of their labor. In such a world, it is no longer the case that inflation is bad news and disinflation is good news, since monopolists care about both the relative price and the quantity of what they sell. Therefore, I impose kinked beliefs on the households without deriving it from their preferences. This is an artifact of the way sticky wages are modelled in the New Keynesian model – in real life, most households do not set their own wages subject to downward sloping labor demand. One could get around this problem by having firms set wages instead, as in the earlier model, however, it is also interesting to put kinked household beliefs into the work-horse New Keynesian model since, independent from the microfoundations, earlier empirical results imply that household

inflation expectations are indeed kinked in the data.

Consider a continuum of households indexed by $i \in [0,1]$. Household seek to maximize

$$\tilde{\mathbb{E}}_t \left(\sum_{t=t_0}^{\infty} \beta^t \frac{C_{it}^{1-\gamma}}{1-\gamma} - \frac{L_{it}^{1+\varphi}}{1+\varphi} \right),$$

where $\tilde{\mathbb{E}}_t$ represents expectation with respect to kinked beliefs in period t. There is a representative firm that produces the consumption good Y using technology

$$Y_t = A_t L_t,$$

where A_t is a productivity shock and L_t is a CES aggregate of labor inputs

$$L_t = \left(\int_0^1 L_{it}^{1-\frac{1}{\eta}} di \right)^{\frac{\eta}{\eta-1}}.$$

This implies that labor demand is given by

$$l_{it} - l_t = -\eta(w_{it} - w_t), \tag{3.17}$$

where lower case variables are in logs, and w_t is the log of the CES wage aggregate. Assume that due to free-entry, the firm makes zero profits in equilibrium, therefore

$$w_t - p_t = a_t. \tag{3.18}$$

Assume that firms are subject to a cash in advance constraint for the labor they purchase

$$w_t + l_t = \theta_t, \tag{3.19}$$

where θ_t is a stochastic process representing the money supply. Let

$$\theta_t = \theta_{t-1} + v_t, \quad v_t \sim \mathcal{N}(0, \sigma_v^2).$$

Assume that agents receive a noisy public signal x_t of the stance of monetary at the start of the period

$$x_t = v_t + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2).$$

As before, households do not know the true signal-to-noise ratio. Suppose agents only know

that $\sigma_\varepsilon \in [\sigma_\varepsilon - d, \sigma_\varepsilon + d]$, where $d > 0$ is a parameter that captures the amount of ambiguity or Knightian uncertainty.

The timing of this model will be similar to Angeletos and La'O (2009): at the beginning of the period, an exogenous fraction $1 - \lambda$ of households set their wages optimally subject to their information set and Calvo frictions. At the end of the period, v_t becomes common knowledge, and consumption and production take place. I assume that a monopoly tax eliminates the markup.

The log-linearized optimal reset wage is given by

$$w_{it} = \tilde{E}_t \left[(1 - \beta\lambda) \sum_{k=0}^{\infty} (\beta\lambda)^k \left(mrs_{i,t+k|t} + p_{t+k} \right) \right], \qquad (3.20)$$

where $mrs_{i,t+k|t}$ is log marginal rate of substitution for household i at time $t + k$ conditional on the wage being set in t. Observe that

$$mrs_{i,t+k|t} = \gamma c_{it} + \varphi l_{it},$$

$$= \gamma c_t + \varphi l_{it},$$

$$= \gamma y_t + \varphi(l_t - \eta(w_{it} - w_t)),$$

$$= \gamma a_t + (\gamma + \varphi)l_t - \varphi\eta w_{it} + \varphi\eta w_t, \qquad (3.21)$$

where the second equality follows from complete insurance markets. Use (3.21), (3.18), and (3.19) to get

$$mrs_{i,t+k|t} + p_{t+k} = \underbrace{\frac{\gamma + \varphi}{1 + \varphi\eta}}_{\alpha} \theta_t + \underbrace{\frac{\varphi\eta - \gamma - \varphi + 1}{1 + \varphi\eta}}_{1-\alpha} w_t + \underbrace{\frac{\gamma - 1}{1 + \varphi\eta} a_t}_{\xi_t}. \qquad (3.22)$$

Note that

$$\tilde{E}_t \theta_t = \theta_{t-1} + \psi_t x_t,$$

where $\psi_t = \overline{\psi} = \sigma_v^2 / (\sigma_v^2 + \sigma_\varepsilon^2 - d)$ if $x_t \geq 0$, and $\psi_t = \underline{\psi} = \sigma_v^2 / (\sigma_v^2 + \sigma_\varepsilon^2 + d)$ if $x_t < 0$.

Conjecture an equilibrium where

$$w_{it} = b_1 w_{t-1} + b_2 \psi_t x_t + b_3 \theta_{t-1} + b_4 \xi_t + b_5. \qquad (3.23)$$

140

Observe that, from aggregation,

$$w_t = \lambda w_{t-1} + (1 - \lambda) w_{it},$$

$$= \lambda w_{t-1} + (1 - \lambda)(b_1 w_{t-1} + b_2 \psi_t x_t + b_3 \theta_{t-1} + b_4 \xi_t + b_5). \tag{3.24}$$

Note that we can write (3.20) recursively

$$
\begin{aligned}
w_{it} &= (1 - \beta\lambda)\, \tilde{E}_t \left(mrs_{it|t} + p_t \right) + \beta\lambda\, \tilde{E}_t \left(w_{i,t+1} \right). \\
&= (1 - \beta\lambda)\, \tilde{E}_t \left(\alpha\theta_t + (1 - \alpha)w_t + \xi_t \right) + \beta\lambda\, \tilde{E}_t \left(w_{i,t+1} \right), \\
&= (1 - \beta\lambda)\, \tilde{E}_t \left(\alpha\theta_t + (1 - \alpha)w_t + \xi_t \right) + \beta\lambda\, \tilde{E}_t \left(b_1 w_t + b_2 \psi_{t+1} x_{t+1} + b_3 \theta_t + b_4 \xi_{t+1} + b_5 \right).
\end{aligned}
$$

$$\tag{3.25}$$

Combine (3.25) and (3.24) and, by matching coefficients, derive expressions for b_1, b_2, b_3, b_4, and b_5. This requires noting that

$$\tilde{E}_t \, \xi_t = \xi_t, \quad E_t(\psi_{t+1} x_{t+1}) = \frac{\sigma_x}{\sqrt{2\pi}} (\overline{\psi} - \underline{\psi}),$$

where σ_x denotes the time-varying variance of x_{t+1}. Matching coefficients gives

$$b_1 = (1 - \alpha)(1 - \beta\lambda)(\lambda + (1 - \lambda)b_1) + \beta\lambda b_1((1 - \lambda)b_1 + \lambda),$$

$$b_2 = (1 - \beta\lambda)\alpha + b_2(1 - \alpha)(1 - \beta\lambda)(1 - \lambda) + b_2\beta\lambda b_1(1 - \lambda) + \beta\lambda b_3,$$

$$b_3 = (1 - \beta\lambda)\alpha + b_3(1 - \alpha)(1 - \beta\lambda)(1 - \lambda) + b_3\beta\lambda b_1(1 - \lambda) + \beta\lambda b_3,$$

$$b_4 = (1 - \beta\lambda) + b_4(1 - \alpha)(1 - \beta\lambda)(1 - \lambda) + b_4\beta\lambda b_1(1 - \lambda),$$

$$b_5 = b_5(1 - \alpha)(1 - \beta\lambda)(1 - \lambda) + b_5\beta\lambda b_1(1 - \lambda) + \beta\lambda b_5 + \beta\lambda \frac{\sigma_x(\overline{\psi} - \underline{\psi})}{\sqrt{2\pi}}.$$

Therefore, the wage-reset rule conjectured in (3.23) is an equilibrium. We can spot two differences between the reset wage (3.23) and a standard New Keynesian model with sticky wages. First is the presence of the asymmetric response of real variables to monetary shocks. Second, is the presence of the constant term b_5. This term is an ambiguity premium, and has the same interpretation as ambiguity premia in asset pricing contexts. Ambiguity-averse households try to insure themselves against monetary shocks in the future by setting higher

wages than they otherwise would. This results in a steady state level of real wages that is higher than, and output that is lower than, in the case with no ambiguity. This ambiguity premium has exactly the same implications as a mark-up and it generates a distortion of the steady state. Some impulse response functions can be seen in figures 3.11 and 3.12. As expected, negative shocks cause larger changes than positive shocks. The persistence of either type of shock, however, is identical, since after the first period, the shock becomes common knowledge.
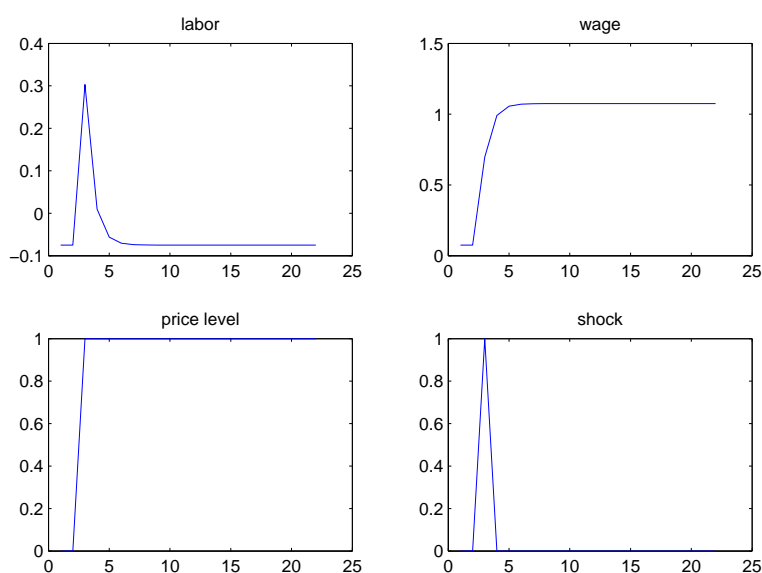
Figure 3.11: *The nominal wage and employment as a function of permanent positive shock to money supply.*

3.6.1 Time-varying ambiguity

Consider a world where the degree of Knightian uncertainty d is a time-varying quantity, suppose for example that d is a random walk. Then in equilibrium,

$$w_{it} = b_1 w_{t-1} + b_2 \psi_t x_t + b_3 \theta_{t-1} + b_4 \xi_t + b_5(d_t), \tag{3.26}$$

where b_5 is increasing in d_t. This results in cost-push shocks, which increase wages and reduce output, giving a new microfoundation for the existence of a meaningful policy
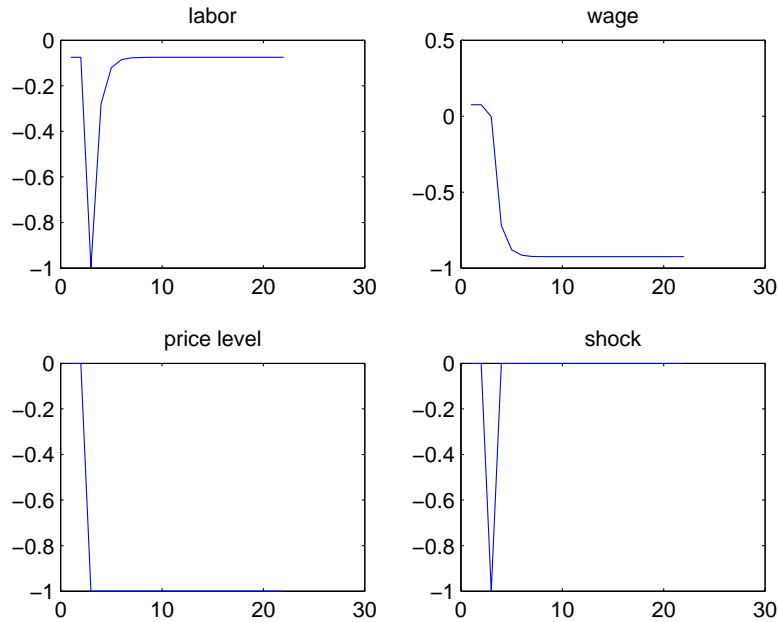
Figure 3.12: *The nominal wage and employment as a function of permanent negative shock to money supply.*

tradeoff between output and inflation for the central bank and violation of the divine coincidence.

3.6.2 New Keynesian Model with imperfect information and no Calvo frictions

If we eliminate the Calvo friction, the dynamics of the model become degenerate since there is perfect information at the end of each period. An alternative way of endowing the model with some persistence is to follow Woodford (2003). In this set up, ambiguity aversion not only makes the shocks asymmetric on impact, but it also changes their persistence. In particular, disinflationary signals take longer to be incorporated into agents' beliefs with the result that recessions are not just deeper, but also longer-lived, than booms.

Each period, agents receive a public signal x_t as before, but now, instead of the true state being revealed after one period, the true state is never revealed. On the other hand, we

dispense with the Calvo friction so that wages can be reset every period. Now

$$\tilde{E}_t(\theta_t) = \sum_{j=0}^{\infty} \left(\prod_{i=0}^{j-1} \psi_{t-i} \right) (1 - \psi_{t-j}) x_{t-j},$$

where ψ_t corresponds to the Kalman gain coefficient under worst case beliefs. We can substitute this into (3.24) to get

$$w_t = \frac{\alpha}{1-\alpha} \left[\sum_{j=0}^{\infty} \left(\prod_{i=0}^{j-1} \psi_{t-i} \right) (1 - \psi_{t-j}) x_{t-j} \right] + \frac{1}{1-\alpha} \xi_t.$$

We see that now, the shock not only affects the magnitude, but also the persistence of the shocks. In particular, negative shocks will on average take longer to be incorporated into the price, which in turn will result in more persistent declines in output. This increases the welfare costs of negative demand shocks.

3.7 Conclusion

In this paper, I argue that information frictions, coupled with ambiguity-aversion, can result in household expectations of the price level that are more sensitive to inflationary news than disinflationary news. The intuition is that households pay closer attention to and respond more strongly to bad news that their purchasing power might be lower than they thought than good news that their purchasing power is higher than they thought. I confirm that this asymmetry exists in survey data of household inflation expectations, and show that such asymmetric beliefs can give rise to downward rigidity in equilibrium wages. A simple general equilibrium model then implies that nominal and real variables respond asymmetrically to monetary policy shocks. In particular, negative monetary shocks cause larger changes to output than positive monetary shocks. On the other hand, negative monetary policy shocks cause smaller changes to wage inflation than positive monetary shocks. I show that these predictions hold in time series data from the United States.

Normatively, the asymmetry induced by ambiguity aversion increases the welfare costs of business cycles. Since positive and negative shocks do not cancel, reductions in variance

reap first-order gains. A back of the envelope calculation shows that these costs are around 20 times higher than the ones in Lucas (1987). Furthermore, it is typically assumed that downward wage rigidity should imbue the central bank with an inflationary bias, for example in Akerlof *et al.* (1996) or Kim and Ruge-Murcia (2009). However, this intuition fails to hold in my model if household conditional expectations respond endogenously to inflationary pressure from the central bank. In other words, the idea that in a world with downward wage rigidity, positive inflation "greases the wheels" of the labor market may be subject to the Lucas critique for reasons similar to the long-run Phillips curve. Finally, I embed this type of ambiguity aversion into a standard New Keynesian model with sticky wages and show that ambiguity about inflation is observationally equivalent to cost push shocks. So one does not need a supply-side, or markups-driven, story to derive a meaningful policy tradeoff between inflation and employment. Furthermore, in a dynamic model, ambiguity aversion means that disinflationary signals take longer to be incorporated into household beliefs and therefore demand-driven recessions are longer-lived than demand-driven booms.

References

ACEMOGLU, D., AUTOR, D., DORN, D., HANSON, G. H. and PRICE, B. (2013). *Import Competition and the Great US Employment Sag of the 2000s*. Tech. rep., Mimeo.

—, CARVALHO, V. M., OZDAGLAR, A. and TAHBAZ-SALEHI, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, **80** (5), 1977–2016.

AKERLOF, G. A. (1982). Labor contracts as partial gift exchange. *The Quarterly Journal of Economics*, **97** (4), 543–569.

—, DICKENS, W. T., PERRY, G. L., GORDON, R. J. and MANKIW, N. G. (1996). The macroeconomics of low inflation. *Brookings Papers on Economic Activity*, **1996** (1), pp. 1–76.

ANGELETOS, G.-M. and LA'O, J. (2009). Incomplete information, higher-order beliefs and price inertia. *Journal of Monetary Economics*, **56**, S19–S37.

ANGRIST, J. D., JORDÀ, s. and KUERSTEINER, G. (2013). Semiparametric estimates of monetary policy: String theory revisited.

ATALAY, E. (2013). How important are sectoral shocks? *Manuscript, University of Chicago*.

BAQAEE, D. R. (2014a). *Asymmetric Inflation Expectations, Downward Rigidity of Wages, and Asymmetric Business Cycles*. Working paper series, SSRN.

— (2014b). *Cascading Failures in Production Networks*. Tech. rep.

— (2014c). Labor intensity in an interconnected economy.

BARATTIERI, A., BASU, S. and GOTTSCHALK, P. (2010). *Some Evidence on the Importance of Sticky Wages*. Working Paper 16130, National Bureau of Economic Research.

BERMAN, E., BOUND, J. and GRILICHES, Z. (1994). Changes in the demand for skilled labor within us manufacturing: Evidence from the annual survey of manufacturers. *The Quarterly Journal of Economics*, pp. 367–397.

BETTENDORF, L. and HEIJDRA, B. J. (2003). *The monopolistic competition revolution in retrospect*.

BEWLEY, T. F. (1999). *Why wages don't fall during a recession*. Harvard University Press.

CAMPBELL, J. Y. and MANKIW, N. G. (1990). Permanent income, current income, and consumption. *Journal of Business & Economic Statistics*, **8** (3), 265–279.

146

CHODOROW-REICH, G., FEIVESON, L., LISCOW, Z. and WOOLSTON, W. G. (2012). Does state fiscal relief during recessions increase employment? evidence from the american recovery and reinvestment act. *American Economic Journal: Economic Policy*, **4** (3), 118–145.

CHRISTIANO, L., EICHENBAUM, M. and REBELO, S. (2011). When is the government spending multiplier large? *Journal of Political Economy*, **119** (1), 78–121.

COIBION, O. and GORODNICHENKO, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, **120** (1), 116–159.

CORREIA, I., FARHI, E., NICOLINI, J. P. and TELES, P. (2013). Unconventional fiscal policy at the zero bound. *American Economic Review*, **103** (4), 1172–1211.

COVER, J. P. (1992). Asymmetric effects of positive and negative money-supply shocks. *The Quarterly Journal of Economics*, **107** (4), 1261–1282.

DALY, M. and HOBIJN, B. (2013). Downward nominal wage rigidities bend the phillips curve. *FRBSF Working Paper*, (8).

DE LONG, J. B. and SUMMERS, L. H. (1988). How does macroeconomic policy affect output? *Brookings Papers on Economic Activity*, **1988** (2), pp. 433–494.

DEARDORFF, A. V. (2006). *Glossary of international economics*. World Scientific.

DICKENS, W. T., GOETTE, L., GROSHEN, E. L., HOLDEN, S., MESSINA, J., SCHWEITZER, M. E., TURUNEN, J. and WARD, M. E. (2007). How wages change: Micro evidence from the international wage flexibility project. *Journal of Economic Perspectives*, **21** (2), 195–214.

EGGERTSSON, G. and MEHROTRA, N. (2014). *A Model of Secular Stagnation*. Tech. rep.

EGGERTSSON, G. B. and KRUGMAN, P. (2012). Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach*. *The Quarterly Journal of Economics*, **127** (3), 1469–1513.

ELLIOTT, M., GOLUB, B. and JACKSON, M. (2012). Financial networks and contagion. *Available at SSRN 2175056*.

ELSBY, M. W., HOBIJN, B. and ŞAHIN, A. (2013). The decline of the us labor share. *Brookings Papers on Economic Activity*, **2013** (2), 1–63.

EPSTEIN, L. G. and SCHNEIDER, M. (2008). Ambiguity, information quality, and asset pricing. *The Journal of Finance*, **63** (1), 197–228.

ESTRADA, N. and VALDEOLIVAS, E. (2012). *The Fall of the Labour Income Share in Advanced Economies*. Occasional Paper 1209, Bank of Spain.

FARHI, E. and WERNING, I. (2012). *Fiscal multipliers: liquidity traps and currency unions*. Tech. rep., National Bureau of Economic Research.

— and — (2013). *A theory of macroprudential policies in the presence of nominal rigidities*. Tech. rep., National Bureau of Economic Research.

Fehr, E. and Tyran, J.-R. (2001). Does money illusion matter? *American Economic Review*, pp. 1239–1262.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, **79** (3), 733–772.

Galí, J., López-Salido, J. D. and Vallés, J. (2007). Understanding the effects of government spending on consumption. *Journal of the European Economic Association*, **5** (1), 227–270.

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, **18** (2), 141 – 153.

Goldin, C. D. and Katz, L. F. (2009). *The race between education and technology*. Harvard University Press.

Goolsbee, A. D. and Krueger, A. B. (2015). *A Retrospective Look at Rescuing and Restructuring General Motors and Chrysler*. Working Paper 588, Princeton University.

Hall, R. E. (2005). Employment fluctuations with equilibrium wage stickiness. *American economic review*, pp. 50–65.

Hansen, L. P. and Sargent, T. J. (2011). *Robustness*. Princeton university press.

Holden, S. and Wulfsberg, F. (2009). Wage rigidity, institutions, and inflation.

Holmstrom, B. (1983). Equilibrium long-term labor contracts. *The Quarterly Journal of Economics*, **98**, pp. 23–54.

Hulten, C. R. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies*, pp. 511–518.

Ilut, C. and Schneider, M. (2012). *Ambiguous business cycles*. Tech. rep., National Bureau of Economic Research.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics*, **107** (1), 35–78.

Kaur, S. (2012). Nominal wage rigidity in village labor markets.

Kempe, D., Kleinberg, J. and Tardos, É. (2003). Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, ACM, pp. 137–146.

Kim, J. and Ruge-Murcia, F. J. (2009). How much inflation is necessary to grease the wheels? *Journal of Monetary Economics*, **56** (3), 365 – 377.

Krugman, P. R. (1998). It's baaack: Japan's slump and the return of the liquidity trap. *Brookings Papers on Economic Activity*, pp. 137–205.

Kuhnen, C. M. (2012). Asymmetric learning from financial information.

Leontief, W. W. (1936). Quantitative input and output relations in the economic systems of the united states. *The Review of Economic Statistics*, pp. 105–125.

Long, J. B. and Plosser, C. I. (1983). Real business cycles. *The Journal of Political Economy*, pp. 39–69.

Lucas, J., Robert E (1973). Some international evidence on output-inflation tradeoffs. *American Economic Review*, **63** (3), 326–34.

Lucas, R. E. (1987). *Models of business cycles*, vol. 26. Basil Blackwell Oxford.

Malinvaud, E. (1977). *The theory of unemployment reconsidered*, vol. 1985. Blackwell Oxford.

Morris, S. (2000). Contagion. *The Review of Economic Studies*, **67** (1), 57–78.

Neiman, B. and Karabarbounis, L. (2014). The global decline of the labor share. *The Quarterly Journal of Economics*, **129** (1), 61–103.

Oberfield, E. and Raval, D. (2012). Micro data and the macro elasticity of substitution. *US Census Bureau Center for Economic Studies Paper No. CES-WP-12-05*.

Patinkin, D. (1965). Money, interest, and prices.

Piketty, T. (2014). Capital in the 21st century.

Romer, C. D. (1986). Is the stabilization of the postwar economy a figment of the data? *The American Economic Review*, **76** (3), pp. 314–334.

— and Romer, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, pp. 1055–1084.

Scheinkman, J. A. and Woodford, M. (1994). Self-organized criticality and economic fluctuations. *The American Economic Review*, pp. 417–421.

Schelling, T. C. (1971). Dynamic models of segregation†. *Journal of mathematical sociology*, **1** (2), 143–186.

Schmitt-Grohé, S. and Uribe, M. (2011). *Pegs and Pain*. Tech. rep., National Bureau of Economic Research.

Shiller, R. J. (1997). Why do people dislike inflation? In *Reducing Inflation: Motivation and Strategy*, University of Chicago Press, pp. 13–70.

Shoag, D. (2010). The impact of government spending shocks: Evidence on the multiplier from state pension plan returns.

Stolper, W. F. and Samuelson, P. A. (1941). Protection and real wages. *The Review of Economic Studies*, **9** (1), 58–73.

Timmer, M., Erumban, A., Gouma, R., Los, B., Temurshoev, U., de Vries, G. and Arto, I. (2012). The world input-output database (wiod): contents, sources and methods. *WIOD Background document available at www. wiod. org*.

TREFLER, D. and ZHU, S. C. (2010). The structure of factor content predictions. *Journal of International Economics*, **82** (2), 195–207.

VALENTINYI, A. and HERRENDORF, B. (2008). Measuring factor income shares at the sectoral level. *Review of Economic Dynamics*, **11** (4), 820–835.

WERNING, I. (2011). *Managing a Liquidity Trap: Monetary and Fiscal Policy*. Tech. rep., National Bureau of Economic Research.

WOODFORD, M. (2003). Imperfect common knowledge and the effects of monetary policy. *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Strother Phelps*, p. 25.

WOODS, J. (1979). A note on hicks' composite commodity theorem. *Zeitschrift für Nationalökonomie*, **39** (1-2), 185–188.

YANG, W., FIDRMUC, J. and GHOSH, S. (2012). *Government Spending Shocks and the Multiplier: New Evidence from the U.S. Based on Natural Disasters*. Tech. rep.

Appendix A

Appendices to Chapter 1

A.1   Appendix I: Proofs

Proof of theorem 1.2.1. The fact that labor's share of income is equal to 1 follows trivially from considering the aggregate budget constraint:

$$wl - \tau + \tau = \sum_i p_i(c_i + g_i).$$

To see that the distribution of government expenditures does not affect equilibrium employment, consider the cost minimization problem of firm i

$$c(x_i; \boldsymbol{p}, w) = \min_{x_{ij}, l_i} \left\{ \sum_j p_j x_{ij} + w l_i : F_i(x_{i1}, \dots, x_{in}, l_i) = x_i \right\}.$$

Once again, note that since the problem of the firm is static, I have suppressed time subscripts. Note that for any $\alpha > 0$,

$$
\begin{aligned}
c_i(\alpha x_i; \boldsymbol{p}, w) &= \min_{x_{ij}, l_i} \left\{ \sum_j p_j x_{ij} + w l_i : F_i(x_{i1}, \ldots, x_{in}, l_i) = \alpha x_i \right\}, \\
&= \min_{x_{ij}, l_i} \left\{ \sum_j p_j x_{ij} + w l_i : F_i(x_{i1}/\alpha, \ldots, x_{in}/\alpha, l_i/\alpha) = x_i \right\}, \\
&= \min_{x_{ij}, l_i} \left\{ \sum_j \alpha p_j \frac{x_{ij}}{\alpha} + w\alpha \frac{l_i}{\alpha} : F_i(x_{i1}/\alpha, \ldots, x_{in}/\alpha, l_i/\alpha) = x_i \right\}, \\
&= \alpha \min_{x_{ij}, l_i} \left\{ \sum_j p_j \frac{x_{ij}}{\alpha} + w\frac{l_i}{\alpha} : F_i(x_{i1}/\alpha, \ldots, x_{in}/\alpha, l_i/\alpha) = x_i \right\}, \\
&= \alpha c_i(x_i; \boldsymbol{p}, w).
\end{aligned}
$$

So the marginal cost of firm i is

$$
\frac{\partial c_i}{\partial x_i} = c_i(1; \boldsymbol{p}, w).
$$

In other words, the marginal cost of firm i depends only on the wage and the prices of firm i's inputs.

In equilibrium, all firms must make zero profits, otherwise they would expand their size to infinity or shrink to zero. In particular, this means that price must equal marginal cost for each good

$$
p_i = c_i(1; \boldsymbol{p}, w). \tag{A.1}
$$

Furthermore, observe that p_i scale one-for-one with the wage w. That is, if $\boldsymbol{p}$ solves (A.1), then $\tilde{\boldsymbol{p}} \equiv \alpha \boldsymbol{p}$ solves

$$
\tilde{p}_i = c_i(1; \tilde{\boldsymbol{p}}, \alpha w).
$$

So (A.1) implies that all prices are pinned down by technologies and the nominal wage. So let $\overline{\boldsymbol{p}}$ solve the following equations:

$$
\overline{p}_i = c_i(1; \overline{\boldsymbol{p}}, 1),
$$

and note that any equilibrium price vector must be $w\overline{\boldsymbol{p}}$.

At the steady-state equilibrium with zero inflation, $1 + i_{t+1} = \frac{\rho_t}{\rho_{t+1}}$. This implies that

consumption of each good and labor are the same at every period. To see this, observe that the household's problem can be written as

$$\max \quad \sum_{\tau=t}^{\infty} \rho_\tau U(c_{1\tau}, \ldots, c_{n\tau}, l_\tau),$$

subject to

$$\sum_{s=1}^{\infty} \left(\sum_i p_{it+s} c_{it+s} + \tau_{t+s} \right) \left(\prod_{\tau=0}^{s-1} \frac{1}{1+i_{t+\tau}} \right) = \sum_{s=1}^{\infty} w_{t+s} l_{t+s} \left(\prod_{\tau=0}^{s-1} \frac{1}{1+i_{t+\tau}} \right),$$

where profits have been dropped since they are always equal to zero. The first order conditions, along with the assumption that $1 + i_{t+1} = \rho_t / \rho_{t+1}$ implies that

$$u_{it}(c_{1t} \ldots, c_{nt}, l_t) = u_{it+1}(c_{1t+1} \ldots, c_{nt+1}, l_t + 1),$$

for every t and i, where u_{it} is the marginal utility of good i in period t. Furthermore,

$$u_{lt}(c_{1t} \ldots, c_{nt}, l_t) = u_{lt+1}(c_{1t+1} \ldots, c_{nt+1}, l_t + 1).$$

These relations imply that consumption and labor supplied are the same in every period. This means that we can collapse the household problem into finding just the steady-state values of consumption and labor. In particular, we simply need to solve

$$\max \quad u(c_1, \ldots, c_n, l)$$

subject to

$$\sum_i p_i c_i = wl - \tau,$$

since perfect consumption-smoothing means that there is no variation across periods.

Define

$$V(C, l) = \max \left\{ u(c_1, \cdots, c_N, l) : \sum \overline{p}_i c_i = C \right\}.$$

Note that the indirect utility of the household, defined as the solution to the household's problem (1.1) will coincide with the solution to

$$\max_{C, l} \quad V(C, l)$$

153

such that

$$wC = wl - \tau$$

The first order condition to this problem is

$$-\frac{V_l(l - \tau/w, l)}{V_c(l - \tau/w, l)} = \frac{w}{w}, \tag{A.2}$$

Now let labor be the numeraire so that $w = 1$. Then this equation pins down l. We see that the only way in which government policy appears in these two expressions is via the size of the government's budget τ. In particular, the distribution of spending is irrelevant. ∎

The assumption of constant returns to scale and marginal cost pricing allow us to use the logic of the Hicks-Leontief composite commodity theorem, see for example Woods (1979).

An extension of the logic of theorem 1.2.1 to demand shocks is possible if we assume that the indirect utility function of the household $V(C, l)$ is quasi-linear in the composite consumption good.

Corollary. *Suppose that*

$$V(C, l) = const\, C - v(l),$$

then equilibrium employment depends only on the disutility of labor.

Proof. The assumption of quasi-linearity allows us to write (A.2) as

$$v'(l) = const.$$

∎

In particular, corollary A.1 implies that changes in the utility function of the household (demand-shocks) do not affect equilibrium employment so long as the indirect utility function remains quasi-linear in consumption. The leading example of this scenario is when the utility function is Cobb-Douglas or has the CES form in consumption. In that case, changes to the share parameters will not affect equilibrium employment.

Proof of Lemma 1.3.2.

$$\mathbf{1} = I\Omega\mathbf{1},$$

$$= (\mathrm{diag}(a+c) + \mathrm{diag}(1-a-c))W\mathbf{1},$$

$$= \mathrm{diag}(a+c)W\mathbf{1} + \mathrm{diag}(1-a-c)W\mathbf{1},$$

$$= a + c + \mathrm{diag}(1-a-c)W\mathbf{1}.$$

Rearrange this to get

$$(I - \mathrm{diag}(1-a-c)W\mathbf{1} = a+c,$$

or

$$\mathbf{1} = (I - \mathrm{diag}(1-a-c)W)^{-1}(a+c).$$

∎

Proof of Proposition 1.3.3. Let s_i denote the sales of firm i. By definition,

$$p_j x_{ij} = w_{ij} s_i,$$

and

$$wl_i = a_i s_i.$$

Goods market clearing implies that

$$y_i = c_i + g_i + \sum_j x_{ji},$$

or equivalently

$$s_i = p_i c_i + p_i g_i + \sum s_j w_{ji}.$$

Solve this system of linear equations for s to get

$$s' = (H+G)'(I-W)^{-1},$$

where H is household expenditures net of taxes and G is government expenditures.

Finally, labor market clearing implies that

$$wl = \sum_i wl_i = s'a. \qquad \text{(A.3)}$$

Hence,

$$wl = (H + G)'(I - W)^{-1}a = (H + G)'\tilde{a}. \qquad \text{(A.4)}$$

∎

Proof of Proposition 1.3.4. For simplicity, assume the government's budget is zero. Note that labor market clearing requires that

$$l^{\theta-1} = \frac{w}{PC}.$$

Rearrange this to get

$$l^\theta = \frac{wl}{PC} = \frac{s'\alpha}{PC} = \frac{\beta'\Psi\alpha PC}{PC} = \beta'\tilde{\alpha}.$$

Now simply differentiate implicitly with respect to the taste parameter β_i to get

$$\theta l^{\theta-1} \frac{\mathrm{d}l}{\mathrm{d}\beta_i} = \left(\frac{e_i}{\sum_j \beta_j} - \frac{\beta}{(\sum_j \beta_j)^2} \right)' \tilde{\alpha}.$$

Divide this by the same expression for β_j to get the desired result:

$$\frac{\mathrm{d}l/\mathrm{d}\beta_i}{\mathrm{d}l/\mathrm{d}\beta_j} = \frac{e_i'\tilde{\alpha} - \beta'\tilde{\alpha}}{e_j'\tilde{\alpha} - \beta'\tilde{\alpha}}.$$

∎

Proof of Proposition 1.5.1. For simplicity, let the consumption taxes be equal to zero so that there is only lump-sum taxation. Firm cost minimization implies that

$$p_j x_{ij} = (1 - \alpha_i - \eta_i)\omega_{ij} p_i x_i,$$

and

$$rk_i = \eta_i p_i y_i,$$

and

$$wl_i = \alpha p_i y_i.$$

Substitute firm i's demand for inputs into its production function to get that

$$p_i = \left(\frac{w}{\alpha_i}\right)^{\alpha_i} \left(\frac{r}{\eta_i}\right)^{\eta_i} \prod_j \left(\frac{(1 - \alpha_i - \eta_i)\omega_{ij}}{p_j}\right)^{(1 - \alpha_i - \eta_i)\omega_{ij}}. \tag{A.5}$$

Note that

$$\log(p_i) = \alpha_i(\log(w) - \log(\alpha_i)) + \eta_i(\log(r) - \log(\eta_i)) + \sum_j (1 - \alpha_i - \eta_i)\omega_{ij}(\log(p_j) - \log(\omega_{ij})).$$

Rearrange this to get

$$\log(p) = (I - \hat{\Omega})^{-1}(\alpha \log(w) + \eta \log(r) - \Theta), \tag{A.6}$$

where

$$\hat{\Omega} = \text{diag}(1 - \alpha - \eta)\Omega,$$

and

$$\theta_i = \alpha_i \log(\alpha_i) + \eta_i \log(\eta_i) + \sum_j (1 - \alpha_i - \eta_i)\omega_{ij} \log((1 - \alpha_i - \eta_i)\omega_{ij}).$$

In the expressions above, logs of a vector or matrix are taken element by element. Equation (A.6) is informative, since it implies that the relative prices of consumption goods in the economy depend solely on the relative cost of the two factors and the technology.

Let $s_i = p_i y_i$. Then labor market clearing implies that

$$l = \sum_i l_i = \frac{s'\alpha}{w}. \tag{A.7}$$

Rearrange this to get

$$w = \frac{s'\alpha}{l}.$$

Plug this into the labor supply equation to get

$$l = \min\left\{ \left[\frac{s'\alpha}{lPC}\right]^{\frac{1}{\theta-1}}, \bar{l} \right\}.$$

Rearrange this to get

$$l = \min\left\{ \left[\prod \left(\frac{\beta_i}{p_i}\right)^{\beta_i} \frac{s'\alpha}{C}\right]^{\frac{1}{\theta}}, \bar{l} \right\}. \tag{A.8}$$

Similarly, capital market clearing implies that

$$K = \sum_i k_i = \frac{s'\eta}{r}. \tag{A.9}$$

Finally, goods market clearing implies that

$$y_i = c_i + g_i + \sum_j x_{ji}.$$

By market clearing,

$$p_i y_i = p_i(c_i + g_i) + \sum p_j x_j (1 - \alpha_j - \eta_j)\omega_{ji}$$

Denote the vector of s_i's by s. Then

$$s' = (H + G)'(I - \hat{\Omega})^{-1}, \tag{A.10}$$

where H is the vector of household expenditure and G is the vector of government expenditure.

Assume that $l \leq \bar{l}$ and substitute (A.10) into (A.8) to get

$$l = \left[(H + G)'(I - \hat{\Omega})^{-1}\alpha\right]^{\frac{1}{\theta}} \left(\frac{1}{PC}\right)^{\frac{1}{\theta}}. \tag{A.11}$$

So, the equilibrium labor is given by (A.11), with prices that satisfy (A.6), (A.9), and the normalization $w = 1$.

By changing the shares of government expenditures, the government affects equilibrium employment through three different channels. First, the government directly changes the demand for labor through its purchases. Second, the government changes demand for labor by affecting the price of labor relative to capital, and therefore the relative prices of more and less labor intensive goods. Lastly, the government changes the income of households. Fortunately, all three of these forces can be expressed as multiples of the relative network-adjusted labor intensities of the various sectors. This makes the clean expression in (1.7) possible.

Let labor to be the numeraire $w = 1$ to form the following expression

$$l^\theta = \frac{(\beta'(l + rK - \tau) + \delta'\tau)\,\Psi\alpha}{PC},$$

note that equilibrium employment l, rental rate of capital r, and household expenditures PC all depend on δ. Implicitly differentiate this expression with respect to δ_i

$$\theta l^{\theta-1}\frac{\mathrm{d}l}{\mathrm{d}\delta_i} = \left(\beta'\Psi\alpha\left(\frac{\mathrm{d}l}{\mathrm{d}\delta_i} + K\frac{\mathrm{d}r}{\mathrm{d}\delta_i}\right) + (e_i - \delta)'\Psi\alpha\tau\right)\frac{1}{PC} \tag{A.12}$$

$$- \frac{1}{(PC)^2}\left(\frac{\mathrm{d}l}{\mathrm{d}\delta_i} + K\frac{\mathrm{d}r}{\mathrm{d}\delta_i}\right)(\beta'PC + \delta'\tau)\,\Psi\alpha, \tag{A.13}$$

where e_i denotes the vector with zeros everywhere except the ith element. From (A.9), note that

$$K\frac{\mathrm{d}r}{\mathrm{d}\delta_i} = \beta'\Psi\eta\left(\frac{\mathrm{d}l}{\mathrm{d}\delta_i} + K\frac{\mathrm{d}r}{\mathrm{d}\delta_i}\right) + (e_i - \delta)'\Psi\eta.$$

Rearrange this to get $K\frac{\mathrm{d}r}{\mathrm{d}\delta_i}$ and substitute that into (A.12). After some rearranging, we get

$$PC\left[\theta l^{\theta-1} - (\beta'\Psi\alpha - \beta'\Psi\eta)\frac{1}{PC} + \frac{l}{(PC)^2}\frac{1}{\beta'\Psi\alpha}\right]\frac{\mathrm{d}l}{\mathrm{d}\delta_i} =$$
$$\left[(e_i - \delta)'\tau\Psi(\eta + \alpha) - \frac{1}{PC}l(e_i - \delta)'\Psi\eta\tau\frac{1}{\beta'\Psi\alpha}\right].$$

Divide the above expression for i by the same expression for k to get

$$\frac{\mathrm{d}l/\mathrm{d}\delta_i}{\mathrm{d}l/\mathrm{d}\delta_k} = \frac{(e_i - \delta)'\tau\left(\Psi\eta + \Psi\alpha - \frac{s'\alpha}{PC}\frac{1}{\beta'\Psi\alpha}\Psi\eta\right)}{(e_k - \delta)'\tau\left(\Psi\eta + \Psi\alpha - \frac{s'\alpha}{PC}\frac{1}{\beta'\Psi\alpha}\Psi\eta\right)}. \tag{A.14}$$

By lemma 1.3.2, note that $\Psi\eta = \mathbf{1} - \Psi\alpha$. Using this, we can simplify (A.14) to be

$$\frac{\mathrm{d}l/\mathrm{d}\delta_i}{\mathrm{d}l/\mathrm{d}\delta_k} = \frac{(e_i - \delta)'(\Psi\eta + \Psi\alpha - \frac{s'\alpha}{PC}\frac{1}{\beta'\Psi\alpha}\Psi\eta)}{(e_k - \delta)'(\Psi\eta + \Psi\alpha - \frac{s'\alpha}{PC}\frac{1}{\beta'\Psi\alpha}\Psi\eta)}$$

$$= \frac{(e_i - \delta)'(\mathbf{1} - \frac{s'\alpha}{PC}\frac{1}{\beta'\Psi\alpha}(\mathbf{1} - \Psi\alpha))}{(e_k - \delta)'(\mathbf{1} - \frac{s'\alpha}{PC}\frac{1}{\beta'\Psi\alpha}(\mathbf{1} - \Psi\alpha))},$$

$$= \frac{(e_i - \delta)'\Psi\alpha}{(e_k - \delta)'\Psi\alpha}.$$

The final line follows from the fact that

$$e_i' 1 \frac{s'\alpha}{PC} \frac{1}{\beta'\Psi\alpha} = \delta' 1 \frac{s'\alpha}{PC} \frac{1}{\beta'\Psi\alpha}.$$

This completes the proof. ∎

Proof of Proposition 1.5.2. For simplicity, assume that the government only uses consumption taxes and ensures a balanced budget with lump-sum taxes so there are no government purchases. Then, as before, the sales industry i are given by

$$p_i y_i = p_i c_i + \sum_j p_i x_{ji}.$$

Substituting household and firm input demands we get

$$p_i y_i = \frac{\beta_i PC}{1 + \tau_i} + \sum_j \omega_{ji} p_j y_j.$$

Denote the vector of sales by s and household expenditure share on good i net of taxes by $\beta_i^* = \beta_i / (1 + \tau_i)$. Then

$$s' = (\beta^*)' \Psi PC.$$

Letting labor be the numeraire, labor demand is then given by

$$l = s'\alpha = (\beta^*)' \Psi \alpha PC = (\beta^*)' \tilde{\alpha} PC.$$

Substitute this into the labor supply equation to get equilibrium labor

$$l^\theta = (\beta^*)' \tilde{\alpha}.$$

Differentiate this expression with respect to τ_i to get

$$\theta l^{\theta-1} \frac{dl}{d\tau_i} = -\frac{\beta_i}{(1 + \tau_i)^2} \tilde{\alpha}_i.$$

Rearrange this expression and divide through by the same expression for j to get

$$\frac{dl/d\tau_i}{dl/d\tau_j} = \frac{\beta_i \tilde{\alpha}_i (1 + \tau_i)^2}{\beta_j \tilde{\alpha}_j (1 + \tau_j)^2}.$$

■

Proof of lemma 1.5.3. From the Euler equation, we have that

$$P_{t+1}C_{t+1} = \rho^*(1 + i_t)P_t C_t.$$

Assume an equilibrium where $C_{t+1} = \overline{C}$, where $\overline{C}$ is the long-run efficient steady state value of consumption. Then the euler equation and the aggregate budget set imply that

$$1 + i_t = \frac{P_{t+1}\overline{C}/\rho^*}{w_t l_t + r_t k - \tau_t}.$$

When the zero-lower bound is not binding, the central bank can ensure full employment and no inflation by setting the nominal rate equal to $1/\rho^*$. However, at the zero lower-bound, we have

$$\frac{P_{t+1}\overline{C}/\rho^*}{w_t l_t + r_t k - \tau_t} = 1. \tag{A.15}$$

Rearrange this for labor earnings to get

$$w_t l_t = \frac{P_{t+1}\overline{C}}{\rho^*} - r_t k + \tau_t.$$

Note that

$$r_t k = (1 - \beta\tilde{\alpha})(w_t l_t + r_t k - \tau_t) + (1 - \delta'\tilde{\alpha})\tau_t.$$

Rearrange this to get

$$r_t k = \frac{(1 - \beta'\tilde{\alpha})w_t l_t + (\beta - \delta)'\tilde{\alpha}\tau_t}{\beta'\tilde{\alpha}}, \tag{A.16}$$

and substitute it into (A.15) to get

$$w_t l_t = \frac{P_{t+1}\overline{C}}{\rho^*}\beta'\tilde{\alpha} - (\beta - \delta)'\tilde{\alpha}\tau_t.$$

So labor earnings today depend on private nominal consumption tomorrow, and government

policy today. Note that

$$
\begin{aligned}
P_{t+1}\overline{C} &= w_{t+1}\overline{l} + r_{t+1}k - \tau_{t+1} \\
&= w_{t+1}\overline{l} + \frac{1 - \beta'\tilde{\alpha}}{\beta'\tilde{\alpha}}w_{t+1}\overline{l} + \frac{(\beta - \delta_{t+1})'\tilde{\alpha}}{\beta'\tilde{\alpha}}\tau_{t+1} - \tau_{t+1}, \\
&= \frac{1}{\beta'\tilde{\alpha}}w_{t+1}\overline{l} - \frac{\delta'_{t+1}\tilde{\alpha}}{\beta'\tilde{\alpha}}\tau_{t+1}.
\end{aligned}
$$

Substitute this into the previous expression to get

$$
w_t l_t = \frac{1}{\rho^*}w_{t+1}\overline{l} - \frac{1}{\rho^*}\delta'_{t+1}\tilde{\alpha}\tau_{t+1} - (\beta - \delta_t)'\tilde{\alpha}\tau_t. \tag{A.17}
$$

This gives the aggregate demand curve for labor. ∎

Proof of Proposition 1.5.4. Since the government cannot commit to future policy τ_{t+1} and δ_{t+1}, we see that the only way the government can boost employment is via $(\delta_t - \beta)'\tilde{\alpha}\tau_t$.

The household Euler equation at the zero lower bound implies that

$$
P_{t+1}C_{t+1} = \rho^*(1 + i_t)P_tC_t = \rho^*P_tC_t.
$$

The economy is back to its efficient full-employment steady state in period $t+1$. So $C_{t+1} = \overline{C}$. Note that as long as current government spending is not so high that it crowds out the private sector from the labor market, the equilibrium features $w_t = w_{t+1}$. Due to lack of commitment for fiscal policy, equation (A.16) pins down the rental rate of capital in period $t+1$ in terms of w_t:

$$
r_{t+1}k = \frac{1 - \beta'\tilde{\alpha}}{\beta'\tilde{\alpha}}w_t\overline{l} + \frac{(\beta - \phi)'\tilde{\alpha}}{\beta'\tilde{\alpha}}\overline{\tau} = \frac{1 - \beta'\tilde{\alpha}}{\beta'\tilde{\alpha}}w_t\overline{l} + \frac{(\beta - \phi)'\tilde{\alpha}}{\beta'\tilde{\alpha}}\frac{\lambda}{1 - \lambda}P_{t+1}\overline{C}.
$$

Therefore, P_{t+1} is also pinned down at its long-run steady state value. Since both C_{t+1} and P_{t+1} do not respond to the shock, this means that in order for the Euler equation to hold, either P_t needs to fall or C_t needs to fall.

Since in period $t+1$, the economy returns to full employment, the government's problem can be separated in two. In period $t+1$, the government spends

$$
p_{i,t+1}g_{i,t+1} = \frac{\lambda}{1 - \lambda}\phi_i P_{t+1}C_{t+1}.
$$

162

In period t, however, the government's problem is different because there is idle labor. The government's problem is

$$\max_{\delta, \tau_t} (1 - \lambda) \log \left(\frac{P_t C_t}{P_t} \right) + \lambda \sum_i \phi_i \log(g_{it}).$$

The Euler equation pins down $P_t C_t$ to be $\overline{PC}/\rho^*$, where $\overline{PC}$ is nominal GDP from period $t + 1$ onwards. So we can rewrite the government's problem in period t as

$$\max_{\delta, \tau_t} -(1 - \lambda) \log(P_t) + \lambda \sum_i \phi_i \log(g_{it}).$$

subject to the following constraints

$$w_t l_t = \beta' \Psi \alpha \frac{\overline{PC}}{\rho^*} + \delta' \Psi \alpha \tau_t, \tag{A.18}$$

$$r_t K = \beta' \Psi \eta \frac{\overline{PC}}{\rho^*} + \delta' \Psi \eta \tau_t, \tag{A.19}$$

$$w_{t-1} \leq w_t, \tag{A.20}$$

$$(l_t - \bar{l})(w_t - w_{t-1}) = 0, \tag{A.21}$$

$$g_{it} = \frac{\delta_i \tau_t}{p_{it}}, \tag{A.22}$$

$$p_{it} = w_t^{\tilde{\alpha}_i} r_t^{\tilde{\eta}_i} \text{const}_i, \tag{A.23}$$

$$P_t = \prod_i^N \left(\frac{p_{it}}{\beta_i} \right)^{\beta_i}, \tag{A.24}$$

$$\sum_i \delta_i = 1. \tag{A.25}$$

The first order condition for δ_i is given by

$$\frac{\lambda \phi_i}{\delta_i} + \mu_1 \tilde{\alpha}_k \tau_t + \mu_2 \tilde{\eta}_k \tau_t + \mu_8 = 0,$$

where μ_1 is the lagrange multiplier on the labor market condition, μ_2 is the lagrange multiplier on the capital market condition, and μ_8 makes sure that the δ_i sum to 1. Rearrange this to get

$$\frac{\delta_i}{\delta_j} = \frac{\phi_i}{\mu_2 + (\mu_1 - \mu_2)\tilde{\alpha}_i + \mu_8} \frac{\mu_2 + (\mu_1 - \mu_2)\tilde{\alpha}_j + \mu_8}{\phi_j}. \tag{A.26}$$

When the labor market clears, μ_1 and μ_2 are equal, therefore, the government simply

equates the marginal returns to various forms of government expenditures. However, during recession μ_2 exceeds μ_1, so that expenditures are tilted in favor of sectors with relatively high network-adjusted labor intensities.

To see this, note that the optimum features $w_t = w_{t-1}$ and $l_t = \bar{l}$. Substituting these into labor market clearing (A.18) for t and $t-1$ gives

$$\beta'\tilde{\alpha}\overline{PC}/\rho^* + \delta'\tilde{\alpha}\tau_t = \beta'\tilde{\alpha}\overline{PC} + \frac{\lambda}{1-\lambda}\overline{PC}\phi'\tilde{\alpha}.$$

Rearrange this to get

$$\delta'\tilde{\alpha} = \beta'\tilde{\alpha}(1-1/\rho^*)\overline{PC} + \frac{\lambda}{1-\lambda}\overline{PC}\phi'\tilde{\alpha}. \tag{A.27}$$

Substitute this into the first order condition for w_t to get

$$
\begin{aligned}
\mu_1 + \frac{\mu_3}{l_t} &= \frac{(1-\lambda)\beta'\Psi\alpha + \lambda\phi'\Psi\alpha}{w_t l_t} - \mu_4\frac{(l_t - \bar{l})}{l_t}, \\
&= \frac{(1-\lambda)\beta'\Psi\alpha + \lambda\phi'\Psi\alpha}{w_t l_t}, \\
&= \frac{(1-\lambda)\beta'\Psi\alpha + \lambda\phi'\Psi\alpha}{\beta'\Psi\alpha\overline{PC}/\rho^* + \delta'\Psi\alpha\tau_t}, \\
&= \frac{(1-\lambda)\beta'\Psi\alpha + \lambda\phi'\Psi\alpha}{\beta'\Psi\alpha\overline{PC} + \frac{\lambda}{1-\lambda}\overline{PC}\phi'\Psi\alpha}, \\
&= \frac{1-\lambda}{\overline{PC}}.
\end{aligned}
$$

Note that the first order condition for τ_t implies that

$$\mu_2\tau_t + (\mu_1 - \mu_2)\delta'\tilde{\alpha}\tau_t = \lambda.$$

Substitute (A.27) into this to get

$$\mu_2\tau_t\delta'\tilde{\eta} = \lambda\phi'\tilde{\eta} - (1-\lambda)\beta'\tilde{\alpha}(1-1/\rho^*). \tag{A.28}$$

Note that the first order equation for r_t gives

$$
\begin{aligned}
\mu_2 &= \frac{(1-\lambda)\beta'\tilde{\eta} + \lambda\phi'\tilde{\eta}}{r_t K}, \\
&= \frac{(1-\lambda)\beta'\tilde{\eta} + \lambda\phi'\tilde{\eta}}{\beta'\tilde{\eta}\overline{PC}/\rho^* + \delta\tilde{\eta}\tau_t}, \\
&= \frac{(1-\lambda)\rho^*}{\overline{PC}} + \frac{(1-\lambda)(\rho^* - 1)}{\overline{PC}}\frac{\beta'\tilde{\alpha}}{\beta'\tilde{\eta}},
\end{aligned}
$$

where going from the second to the third line requires substituting in (A.28).

Now note that

$$
\begin{aligned}
\mu_1 - \mu_2 &= \frac{1-\lambda}{\overline{PC}} - \frac{(1-\lambda)\rho^*}{\overline{PC}} - \frac{(1-\lambda)(\rho^* - 1)}{\overline{PC}}\frac{\beta'\tilde{\alpha}}{\beta'\tilde{\eta}} - \frac{\mu_3}{l_t}, \\
&= -\frac{(1-\lambda)(\rho^* - 1)}{\overline{PC}} - \frac{(1-\lambda)(\rho^* - 1)}{\overline{PC}}\frac{\beta'\tilde{\alpha}}{\beta'\tilde{\eta}} - \frac{\mu_3}{l_t}, \\
&< 0,
\end{aligned}
$$

as required. Note that $\mu_3 \geq 0$ by the KKT conditions.

We now see that the higher the labor-intensity of the household's consumption, the larger the required tilting by the government in its stimulus. It is also easy to verify that if the zero-lower bound does not bind ($\rho^* \leq 1$), then no intervention, and no tilting, is necessary. Indeed, when the zero-lower bound does not bind, we have that $\mu_1 = \mu_2$. ∎

Proof of lemma 1.5.5. I analyze the equilibrium where the steady-state equilibrium features zero inflation, full employment, and no government spending after the first period and constant government taxes starting in period $t + 1$. In period t, the intertemporal budget constraint of the saver and the Euler equation pin down his consumption in period t

$$
p_t c_t^s = \frac{p_{t+1} c_{t+1}^s}{\rho(1 + i_t)} = \frac{(w_{t+1}l_{t+1} + r_{t+1}K_{t+1})(1 - \chi) + (1 - \rho)\left[D^l \frac{p_{t+1}}{p_t} + B_t(1 + i_t) - p_t^g g_t(1 - \chi)\right]}{\rho(1 + i_t)}.
$$

The budget constraint for the borrower pins down his consumption in period t

$$
p_t c_t^b = (w_t l_t + r_t K_t)\chi + \frac{D^l}{1 + i_t}\frac{p_{t+1}}{p_t} - D^h \frac{p_t}{p_{t-1}}.
$$

To find the real rate of interest R_t, we solve for the equilibrium interest rate that equates

supply and demand. This requires that

$$c_t^s = \frac{w_t l_t + r_t k_t}{p_t} - c_t^b - \frac{p_t^g g_t}{p_t}.$$

Substituting the Euler equation into this and solving for the interest rate gives

$$1 + R_t = \frac{\frac{(w_{t+1}\bar{l} + r_{t+1}\bar{k})(1-\chi)}{p_{t+1}} + \frac{D^l}{p_{t+1}} - \frac{p_t^g g_t (1-\chi)(1-\rho)}{p_{t+1}}}{\rho\left(\frac{(w_t l_t + r_t \bar{k})(1-\chi)}{p_t} + \frac{D^h}{p_t}\right) - \frac{p_t^g g_t}{p_t}}.$$

The natural rate of interest, on the other hand, is given by

$$1 + R_t^n = \frac{\frac{(w_{t+1}\bar{l} + r_{t+1}\bar{k})(1-\chi)}{p_{t+1}} + \frac{D^l}{p_{t+1}} - \frac{p_t^g g_t (1-\chi)(1-\rho)}{p_{t+1}}}{\rho\left(\frac{(w_t \bar{l} + r_t \bar{k})(1-\chi)}{p_t} + \frac{D^h}{p_t}\right) - \frac{p_t^g g_t}{p_t}}.$$

The no-arbitrage condition between nominal government bonds and household debt implies the Fisher equation

$$1 + R_t = (1 + i_t)\frac{p_t}{p_{t+1}}.$$

If $R_t^n > 0$, then the central bank can maintain full employment with zero inflation, and government spending is unambiguously bad since it wastes resources that would otherwise be going to the households. We can see that whether or not the zero lower bound binds depends on the inflation rate, the amount of government spending, and the size of the deleveraging shock. When the central bank is able to set $i_t = R_t^n$, we have full employment as in figure A.1.
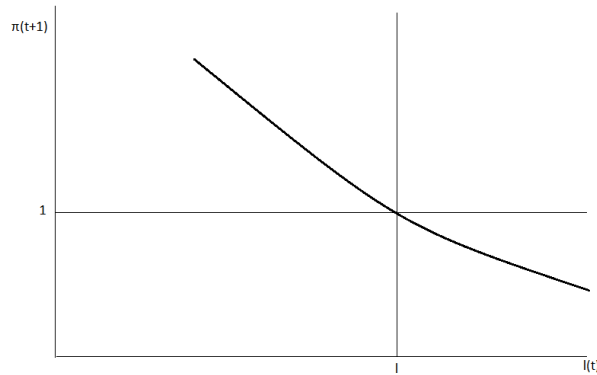
Figure A.1: *Zero-lower bound is not binding and aggregate demand is downward sloping*

However, when the zero lower bound on the nominal interest rate binds, then the Fisher equation implies that

$$(1 + R_t)\frac{p_{t+1}}{p_t} = 1.$$

Substituting the expression we have for R_t and solving for $w_t l_t$ we get

$$w_t l_t = \frac{1}{\rho}\left(w_{t+1}\bar{l} + r_{t+1}\bar{k}\right) + \frac{D^l - \rho D^h}{\rho(1-\chi)} + \frac{1 - (1-\rho)(1-\chi)}{\rho(1-\chi)}p_t^g g_t - r_t\bar{k}.$$

Now substitute

$$r_t\bar{k} = \frac{(1 - \beta'\tilde{\alpha})}{\beta'\tilde{\alpha}}w_t l_t + \frac{\beta'\tilde{\alpha} - \delta'\tilde{\alpha}}{\beta'\tilde{\alpha}}p_t^g g_t \qquad (A.29)$$

into this expression and solve for $w_t l_t$ to get

$$w_t l_t = \frac{1}{\rho}\beta'\tilde{\alpha}\left(w_{t+1}\bar{l} + r_{t+1}\bar{k}\right) + \beta'\tilde{\alpha}\frac{D^l - \rho D^h}{\rho(1-\chi)} + \left(\beta'\tilde{\alpha}\left[\frac{1 - (1-\rho)(1-\chi)}{\rho(1-\chi)} - 1\right] + \delta'\tilde{\alpha}\right)p_t^g g_t.$$

$$(A.30)$$

∎

Proof of proposition 1.5.6. Societal welfare is given by real GDP net of government expenditures. That is,

$$\frac{w_t l_t + r_t\bar{k} - p_t^g g_t}{p_t} = \frac{1}{\rho}\frac{\left(w_{t+1}\bar{l} + r_{t+1}\bar{k}\right) + \frac{D^l - \rho D^h}{\rho(1-\chi)} + \left[\frac{1-(1-\rho)(1-\chi)}{\rho(1-\chi)} - 1\right]p_t^g g_t}{c_1 w_t\left(\frac{1-\beta'\tilde{\alpha}}{\beta'\tilde{\alpha}}\frac{l_t}{\bar{k}} + \frac{\beta'\tilde{\alpha}-\delta'\tilde{\alpha}}{\beta'\tilde{\alpha}}\frac{p_t^g g_t}{\bar{k}w_t}\right)^{1-\beta'\tilde{\alpha}}}, \qquad (A.31)$$

where the numerator comes from combining (A.29) with (1.10) and c_1 is a constant. The denominator comes from combining (A.29) with the household's price index

$$p_t \propto \prod\left(w_t^{\tilde{\alpha}_i}r_t^{1-\tilde{\alpha}_i}\right)^{\beta_i}.$$

Equation (A.31) gives real GDP in period t net of government consumption, so it warrants close inspection. The term

$$\left[\frac{1 - (1-\rho)(1-\chi)}{\rho(1-\chi)} - 1\right] > 0$$

in the numerator is the government multiplier on nominal private GDP. This term does not depend on the composition of government spending since income from either factor

167

is distributed between the two households uniformly. The denominator, which gives the price level, does however the depend on the composition of spending, since government purchases of capital-intensive goods directly crowds out the household.

The numerator of (A.31) does not depend on the composition of spending, so we can focus on minimizing the denominator. Note that w_t is fixed as long as the zero lower bound is binding, so we can substitute l_t using (1.10) into the denominator of (A.31) and treat w_t as a constant, to get

$$\frac{1}{p_t} \propto \left[\frac{1 - \beta'\tilde{\alpha}}{\beta'\tilde{\alpha}} \left(\frac{c_2}{w_t \overline{k}} + \delta'\tilde{\alpha} \frac{p_t^g g_t}{w_t \overline{k}} \right) + p_t^g g_t - \frac{\delta'\tilde{\alpha}}{\beta'\tilde{\alpha}} \frac{p_t^g g_t}{\overline{k} w_t} \right]^{\beta'\tilde{\alpha}-1} = \left(c_3 - \beta'\tilde{\alpha}\delta'\tilde{\alpha} p_t^g g_t \right)^{\beta'\tilde{\alpha}-1},$$

where c_3 and c_2 are constants not affected by $p_t^g g_t$ or δ. This is maximized when the inner term is minimized because the exponent is less than one. The inner term is minimized when $\delta'\tilde{\alpha}$ is maximized, as required. ∎

A.2 Appendix II: World Trade

In this section, I augment the model in section 1.3 with trade in goods and services but immobile labor and capital. Assume capital is inelastically supplied at quantity K. The household chooses

$$\max \quad \sum_{t=0}^{\infty} \rho^t \left(\log(C_t) - \frac{l_t^\theta}{\theta} \right),$$

where

$$C_t = \prod_{i=1}^{N} c_{it}^{\beta_i},$$

where

$$c_{it} = \left(c_{it}^h \right)^{\kappa_i} \left(c_{it}^f \right)^{1-\kappa_i},$$

subject to budget constraint

$$\sum \left(p_{it}^h c_{it}^h + p_{it}^f c_{it}^f \right) + q_t B_t = w_t l_t + r_t K + B_{t-1} + \Pi_t - \tau_t,$$

where B_t is a nominal bond (in zero net supply), Π_t is firm profits, τ_t is lump sum taxes, r_t is the rental rate of capital, and c_{it}^h is quantity of domestically produced good i and c_{it}^f is quantity of foreign produced good i consumed. Suppose that there is a physical limit to the number of hours that can be worked

$$l_t \leq \bar{l}.$$

A.2.1 Firms' problem

Firms rent capital and labor on spot markets from the household, and reoptimize every period. Therefore, their problems are static, so I suppress time-subscripts. Since, in a competitive equilibrium with constant returns to scale, firm size is indeterminate, I simply state the problem of the representative firm in industry i:

$$\max_{y_i^h, l_i, x_{ij}} \quad p_i^h y_i^h - \sum_j p_j x_{ij}^h - \sum_j p_j^f x_{ij}^f - w l_i - r k_i$$

subject to the production function

$$y_i^h = l_i^{\alpha_i} k_i^{\eta_i} \prod x_{ij}^{(1-\alpha_i-\eta_i)\omega_{ij}},$$

with

$$x_{ij} = \left(x_{ij}^h\right)^{\kappa_{ij}} \left(x_{ij}^f\right)^{1-\kappa_{ij}},$$

where superscript h denotes domestic and f foreign use of input j by firm i. The network structure of the economy is captured by the parameters of the Cobb-Douglas production function.

A.2.2 Government behavior

Let government run balanced budgets every period

$$\sum p_i g_i = \tau. \tag{A.32}$$

169

Also, define the fraction of government expenditures on industry i to be

$$\delta_i = \frac{p_i^h g_i}{\sum_i p_i^h g_i}.$$

To keep the notation simple, the government is assumed to only make domestic purchases.

A.2.3 The Rest of the World

The rest of the world's behavior is treated as being exogenous. The world simply spends its earnings from trading with home on buying goods and services from home, so that

$$\sum_i p_{it}^h e_{it} = \sum_j \left(\sum_i p_{jt}^f x_{jt}^f + p_{jt}^f c_{jt}^f \right),$$

where $e_i t$ is exports of good i to foreign.

A.2.4 Market Clearing

The market for good or service i clears so that

$$p_{it}^h e_{it} + p_{it}^h c_{it}^h + p_{it}^h g_{it} + \sum_j p_{it}^h x_{jit}^h = p_{it}^h y_{it}^h.$$

The variable d_{it} is foreign demand for good i in domestic currency. The expenditures of foreigners on each good and service is treated as being exogenous. This would follow from a Cobb-Douglas utility function for the rest of the world.

A.2.5 Equilibrium

I will focus on the steady state of this model, and will therefore suppress time-subscripts.

Definition A.2.1. The steady-state competitive equilibrium of this economy is a collection of prices $\{p_i^h\}_{i=1}^N$, wage w, quantities $\{x_{ij}, x_{ij}^h, x_{ij}^f, c_i, c_i^h, c_i^f\}_{ij}$, and labor supply l and labor demands $\{l_i\}_i$ such that for a given δ and τ,

 (i) Each firm maximizes its profits given prices,

(ii) the representative household chooses consumption basket $\{c_i\}$ and labor supply l every period to maximizes utility,

(iii) the government runs a balanced budget,

(iv) and markets for each good, labor, and capital clear.

Definition A.2.2. The employment multiplier of government spending in industry i is defined as dl/dG_i, where l is equilibrium employment and G_i is government expenditures in industry i.

Since I am focusing on perturbations to the steady state of this model, the changes in government policy are permanent changes to the steady state of the model. This implies that government spending has very strong crowding-out effects.

Definition A.2.3. The relative employment multiplier of government spending in industry i is defined as $dl/d\delta_i$, where l is equilibrium employment and δ_i is the share of government expenditures in industry i, holding fixed the total size of the governments' budget.

The presence of trade with the rest of the world means that we must adjust the influence matrix for trade. To that end, let

$$\Psi^* \equiv (I - \text{diag}(1 - \alpha - \eta)(\kappa \cdot \Omega))^{-1},$$

represent the influence matrix with trade and ψ_{ij}^* represent the jth element of the ith row of this matrix. Similarly, let

$$\beta_i^* \equiv \beta_i \kappa_i,$$

for each i and β^* denote the column vector β_i^*.

Proposition A.2.1. *The relative government multiplier for shares of expenditure satisfy*

$$\frac{dl/d\delta_i}{dl/d\delta_k} = \frac{e_i' \Psi^* \alpha - \delta' \Psi^* \alpha}{e_k' \Psi^* \alpha - \delta' \Psi^* \alpha}. \tag{A.33}$$

Furthermore, labor's share of income is equal to

$$\frac{wl}{GDP} = \frac{(PC\beta^* + \tau\delta + E)' \Psi^* \alpha}{GDP},$$

171

where E is the vector of expenditures on foreign exports.

The proof for this proposition is very similar to that of proposition 1.3.3 and 1.5.1.

Proof. Let s_i denote sales of domestic industry i and note that

$$l = \sum_i l_i = \frac{s'\alpha}{w}.$$

Market clearing implies that

$$p_i^h y_i^h = p_i^h c_i + p_i^h g_i + p_i^h e_i + \sum_j p_j^h y_j^h (1 - \alpha_j - \eta_j)\kappa_{ji}\omega_{ji}.$$

Rearrange this to get

$$s' = (PC\beta^* + \tau\delta + E)'\Psi^*.$$

Equating Labor supply and labor demand gives

$$l = \left((PC\beta^* + \tau\delta + E)'\Psi^*\alpha\right)^{\frac{1}{\theta}} \left(\frac{1}{PC}\right)^{\frac{1}{\theta}}.$$

Take derivatives, rearrange, and use lemma 1.3.2 to get the desired result. ∎

A.3 Appendix III: Non-unit Elasticity of Substitution

A.3.1 Response to Demand Shocks

Recent work by Atalay (2013) suggests that at business cycle frequencies, the elasticity of substitution may be significantly less than one. In this section, I sketch how the optimal policy results can be generalized to cases with non-unitary elasticities.

Suppose that household utility is given by

$$\max_{c_i} \quad \left(\sum_{i=1}^N \beta_i^{\frac{1}{\varepsilon}} c_i^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}},$$

where ε is the elasticity of substitution, β_i is the share parameter for good i. The household faces the same budget set as before

$$\sum p_i c_i = wl + rk - \tau.$$

The government runs a balanced budget

$$\sum_i^N p_i g_i = \tau.$$

Government purchases are given by

$$g_i = \left(\frac{p_i}{P_G} \right)^{-\varepsilon} \delta_i G, \tag{A.34}$$

where

$$G = \left(\sum_i \delta_i^{\frac{1}{\varepsilon}} g_i^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{A.35}$$

and

$$P_G = \left(\sum_i^N \delta_i p_i^{1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}}. \tag{A.36}$$

The instrument of government policy are the share parameters δ_i. Observe that as $\varepsilon \to 1$ we recover the Cobb-Douglas production function as before.

The representative firm in each industry is competitive. It chooses inputs and prices to maximize profits:

$$\max_{p_i, l_i, k_i, x_{ij}} \quad p_i y_i - \sum_j p_j x_{ij} - w l_i - r k_i$$

using the production technology

$$y_i \leq \left(\alpha_i^{\frac{1}{\varepsilon}} l_i^{\frac{\varepsilon-1}{\varepsilon}} + \eta_i^{\frac{1}{\varepsilon}} k^{\frac{\varepsilon-1}{\varepsilon}} + \sum_i^N \omega_{ij}^{\frac{1}{\varepsilon}} x_{ij}^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

We have the following proposition:

Proposition A.3.1. *In the presence of downward sticky wages, when the labor market fails to clear, we have*

$$\frac{\mathrm{d}l/\mathrm{d}\delta_i}{\mathrm{d}l/\mathrm{d}\delta_i} = \frac{\tilde{\alpha}_i - \delta'\tilde{\alpha}}{\tilde{\alpha}_j - \delta'\tilde{\alpha}},$$

This proposition shows that the qualitative logic of the previous results carry-over without change to the case with non-unitary elasticities. Of course, quantitatively, the share parameters Ω and α no longer correspond to expenditure shares. For simplicity, I have assumed that the household's elasticity of substitution across goods is the same as the firms' elasticity of substitution across inputs. This assumption can be relaxed without losing

analytical tractability.

Proof. Note that the price vector P raised element-wise by $1 - \varepsilon$ satisfies

$$P^{1-\varepsilon} = \Psi(\alpha w^{1-\varepsilon} + \eta r^{1-\varepsilon}).$$

Let $s_i = p_i^{\varepsilon} y_i$. Then

$$s' = (P_c^{\varepsilon} C \beta + P_G^{\varepsilon} G \delta)' \Psi.$$

Without loss of generality, suppose that the wage is stuck at $w = 1$ and that the labor market is not clearing. Then, labor is determined by demand, so we have

$$l = s' \alpha.$$

Therefore,

$$\frac{\mathrm{d}l}{\mathrm{d}\delta_i} = \left(\beta \frac{\mathrm{d}P_c^{\varepsilon} C}{\mathrm{d}\delta_i} + \delta \frac{\mathrm{d}P_G^{\varepsilon} G}{\mathrm{d}\delta_i} + e_i P_G^{\varepsilon} G \right)' \Psi \alpha. \tag{A.37}$$

Note that

$$P_c^{\varepsilon} C = \frac{l + rK - \tau}{P_c^{1-\varepsilon}},$$

$$= \frac{\mathrm{d}l/\mathrm{d}\delta_i + \mathrm{d}rK/\mathrm{d}\delta_i}{P_c^{1-\varepsilon}} + (\varepsilon - 1) P_c C P_c^{\varepsilon} \left(\beta' \Psi \eta \mathrm{d}r^{1-\varepsilon}/\mathrm{d}\delta_i \right),$$

where the last line uses the fact that

$$P_c^{1-\varepsilon} = \beta \Psi (\alpha + \eta r^{1-\varepsilon}).$$

Market clearing for capital implies that

$$rK = (1 - s'\alpha) r^{1-\varepsilon}.$$

Therefore,

$$r^{1-\varepsilon} = \left(\frac{1 - s'\alpha}{K} \right)^{\frac{1-\varepsilon}{\varepsilon}}.$$

So,

$$\frac{\mathrm{d}r^{1-\varepsilon}}{\mathrm{d}\delta_i} = \left(\frac{\varepsilon - 1}{\varepsilon} \right) \left(\frac{s'\eta}{K} \right)^{\frac{1-\varepsilon}{\varepsilon} - 1} \frac{\mathrm{d}s'\alpha}{\mathrm{d}\delta_i} \frac{1}{K}.$$

174

Furthermore,

$$\frac{\mathrm{d}rK}{\mathrm{d}\delta_i} = \frac{\mathrm{d}r^{1-\varepsilon}}{\mathrm{d}\delta_i} - r^{1-\varepsilon}\frac{\mathrm{d}s'\alpha}{\mathrm{d}\delta_i}.$$

Combine these two equations with (A.37) to get

$$\left[\left(1 - \beta'\tilde{\alpha}P_c^{\varepsilon-1}(1 - r^{1-\varepsilon})\right) - \beta'\tilde{\alpha}P_c^{\varepsilon-1}\left(1 + (\varepsilon-1)P_c^2(1 - s'\alpha)\right)\frac{\varepsilon-1}{\varepsilon}\left(\frac{s'\eta}{K}\right)^{\frac{1-\varepsilon}{\varepsilon}-1}\frac{1}{K}\right]\frac{\mathrm{d}s'\alpha}{\mathrm{d}\delta_i}$$

$$= \tilde{\alpha}_i P_G^\varepsilon G + \delta'\tilde{\alpha}\frac{\mathrm{d}P_G^\varepsilon G}{\mathrm{d}\delta_i},$$

$$= \tilde{\alpha}_i P_G^\varepsilon G + \delta'\tilde{\alpha}\frac{\mathrm{d}\tau/P_G^{1-\varepsilon}}{\mathrm{d}\delta_i},$$

$$= \tilde{\alpha}_i P_G^\varepsilon G + \delta'\tilde{\alpha}\frac{P_G^\varepsilon G}{P_G^{\varepsilon-1}}\left[\tilde{\alpha}_i + \tilde{\eta}_i r^{1-\varepsilon} + \delta'\tilde{\eta}\frac{\mathrm{d}r^{1-\varepsilon}}{\mathrm{d}\delta_i}\right].$$

Rearrange this to get

$$\Theta\frac{\mathrm{d}s'\alpha}{\mathrm{d}\delta_i} = \tilde{\alpha}_i - \frac{\delta\tilde{\alpha}}{P_G^{1-\varepsilon}}\left(\tilde{\alpha}_i + \tilde{\eta}_i r^{1-\varepsilon}\right),$$

where Θ does not depend on δ_i. Recall that $\mathrm{d}s'\alpha/\mathrm{d}\delta_i = \mathrm{d}l/\mathrm{d}\delta_i$. Divide the derivative of labor with respect to δ_i by the derivative with respect to δ_j to get

$$\frac{\mathrm{d}l/\mathrm{d}\delta_i}{\mathrm{d}l/\mathrm{d}\delta_j} = \frac{\tilde{\alpha}_i - \frac{\delta\tilde{\alpha}}{P_G^{1-\varepsilon}}\left(\tilde{\alpha}_i + \tilde{\eta}_i r^{1-\varepsilon}\right)}{\tilde{\alpha}_j - \frac{\delta\tilde{\alpha}}{P_G^{1-\varepsilon}}\left(\tilde{\alpha}_j + \tilde{\eta}_j r^{1-\varepsilon}\right)},$$

$$= \frac{\tilde{\alpha}_i - \frac{\delta\tilde{\alpha}}{P_G^{1-\varepsilon}}\left(\tilde{\alpha}_i + \tilde{\eta}_i r^{1-\varepsilon}\right)}{\tilde{\alpha}_j - \frac{\delta\tilde{\alpha}}{P_G^{1-\varepsilon}}\left(\tilde{\alpha}_j + \tilde{\eta}_j r^{1-\varepsilon}\right)},$$

$$= \frac{\tilde{\alpha}_i - \frac{\delta\tilde{\alpha}/P_G^{1-\varepsilon}r^{1-\varepsilon}}{1-\delta\tilde{\alpha}/P_G^{1-\varepsilon}(1-r^{1-\varepsilon})}}{\tilde{\alpha}_j - \frac{\delta\tilde{\alpha}/P_G^{1-\varepsilon}r^{1-\varepsilon}}{1-\delta\tilde{\alpha}/P_G^{1-\varepsilon}(1-r^{1-\varepsilon})}}, \tag{A.38}$$

Note that

$$\frac{\delta\tilde{\alpha}/P_G^{1-\varepsilon}r^{1-\varepsilon}}{1-\delta\tilde{\alpha}/P_G^{1-\varepsilon}(1-r^{1-\varepsilon})} = \frac{\delta'\tilde{\alpha}r^{1-\varepsilon}}{P_G^{\varepsilon-1}-\delta'\tilde{\alpha}(1-r^{1-\varepsilon})},$$

$$= \frac{\delta'\tilde{\alpha}r^{1-\varepsilon}}{\delta'(\tilde{\alpha}+\tilde{\eta}r^{1-\varepsilon})-\delta'\tilde{\alpha}(1-r^{1-\varepsilon})},$$

$$= \frac{\delta'\tilde{\alpha}r^{1-\varepsilon}}{r^{1-\varepsilon}},$$

$$= \delta'\tilde{\alpha}.$$

Substitute this fact into expression (A.38) to get the desired result.

∎

A.3.2 Labor Share of Income with CES

The results of section 1.4 take advantage of the Cobb-Douglas form of the production functions, but they can be viewed in a more general reduced-form manner. Whatever the underlying production functions of the different industries, we can define Ω, α, η, and β to be the observed expenditure shares. Market clearing will then imply that equation (1.3) must hold. Which in turn allows us to carry out the decomposition in (1.5). Therefore, since these depend only on accounting identities, the resulting calculations still tell us how changing expenditure shares are changing labor's share of income regardless of the underlying production functions. Of course, without a structural model, it is impossible to know the causes of these changes.

On the structural front, the benchmark model can be extended to allow for non-unit symmetric elasticity of substitution. Let the composite consumption of household be given by

$$C = \left(\sum_k \beta_k^{\frac{1}{\varepsilon_h}} c_k^{\frac{\varepsilon_h-1}{\varepsilon_h}}\right)^{\frac{\varepsilon_h}{\varepsilon_h-1}}.$$

Note that as $\varepsilon_h \to 1$, we recover the utility function of the benchmark model. Let the

production function of the representative firm in industry i be given by

$$y_i \leq \left(\alpha_i^{\frac{1}{\varepsilon}} l_i^{\frac{\varepsilon-1}{\varepsilon}} + \eta_i^{\frac{1}{\varepsilon}} k^{\frac{\varepsilon-1}{\varepsilon}} + \sum_i^N \omega_{ij}^{\frac{1}{\varepsilon}} x_{ij}^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

Once again, note that letting $\varepsilon \to 1$ recovers the production functions of the benchmark model.

Now, labor's share of private GDP can be written as

$$\frac{wl}{GDP} = \left(\beta \circ P^{\varepsilon - \varepsilon_h} \right)' \tilde{\alpha} w^{1-\varepsilon} P_c^{\varepsilon_h - 1}, \tag{A.39}$$

where P is a column vector of the price of each good, w is the nominal wage, and P_c is the price level of aggregate consumption. The network-adjusted labor intensities are still defined as $\tilde{\alpha} = (I - \hat{\Omega})^{-1} \alpha$, but now these numbers pertain to the share parameters rather than the observed expenditure shares. Equation (A.39) allows us to carry out a decomposition similar to (1.5), although now we need both price and quantity data to identify the relevant parameters. A further restriction of $\varepsilon_h = \varepsilon$ gives us the even simpler expression

$$\frac{wl}{GDP} = \beta' \tilde{\alpha} \left(\frac{w}{P_c} \right)^{1-\varepsilon}.$$

Note that these equations will hold as long as consumption and production have the CES form, and do not depend on assumptions about other aspects of the model like labor supply, intertemporal decision-making, and capital accumulation.

A.4   Appendix IV: Profits

In this section, I show that firm profits, rather than inelastically supplied capital, can play the role of a non-labor sink and break the irrelevance result in section 1.2. A tractable way of showing this is to use Dixit-Stiglitz monopolistic competition.

Let the representative household maximize

$$\max_{c_{it}, l_t, B_t} \quad \sum_{t=0}^{\infty} \rho^t \left(\log(C_t) - \frac{l_t^\theta}{\theta} \right),$$

where

$$C_t = \prod_{i=1}^{N} c_{it}^{\frac{\beta_i}{\sum_j \beta_j}},$$

subject to budget constraint

$$\sum (1 + \tau_{it}) p_{it} c_{it} + q_t B_t = w_t l_t + B_{t-1} + \Pi_t - \tau_t,$$

where p_{it} is the price of good i in period t, B_t is a nominal bond (in zero net supply), Π_t is firm profits, τ_t is lump sum taxes in period t. Note that we no longer have capital income.

The government runs balanced budgets every period

$$\sum p_i g_i = \tau + \sum_i \tau_{it} p_{it} c_{it}, \tag{A.40}$$

and the fraction of government expenditures on industry i is

$$\frac{\delta_i}{\sum_j \delta_j} = \frac{p_i g_i}{\sum_i p_i g_i}.$$

Without loss of generality, suppose that there is a unit mass of firms in each industry. Assume that these firms are monopolistically competitive so that they make positive profits in equilibrium, and the elasticity of substitution across firms producing different varieties in industry i is given by $\varepsilon_i > 1$. The representative firm in industry i maximizes profits

$$\max_{y_i, l_i, x_{ij}} \quad p_i y_i - \sum_j p_j x_{ij} - w l_i$$

subject to the production function

$$y_i = (l_i)^{\alpha_i} \prod x_{ij}^{(1-\alpha_i)\omega_{ij}},$$

where ω_{ij} is the intensity with which the representative firm in industry i uses inputs from industry j. Assume that $\sum_j \omega_{ij} = 1$ for all i to maintain constant returns to scale. In equilibrium, firm i sets its price equal to

$$p_i = \frac{\varepsilon_i}{\varepsilon_i - 1} \lambda_i,$$

where λ_i is its marginal cost. Let μ_i denote the reciprocal of the markup of industry i.

Cost minimization by the firm implies that

$$wl_i = \alpha_i \lambda_i y_i = \alpha_i \mu_i p_i y_i,$$

and

$$p_j x_{ij} = \omega_{ij} \lambda_i y_i = \omega_{ij} \mu_i p_i y_i.$$

Denote the sales of industry i by s_i. Then market clearing for industry i's goods implies

$$s' = H + G + s'\mu \text{diag}(1 - \alpha)\Omega = \beta'(H + G) + s'\mu\hat{\Omega},$$

where μ is a diagonal matrix whose ith diagonal element is μ_i. This implies that

$$s' = (H + G)'(I - \mu\hat{\Omega})^{-1}\alpha.$$

Market clearing for labor implies that

$$wl = s'\alpha = (H + G)'(I - \mu\hat{\Omega})^{-1}\alpha.$$

The network-adjusted labor intensities are now given by

$$\tilde{\alpha} = (I - \mu\hat{\Omega})^{-1}\alpha.$$

Labor supply is the same as before

$$l = \left(\frac{wl}{PC}\right)^{\frac{1}{\theta}}.$$

Combining these two equations we get that equilibrium employment must equal

$$l = \left((H + G)'\tilde{\alpha}P\right)^{\frac{1}{\theta}}.$$

It is easy to verify that this model behaves in the same way as the benchmark model in section 1.3.

179

A.5 Appendix V: Sticky Prices

In this section, I sketch how the basic intuition of the case with sticky wages can be extended to sticky prices, as long as labor (and not capital) is the factor that falls during the recession. Let household utility be given by

$$\sum_t \rho^t \left[(1 - \lambda) \log(c_t) - \frac{l_t^\theta}{\theta} + \lambda \log(G_t) \right], \quad \lambda \in (0,1)$$

where

$$c_t = \sum_k (c_{t,k})^{\beta_k},$$

is private consumption, and

$$G_t = \prod_i g_{it}^{\phi_i}.$$

is government consumption services. We maintain the assumption that $\sum_k \beta_k = 1$. Household's budget is

$$\sum_k p_{t,k} c_{t,k} + B_t = (w_t l_t + r_t K_t) + (1 + i_{t-1}) B_{t-1} + \Pi_t - \tau_t,$$

where $p_{t,k}$ is the price of good k in time t. Nominal government bonds are B_t and Π_t is firm profits in period t. The nominal net interest rate on debt is i_t. The household receives labor income $w_t l_t$ and capital income $r_t K_t$. Households are endowed with an exogenous amount of capital, and both the wage and the rental rate of capital are flexible. Finally, savers face lump sum taxes τ_t.

Suppose that each industry consists of a fraction ξ of firms who set their prices ever period and $1 - \xi$ whose prices are pre-determined. The production function of firms in industry i are given by

$$y_{it} = (l_{it})^{\alpha_i} k_{it}^{\eta_i} \prod x_{ijt}^{(1-\alpha_i-\eta_i)\omega_{ij}}.$$

I assume that demand for goods from industry i are a CES bundle of goods from the firms with pre-determined prices and firms with flexible prices.

A.5.1 Discount Factor Shock

Suppose that there is an unexpected discount factor shock so that for the next period, $\rho^* > 1$. I analyze the government's fiscal policy without commitment when interest rates are at the zero lower bound.

Proposition A.5.1. *The relative employment multiplier of government spending satisfies*

$$\frac{dl/d\delta_i}{dl/d\delta_i} = \frac{\tilde{\alpha}_i - \delta'\tilde{\alpha}}{\tilde{\alpha}_j - \delta'\tilde{\alpha}}. \tag{A.41}$$

Proof. As before, the Euler equation pins down current household expenditures to be

$$P_t C_t = \frac{\overline{PC}}{\rho^*}.$$

Labor supply then implies that

$$l_t^{\theta-1} = \frac{w_t \rho^*}{\overline{PC}}.$$

Furthermore, combining labor supply and labor demand implies that the equilibrium wage is given by

$$w_t = \left(\left(\beta' \frac{\overline{PC}}{\rho^*} + \delta'\tau \right) \tilde{\alpha} \right)^{\frac{\theta-1}{\theta}} (\overline{PC})^{\theta}.$$

Define government expenditure shares on industry i to be

$$\frac{\delta_i}{\sum_k \delta_k},$$

where we assume that $\sum_k \delta_k = 1$. Now Implicitly differentiate equilibrium employment with respect to δ_i evaluated at $\sum_k \delta_k = 1$ to get

$$(\theta-1)\frac{\overline{PC}}{\rho^*} l^{\theta-2} \frac{dl}{d\delta_i} = \frac{\theta-1}{\theta} (w_t l_t)^{-\frac{1}{\theta}} (\overline{PC})^{\theta} \tau(\tilde{\alpha}_i - \delta'\tilde{\alpha}).$$

Divide this expression through by its counterpart for δ_j to get the desired result. ∎

Proposition A.5.1 shows that the government multipliers behave much in the same way with sticky prices and elastic labor supply as with sticky wages and inelastic labor supply. Welfare analysis in this context is considerably less clean however, since unlike the case

with downward sticky wages, in this case, the efficient level of employment depends on the marginal utility of consumption, so there is no "employment-targeting".

Proposition A.5.2. *The optimal share of expenditures by the government in industry i relative to industry j satisfies*

$$\frac{\delta_i}{\delta_j} = \frac{\phi_i}{\phi_j} \frac{const - \tilde{\alpha}_j}{const - \tilde{\alpha}_i}, \tag{A.42}$$

where const > 1 . So the government tilts spending according to network-adjusted labor intensities.

Proof. As before, the Euler equation pins down current household expenditures to be

$$P_t C_t = \frac{\overline{PC}}{\rho^*}.$$

Therefore current consumption is

$$C_t = \frac{\overline{PC}}{\rho^*} \frac{1}{P_t}.$$

Using the above expression for household consumption, the government's optimization problem in period t can be written as

$$\max \quad \frac{1}{\varepsilon - 1} \sum_i (\beta_i + 1) \log(\xi p_{it}^* + (1 - \xi) p_{it}) + \lambda \sum_i \phi \log(\delta_i) + \lambda \log(\tau_t) - \frac{l_t^\theta}{\theta},$$

subject to

$$p_{it}^* = c_i w_t^{\tilde{\alpha}_i} r_t^{\tilde{\eta}_i}, \quad \text{for each } i$$

$$l_t^{\theta - 1} = \rho^* \frac{w_t}{\overline{PC}},$$

$$w_t = \left(\left(\beta' \frac{\overline{PC}}{\rho^*} + \delta' \tau \right) \tilde{\alpha} \right)^{\frac{\theta - 1}{\theta}} (\overline{PC})^\theta,$$

$$r_t = \beta' \tilde{\eta} + \delta' \tilde{\eta} \frac{\tau}{\bar{k}},$$

$$1 = \sum_i \delta_i.$$

The first order conditions for this can be written as

$$\frac{\delta_j}{\delta_i} = \frac{\phi_j}{\phi_i} \frac{const - \tilde{\alpha}_i}{const - \tilde{\alpha}_j},$$

where

$$\text{const} = -\frac{\mu_4 \tau_t/\bar{k} + \mu_5}{\mu_3 \frac{\theta-1}{\theta}(wl)^{-1/\theta} - \mu_4 \tau_t/\bar{k}},$$

where μ_k is the lagrange multiplier corresponding to the kth constraint. The term $\mu_4 \tau_t/\bar{k}$ captures the scarcity of capital, while the term $\mu_3 \frac{\theta-1}{\theta}(wl)^{-1/\theta}$ captures the scarcity of labor.

∎

A.6 Appendix VI: Tables and Graphs

Table A.1: *Cumulative changes to each component of labor share from 1996-2009. Rows represent the different countries in the sample. Columns 1-4 show the fraction of change in the labor share for low-skilled labor attributable to each component. The fifth column gives the change in the labor share for low-skilled labor. The columns 6-9 show the fraction of change in the labor share for medium-skilled labor attributable to each component. The tenth column gives the change in the labor share for medium-skilled labor. Columns 11-14 show the fraction of change in the labor share for high-skilled labor attributable to each component. The final column gives the change in the labor share for high-skilled labor.*

	cons	supply	labor	lowskill	total change lowskill	cons	supply	labor	medskill	total change medskill	cons	supply	labor	highskill	total change highskill
AUS	0.142	0.128	-0.014	0.745	-0.055	4.232	1.019	0.524	-4.774	-0.002	0.074	0.046	-0.043	0.924	0.036
AUT	0.100	-0.038	0.207	0.731	-0.038	0.235	-0.226	0.734	0.257	-0.050	-0.056	0.248	-0.234	1.042	0.039
BEL	0.040	-0.024	0.107	0.877	-0.078	0.109	0.072	-0.296	1.115	0.033	0.208	0.106	-0.036	0.721	0.043
BGR	0.339	0.145	0.054	0.462	-0.069	0.201	0.383	-0.210	0.625	0.020	0.213	0.294	-0.124	0.618	0.031
BRA	0.030	-0.039	0.032	0.977	-0.052	-0.079	0.056	0.105	0.918	0.019	-0.069	-0.044	-0.033	1.146	0.029
CAN	0.013	-0.022	0.091	0.918	-0.011	-0.126	-0.647	0.786	0.987	-0.028	0.068	0.186	-0.192	0.938	0.040
CZE	0.423	0.566	-2.379	2.391	-0.003	-0.405	-0.200	1.949	-0.343	0.047	-0.058	0.107	0.542	0.409	0.056
DEU	0.116	0.132	0.243	0.509	-0.016	0.138	0.102	0.319	0.440	-0.066	-0.111	0.098	-0.225	1.239	0.030
DNK	-0.425	0.098	0.366	0.960	0.013	0.248	-0.065	-0.150	0.967	-0.055	0.105	0.184	0.023	0.688	0.058
ESP	0.025	-0.012	0.172	0.815	-0.113	0.097	0.037	-0.584	1.449	0.024	0.223	0.228	-0.860	1.409	0.040
EST	-0.567	-0.001	0.756	0.812	0.009	-1.033	0.013	1.481	0.540	0.017	0.453	0.564	0.770	-0.787	0.020
FIN	0.054	-0.100	0.146	0.900	-0.060	-0.914	2.916	-3.186	2.184	0.004	0.060	0.278	-0.199	0.860	0.052
FRA	0.007	0.023	0.080	0.890	-0.060	2.209	0.886	5.022	-7.118	-0.001	-0.060	0.131	0.012	0.916	0.048
GBR	0.047	0.021	-0.086	1.018	-0.069	-0.704	-0.352	1.707	0.349	0.007	0.099	0.100	0.029	0.771	0.088
GRC	0.322	0.512	-0.860	1.026	-0.059	-0.146	-0.242	0.661	0.728	0.052	0.229	-0.080	0.501	0.351	0.066
HUN	0.228	0.188	0.083	0.500	-0.036	0.326	0.272	0.144	0.258	-0.079	-0.346	-0.086	0.210	1.223	0.031
IND	0.264	0.199	0.181	0.355	-0.074	0.086	0.345	0.264	0.305	-0.026	0.400	-0.104	-0.523	1.228	0.028
IRL	0.228	0.163	-0.111	0.720	-0.079	0.446	0.372	-0.129	0.310	-0.040	0.162	-0.103	0.033	0.908	0.076
ITA	0.086	-0.013	0.070	0.857	-0.114	-0.096	0.201	-0.184	1.079	0.059	0.099	0.192	-0.093	0.802	0.042
JPN	0.122	0.033	0.062	0.784	-0.044	0.331	-0.020	0.508	0.181	-0.045	-0.117	0.173	-0.375	1.319	0.032
KOR	0.146	-0.080	0.159	0.776	-0.075	0.218	-0.127	0.566	0.343	-0.087	-0.210	0.482	-2.882	3.609	0.024
LTU	2.320	2.251	-3.948	0.376	-0.003	-5.861	-5.846	14.188	-1.480	0.004	0.138	0.003	0.726	0.134	0.054
MEX	-0.009	0.106	0.166	0.737	-0.028	0.197	-0.131	-0.641	1.575	0.023	-0.492	0.100	0.319	1.073	-0.014
NLD	0.175	-0.011	0.063	0.772	-0.040	0.047	-0.082	0.212	0.823	-0.039	0.132	0.058	-0.066	0.877	0.072
POL	0.273	0.290	0.250	0.187	-0.123	0.224	0.257	0.422	0.097	-0.330	-0.258	-0.113	-1.358	2.730	0.020
PRT	0.521	0.105	-0.600	0.974	-0.045	0.159	0.133	0.003	0.706	0.026	0.261	0.094	0.057	0.588	0.045
ROU	2.388	1.658	-3.950	0.904	-0.039	-0.096	-0.179	0.272	0.811	0.020	0.307	0.025	0.134	0.534	0.035
RUS	0.107	-0.008	0.147	0.754	-0.014	-0.099	-1.336	2.057	0.379	-0.018	0.261	0.417	-0.031	0.352	0.049
SVK	0.106	0.199	-0.074	0.769	-0.010	1.317	1.905	-2.862	0.640	-0.011	-0.052	-0.019	0.424	0.646	0.023
SVN	0.272	0.171	0.218	0.339	-0.052	0.196	0.134	0.347	0.323	-0.072	0.099	0.016	0.054	0.832	0.049
SWE	0.028	0.061	0.094	0.816	-0.050	-0.013	0.141	0.498	0.374	-0.035	-0.035	0.041	-0.175	1.168	0.046
TUR	0.939	-0.851	-0.337	1.249	-0.033	0.101	0.497	-0.174	0.577	0.026	0.467	0.244	-0.304	0.593	0.043
TWN	0.165	0.043	0.235	0.557	-0.098	0.355	-0.108	0.717	0.036	-0.028	-0.175	0.141	-0.532	1.567	0.036
USA	0.175	0.159	-0.055	0.721	-0.013	0.137	0.221	0.001	0.641	-0.055	0.156	0.009	-0.042	0.878	0.051
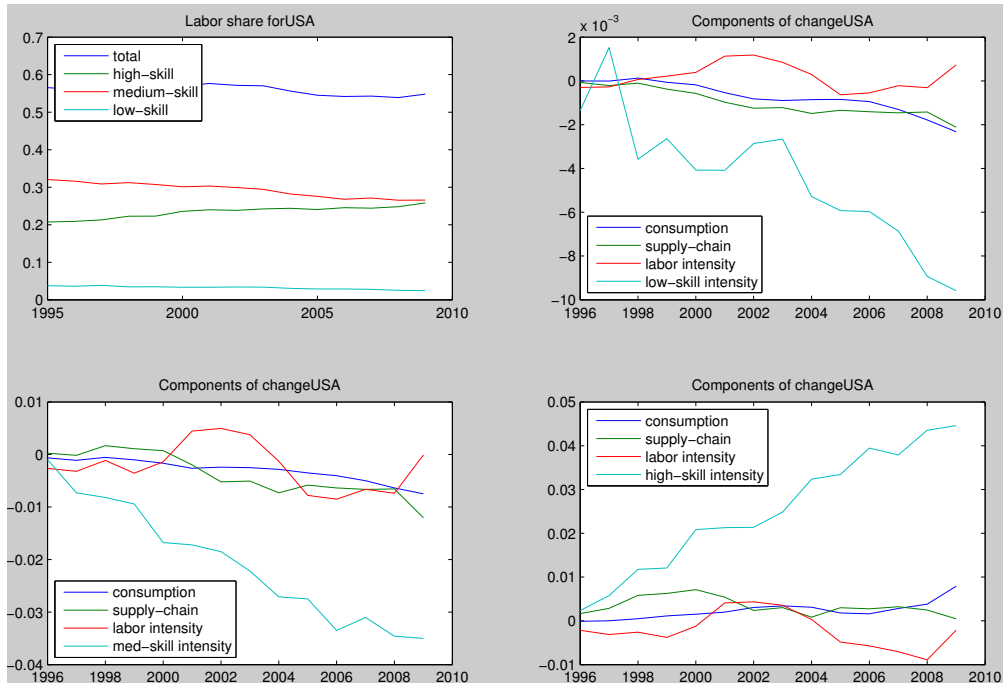
Figure A.2: *Evolution of the labor share over time for the US using the WIOD.*
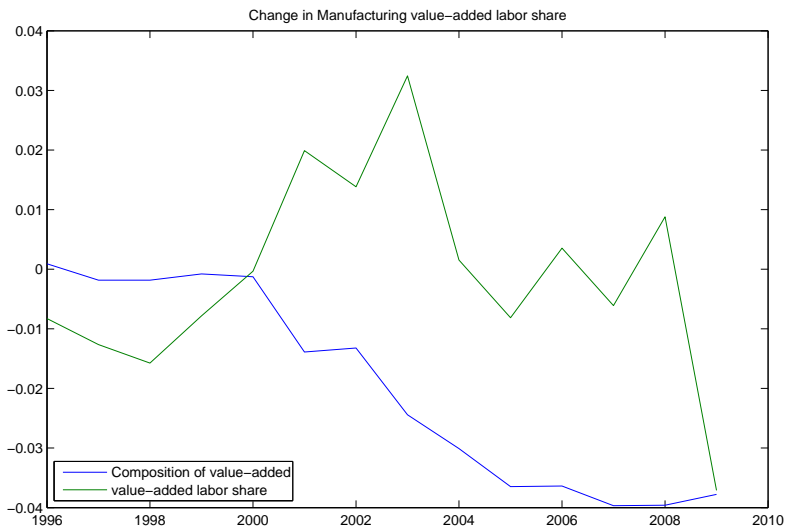
Figure A.3: *Evolution of labor use by the manufacturing industries of the US using the WIOD data from 1995-2009.*

Appendix B

Appendices to Chapter 2

B.1 Appendix I: Proofs

Lemma B.1.1. *Demand for firm i in industry k's output is*

$$y(k,i) = c(k,i) + \sum_l \int^{M_l} x(l,j,k,i)\mathrm{d}j,$$

$$= \beta_k \left(\frac{p(k,i)}{p_k}\right)^{-\varepsilon_k} \left(\frac{p_k}{p_c}\right)^{-\sigma} C + \sum_l M_l \omega_{lk} \left(\frac{p(k,i)}{p_k}\right)^{-\varepsilon_k} \left(\frac{p_k}{\lambda_l}\right)^{-\sigma} y(l,j).$$

Proof. Cost minimization by each firm implies firm j in industry l's demand for inputs from firm i in industry k is given by

$$x(l,j,k,i) = \omega_{lk} \left(\frac{p(k,i)}{p_k}\right)^{-\varepsilon_k} \left(\frac{p_k}{\lambda_l}\right)^{-\sigma} y(l,j),$$

where λ_l is the marginal cost of firms in industry l,

$$\lambda_k = \left(\alpha_k z_k^{\sigma-1} w^{1-\sigma} + \sum_l \omega_{kl} p_l^{1-\sigma}\right)^{\frac{1}{1-\sigma}},$$

and p_k is the price index for industry k

$$p_k = \left(\int^{M_k} p(k,i)^{1-\varepsilon_k}\mathrm{d}i\right)^{\frac{1}{1-\varepsilon_k}}.$$

Household demand for goods from firm i in industry k are

$$c(k,i) = \beta_k \left(\frac{p(k,i)}{p_k} \right)^{-\varepsilon_k} \left(\frac{p_k}{P_c} \right)^{-\sigma} C,$$

where P_c is the consumer price index

$$P_c = \left(\sum_k \beta_k p_k^{1-\sigma} \right)^{\frac{1}{1-\sigma}}.$$

Adding the household and firm's demands together gives the result. $\blacksquare$

Proof of lemma 2.3.1. By lemma B.1.1

$$y(k,i) = \beta_k p(k,i)^{\varepsilon_k} p_k^{\varepsilon_k - \sigma} p_c^\sigma C + \sum_l M_l p(k,i)^{\varepsilon_k} p_k^{\varepsilon_k - \sigma} \lambda_l^\sigma \omega_{lk} y(l,j).$$

Substitute $p_k = M_k^{\frac{1}{1-\varepsilon_k}} p(k,i)$ to get

$$M^{\frac{\varepsilon_k}{\varepsilon_k - 1}} p_k^\sigma y(k,i) = \beta_k p_c^\sigma C + \sum_l M_l \lambda_l^\sigma \omega_{lk} y(l,j).$$

Observe that

$$y_k = \left(\int^{M_k} y(k,i)^{\frac{\varepsilon_k - 1}{\varepsilon_k}} di \right)^{\frac{\varepsilon_k}{\varepsilon_k - 1}} = M_k^{\frac{\varepsilon_k}{\varepsilon_k - 1}} y(k,i).$$

Substitute this into the previous equation to get

$$p_k^\sigma y_k = \beta_k P_c^\sigma C + \sum \omega_{lk} M_l^{\frac{\sigma-1}{\varepsilon_l - 1}} \left(\frac{\varepsilon_l}{\varepsilon_l - 1} \right)^{-\sigma} p_l^\sigma y_l.$$

Define $\tilde{M}$ to be the diagonal matrix whose kth diagonal element is $M_k^{\frac{1}{\varepsilon_k - 1}}$, and μ to be the diagonal matrix whose kth element is industry k's markup $\varepsilon_k/(\varepsilon_k - 1)$. Now denote $s_k = p_k^\sigma y_k$. This means that we can write

$$s' = \beta' P_c^\sigma C + s' \tilde{M}^{\sigma-1} \mu^{-\sigma} \Omega.$$

Rewrite this to get

$$s' = \beta'(I - \tilde{M}^{\sigma-1} \mu^{-\sigma} \Omega)^{-1} P_c^\sigma C$$

$$= \tilde{\beta}' P_c^\sigma C.$$

■

Proof of lemma 2.3.2. Since all firms have constant returns to scale on the margin, firm i's problem in industry k, conditional on entry, can be written as

$$\max p(k,i)y(k,i) - \lambda_k y(k,i),$$

where, by lemma B.1.1,

$$y(k,i) = \mathrm{const}_k p(k,i)^{-\varepsilon_k},$$

because the firm does not internalize its effects on the aggregate price indices. This optimization problem gives

$$p(k,i) = \frac{\varepsilon_k}{\varepsilon_k - 1}\lambda_k.$$

So, as in standard monopolistic competition models, markups are constant.

Note that

$$p_k = \left(\frac{\varepsilon_k}{\varepsilon_k - 1}\right) M_k^{\frac{1}{1-\varepsilon_k}}\lambda_k.$$

Substituting this into the definition of λ_k we get

$$\left(\frac{\varepsilon_k - 1}{\varepsilon_k}\right) M_k^{\frac{1}{\varepsilon_k - 1}} p_k = \left(\alpha_k w^{1-\sigma} + \sum_l \omega_{kl} p_l^{1-\sigma}\right)^{\frac{1}{1-\sigma}}.$$

A column vector an exponent denotes element-wise exponentiation. Then if we let P be the vector of $p_k^{1-\sigma}$ and α the vector of α_k, then this system of equations can be written as

$$\mu^{\sigma-1}\tilde{M}^{1-\sigma}P = \alpha w^{1-\sigma} + \Omega P.$$

We can rearrange this to get

$$P = \mu^{1-\sigma}\tilde{M}^{\sigma-1}\alpha w^{1-\sigma} + \mu^{1-\sigma}\tilde{M}^{\sigma-1}\Omega P.$$

Rearrange this to get

$$P = (I - \mu^{1-\sigma}\tilde{M}^{\sigma-1}\Omega)^{-1}\mu^{1-\sigma}\tilde{M}^{\sigma-1}\alpha w^{1-\sigma} = \tilde{\alpha}w^{1-\sigma}.$$

This implies that for each industry k

$$\left(\frac{p_k}{w_k}\right)^{1-\sigma} = \tilde{\alpha}_k.$$

■

Proof of theorem 2.3.3. Note that the profits of firm i in industry k are

$$\pi(k,i) = p(k,i)y(k,i) - \lambda_k y(k,i) - wf_k.$$

This is equivalent to

$$\pi(k,i) = p(k,i)y(k,i) - \frac{\varepsilon_k - 1}{\varepsilon_k}p(k,i)y(k,i) - wf_k,$$
$$= \frac{1}{\varepsilon_k}p(k,i)y(k,i) - wf_k.$$

Since all active firms in industry k are identical this is

$$\pi(k,i) = \frac{1}{\varepsilon_k M_k}p_k y_k - wf_k.$$

By lemmas 2.3.1 and 2.3.2,

$$p_k y_k = p_k^\sigma y_k p^{1-\sigma} = \tilde{\beta}_k \tilde{\alpha}_k P_c^\sigma C w^{1-\sigma},$$

and so

$$\pi(k,i) = \frac{1}{\varepsilon_k M_k}\tilde{\beta}_k \tilde{\alpha}_k P_c^\sigma C w^{1-\sigma} - wf_k.$$

■

Proof of proposition 2.4.1. Note that real GDP can be written as

$$C = \frac{P_c C}{P_c} = \frac{wl + \pi}{P_c},$$

where π is total profits. By free entry, profits are zero in equilibrium. Normalize $w = 1$. Then

$$\log(C) = -\log(P_c).$$

188

The marginal costs of firms in industry k are given by

$$\lambda_k = \left(\frac{z_k}{w}\right)^{-\alpha_k} \prod_l p_l^{\omega_{kl}}.$$

Since industries are perfectly competitive, firms set prices equal to their marginal costs. Let P denote the vector of industry prices. Then, in equilibrium,

$$\log(P) = (I - \Omega)^{-1} \left(\alpha \circ (\log(w) - \log(z))\right).$$

Therefore,

$$
\begin{aligned}
\log(C) &= -\log(P_c),\\
&= -\beta' \log(P),\\
&= -\beta'(I - \Omega)^{-1} \left(-\alpha \circ \log(z)\right),\\
&= \tilde{\beta}'\alpha \circ \log(z).
\end{aligned}
$$

Note that $\tilde{\beta}$ is sales as a share of GDP, and $\tilde{\beta}_k \alpha_k$ is therefore industry k's wage bill as a share of GDP. ∎

Proof of proposition 2.4.3. Note that industry k's sales are given by

$$p_k y_k = p_k^\sigma y_k p^{1-\sigma} = \tilde{\beta}_k \tilde{\alpha}_k P_c^\sigma C w^{1-\sigma}.$$

Observe that in equilibrium with no shocks,

$$\tilde{\alpha} = (I - \Omega)^{-1}\alpha = 1.$$

Therefore,

$$p_k y_k = \tilde{\beta}_k P_c^\sigma C w^{1-\sigma}.$$

So an industry's shares of sales relative to other industries is determined solely by $\tilde{\beta}_k$. ∎

Proof of theorem 2.4.5. Real consumption is given by

$$\frac{wl + \pi}{P_c} = \frac{1}{P_c}.$$

We have shown that

$$P_c = \left(\beta'(I - \Omega)^{-1} \alpha \right)^{\frac{1}{1-\sigma}}.$$

Let $\tilde{\beta}' = \beta'(I - \Omega)^{-1}$ to get the desired result. ∎

Proof of proposition 2.4.7. The sales of industry j are given by

$$\tilde{\beta}_j e_j' \Psi \left(\alpha \circ z^{\sigma-1} \right) P_c^\sigma C.$$

Therefore

$$\frac{\mathrm{d} \log(\text{sales}_j)}{\mathrm{d} z_i^{\sigma-1}} = \frac{1}{\tilde{\beta}_j} \frac{\mathrm{d} \tilde{\beta}_j}{\mathrm{d} z_i^{\sigma-1}} + \frac{1}{\tilde{\alpha}_j} \frac{\mathrm{d} e_j' \Psi \left(\alpha \circ z^{\sigma-1} \right)}{\mathrm{d} z_i^{\sigma-1}} + \frac{1}{P_c^\sigma C} \frac{\mathrm{d} P_c^\sigma C}{\mathrm{d} z_i^{\sigma-1}},$$

$$= 0 + \frac{1}{\tilde{\alpha}_j} \frac{\mathrm{d} e_j' \Psi \left(\alpha \circ z^{\sigma-1} \right)}{\mathrm{d} z_i^{\sigma-1}} + \frac{1}{P_c^\sigma C} \frac{\mathrm{d} P_c^\sigma C}{\mathrm{d} z_i^{\sigma-1}}.$$

Note that

$$\left. \frac{\mathrm{d} \Psi \left(\alpha \circ z^{\sigma-1} \right)}{\mathrm{d} z_i^{\sigma-1}} \right|_{z=1} = \Psi \left. \frac{\mathrm{d} \left(\alpha \circ z^{\sigma-1} \right)}{\mathrm{d} z_i^{\sigma-1}} \right|_{z=1} = \Psi \alpha_i e_i.$$

Substitute this in to get

$$\frac{\mathrm{d} \log(\text{sales}_j)}{\mathrm{d} z_i^{\sigma-1}} = e_j' \Psi \alpha_i e_i + \frac{1}{P_c^\sigma C} \frac{\mathrm{d} P_c^\sigma C}{\mathrm{d} z_i^{\sigma-1}},$$

where we use the fact that $\tilde{\alpha} = 1$. This means that

$$\frac{\mathrm{d} \log(\text{sales}_i)}{\mathrm{d} z_i^{\sigma-1}} - \frac{\mathrm{d} \log(\text{sales}_j)}{\mathrm{d} z_i^{\sigma-1}} = \alpha_i \left(e_i' \Psi e_i - e_j' \Psi e_i \right),$$

as required. Lemma 2.4.6 then implies that this is always greater than zero. ∎

Proof of proposition 2.4.8. By theorem 2.3.3,

$$\sum_k \pi_k = \frac{\tilde{\beta}'}{\varepsilon} \tilde{\alpha} P_c C - 1' M f,$$

where division by ε is elementwise. Observe that

$$P_c C = (1 + \sum_k \pi_k).$$

Therefore,

$$\sum_k \pi_k = \left(\frac{\tilde{\beta}'}{\varepsilon} \tilde{\alpha} - \mathbf{1}' M f \right) \frac{1}{1 - \frac{\tilde{\beta}'}{\varepsilon} \tilde{\alpha}}.$$

Therefore, nominal GDP

$$P_c C = 1 + \left(\frac{\tilde{\beta}'}{\varepsilon} \tilde{\alpha} - \mathbf{1}' M f \right) \frac{1}{1 - \frac{\tilde{\beta}'}{\varepsilon} \tilde{\alpha}},$$

does not respond to shocks. Therefore, real GDP is given by

$$\log(C) = \text{const} - \log(P_c).$$

Since

$$\log(P_c) = -\beta'(I - \Omega)^{-1} (\alpha \circ \log(z)) + \text{const},$$

we can write

$$\log(C) = \text{const} + \beta'(I - \Omega)^{-1} (\alpha \circ \log(z)).$$

∎

Proof of proposition 2.4.9.

$$\sum_k \pi_k = \frac{\tilde{\beta}'}{\varepsilon} \tilde{\alpha} P_c^\sigma C - \mathbf{1}' M f,$$

where division by ε is elementwise. Observe that

$$P_c^\sigma C = P_c C P^{\sigma-1} = \frac{(1 + \sum_k \pi_k)}{\beta' \tilde{\alpha}}.$$

Combine these two expressions to get the desired result. ∎

Proof of proposition 2.4.10. Using the same argument as in the proof of proposition 2.4.7, we can show that Substitute this in to get

$$\frac{d \log(\text{sales}_j)}{dz_i^{\sigma-1}} = \frac{1}{\tilde{\alpha}_j} e_j' \Psi \alpha_i e_i + \frac{1}{P_c^\sigma C} \frac{dP_c^\sigma C}{dz_i^{\sigma-1}},$$

This means that

$$\frac{d \log(\text{sales}_i)}{dz_i^{\sigma-1}} - \frac{d \log(\text{sales}_j)}{dz_i^{\sigma-1}} = \alpha_i \left(\frac{e_i' \Psi e_i}{\tilde{\alpha}_j} - \frac{e_j' \Psi e_i}{\tilde{\alpha}_j} \right),$$

as required. Lemma 2.4.6 then implies that this is always greater than zero. ∎

Proof of proposition 2.5.1. Note that real GDP can be written as

$$C = \frac{P_c C}{P_c} = \frac{wl + \pi}{P_c},$$

where π is total profits. By free entry, profits are zero in equilibrium. Normalize $w = 1$.
Then

$$\log(C) = -\log(P_c).$$

The marginal costs of firms in industry k are given by

$$\lambda_k = \left(\frac{\alpha_k z_k}{w}\right)^{-\alpha_k} \prod_l \left(\frac{\omega_{kl}}{p_l}\right)^{-\omega_{kl}},$$

substitute

$$\lambda_k = M_k^{\frac{1}{\varepsilon_k - 1}} \frac{\varepsilon_k - 1}{\varepsilon_k} p_k,$$

and let P denote the vector of industry prices. Then, in equilibrium,

$$\log(P) = (I - \Omega)^{-1} \left(-\alpha \circ \log(z) + \log(\mu \mathbf{1}) - \log(\tilde{M}\mathbf{1})\right).$$

Free entry implies that

$$\tilde{M}_k = \left(\frac{\tilde{\beta}_k P_c C}{f_k \varepsilon_k}\right)^{\frac{1}{\varepsilon_k - 1}}.$$

Substituting this in to $\tilde{M}$ and combining with the fact that

$$\log(P_c) = \beta' \log(P),$$

gives

$$\log(C) = -\log(P_c) = \tilde{\beta}'\alpha \circ \log(z) - \sum_k \frac{1}{\varepsilon_k - 1} \tilde{\beta}_k \log(f_k) + \text{const},$$

where

$$\text{const} = \tilde{\beta}' \left(\log(\mu \mathbf{1}) + \frac{1}{\varepsilon - 1} \circ \log(\tilde{\beta}) - \log(\varepsilon)\right).$$

In the above expression $1/(\varepsilon - 1)$ is the vector of $1/(\varepsilon_k - 1)$. By lemma 2.3.1, $\tilde{\beta}$ is proportional
to the sales vector and nominal GDP is always equal to 1, we have the desired result. ∎

Proof of lemma 2.5.2. Recall that

$$\tilde{\beta}' = \beta'(I - \tilde{M}^{\sigma-1}\mu^{-\sigma}\Omega)^{-1}.$$

So

$$\frac{\mathrm{d}\tilde{\beta}'}{\mathrm{d}\log(M_i)} = M_i \frac{\mathrm{d}\tilde{\beta}'}{\mathrm{d}M_i},$$

$$= -M_i\beta'\Psi_s \frac{\mathrm{d}(I - \tilde{M}^{\sigma-1}\mu^{-\sigma}\Omega)}{\mathrm{d}M_i}\Psi_s,$$

$$= M_i\beta'\Psi_s \frac{\mathrm{d}\tilde{M}^{\sigma-1}}{\mathrm{d}M_i}\mu^{-\sigma}\Omega\Psi_s,$$

$$= M_i\beta'\Psi_s \frac{\mathrm{d}\tilde{M}^{\sigma-1}}{\mathrm{d}M_i}\tilde{M}^{\sigma-1}\tilde{M}^{1-\sigma}\mu^{-\sigma}\Omega\Psi_s,$$

$$= M_i\beta'\Psi_s \frac{\mathrm{d}\tilde{M}^{\sigma-1}}{\mathrm{d}M_i}\tilde{M}^{\sigma-1}(\Psi_s - I),$$

The *k*th element of this vector is

$$\frac{\mathrm{d}\tilde{\beta}_k}{\mathrm{d}\log(M_i)} = \frac{\sigma - 1}{\varepsilon_i - 1}\tilde{\beta}_i(\Psi_s - I)e_k.$$

Putting this all into a matrix gives

$$\frac{\mathrm{d}\tilde{\beta}'}{\mathrm{d}\log(M)} = \mathrm{diag}(\tilde{\beta})\mathrm{diag}\left(\frac{\sigma - 1}{\varepsilon - 1}\right)(\Psi_s - I).$$

∎

Proof of lemma 2.5.3. Recall that

$$\tilde{\alpha} = (I - \tilde{M}^{\sigma-1}\mu^{1-\sigma}\Omega)^{-1}\tilde{M}^{\sigma-1}\mu^{1-\sigma}\alpha.$$

To simplify the notation, for this proof, let

$$B = (I - \tilde{M}^{\sigma-1}\mu^{1-\sigma}\Omega)^{-1}.$$

So

$$\frac{1}{M_i}\frac{d\tilde{\alpha}}{d\log(M_i)} = \frac{d\tilde{\alpha}}{dM_i},$$

$$= B\frac{d\tilde{M}^{\sigma-1}}{dM_i}\mu^{1-\sigma}\alpha + B\frac{d(\tilde{M}^{\sigma-1})}{dM_i}\left(\tilde{M}^{\sigma-1}\right)^{-1}(\tilde{M}^{\sigma-1})\mu^{1-\sigma}\Omega B(\tilde{M}^{\sigma-1})\mu^{1-\sigma}\alpha,$$

$$= B\frac{d\tilde{M}^{\sigma-1}}{dM_i}\mu^{1-\sigma}\alpha + B\frac{d(\tilde{M}^{\sigma-1})}{dM_i}\left(\tilde{M}^{\sigma-1}\right)^{-1}(B-I)(\tilde{M}^{\sigma-1})\mu^{1-\sigma}\alpha,$$

$$= B\frac{d(\tilde{M}^{\sigma-1})}{dM_i}\left(\tilde{M}^{\sigma-1}\right)^{-1}B(\tilde{M}^{\sigma-1})\mu^{1-\sigma}\alpha,$$

$$= B\frac{d(\tilde{M}^{\sigma-1})}{dM_i}\left(\tilde{M}^{\sigma-1}\right)^{-1}\tilde{\alpha},$$

$$= \Psi_d\left(\tilde{M}\right)^{1-\sigma}\mu^{\sigma-1}\frac{d(\tilde{M}^{\sigma-1})}{dM_i}\left(\tilde{M}^{\sigma-1}\right)^{-1}\tilde{\alpha},$$

The kth element of this vector is

$$\frac{d\tilde{\alpha}_k}{d\log(M_i)} = \left(\frac{\sigma-1}{\varepsilon_i-1}\right)\left(\frac{\varepsilon_i}{\varepsilon_i-1}\right)^{\sigma-1}e_k'\Psi_d e_i\tilde{\alpha}_i\frac{1}{M_i^\sigma}.$$

Putting this all into a matrix gives

$$\frac{d\tilde{\alpha}}{dM} = \Psi_d\text{diag}(\tilde{\alpha})\mu^{\sigma-1}\text{diag}(M)^{1-\sigma}\text{diag}\left(\frac{\sigma-1}{\varepsilon-1}\right).$$

∎

Proof of proposition 2.5.4. Note that

$$M_i = \frac{\tilde{\beta}_i\tilde{\alpha}_i}{\varepsilon_i f_i}P_c^\sigma C.$$

Therefore,

$$\frac{d\log(M)}{d\alpha_i} = M\mathbf{1}\frac{dP_c^\sigma C/d\alpha_i}{P_c^\sigma C} + M\text{diag}(\tilde{\beta})^{-1}\frac{d\tilde{\beta}}{dM}\frac{dM}{d\alpha_i} + M\text{diag}(\tilde{\alpha})^{-1}\frac{d\tilde{\alpha}}{dM}\frac{dM}{d\alpha_i} + M\text{diag}(\tilde{\alpha})^{-1}\Psi_d e_i.$$

Rewrite this as

$$\frac{d\log(M)}{d\alpha_i} = M\mathbf{1}\frac{dP_c^\sigma C/d\alpha_i}{P_c^\sigma C} + M\left(\text{diag}(\tilde{\beta})^{-1}\Psi_1 + \text{diag}(\tilde{\alpha})^{-1}\Psi_2\right)\frac{dM}{d\alpha_i} + M\text{diag}(\tilde{\alpha})^{-1}\Psi_d e_i.$$

Rearrange this to get the desired result. ∎

Proof of proposition 2.5.6. Note that

$$C = \frac{P_c C}{P_c} = \frac{1 + \sum_k \pi_k}{P_c} = \left(\frac{1}{\beta' \Psi_d (\alpha \circ z^{\sigma-1})} \right)^{\frac{1}{1-\sigma}}.$$

Therefore,

$$\frac{\mathrm{d}\log(C)}{\mathrm{d}z_i} = \frac{1}{\sigma - 1} \frac{\mathrm{d}\left(\log \left(\beta' \Psi_d (\alpha \circ z^{\sigma-1}) \right) \right)}{\mathrm{d}z_i},$$

$$= \frac{1}{\sigma - 1} \frac{1}{\beta' \tilde{\alpha}} \beta' \frac{\partial \tilde{\alpha}}{\partial M} \frac{\partial M}{\partial z_i} + \beta' \Psi_d \frac{\mathrm{d}(\alpha \circ z^{\sigma-1})}{\mathrm{d}z_i},$$

by lemma 2.5.3 and proposition 2.5.4,

$$= \frac{1}{\sigma - 1} \frac{1}{\beta' \tilde{\alpha}} \beta' \Psi_2 M M^{-1} (I - \Lambda)^{-1} \left(1 \frac{\mathrm{d}P_c^{\sigma} C / \mathrm{d}z_i}{P_c^{\sigma} C} + \mathrm{diag}(\tilde{\alpha})^{-1} \Psi_d e_i \right) + \frac{1}{\beta' \tilde{\alpha}} \beta' \Psi_d e_i \alpha_i z_i^{\sigma-2}.$$

Note that

$$\frac{\mathrm{d}P_c^{\sigma} C / \mathrm{d}z_i}{P_c^{\sigma} C} = \frac{\mathrm{d}\log(P_c^{\sigma} C)}{\mathrm{d}z_i} = -(\sigma - 1) \frac{\mathrm{d}\log(C)}{\mathrm{d}z_i}.$$

Substituting this into the previous expression and rearranging gives

$$\left(1 + \frac{\beta' \Psi_2 (I - \Lambda)^{-1} 1}{\beta' \tilde{\alpha}} \right) \frac{\mathrm{d}\log(C)}{\mathrm{d}z_i} = \frac{1}{\beta' \tilde{\alpha}} \left(\beta' \Psi_2 (I - \Lambda)^{-1} \left(\mathrm{diag}(\tilde{\alpha})^{-1} + I \right) \Psi_d e_i \right) \alpha_i z_i^{\sigma-2}.$$

Rearranging this gives the desired result.

∎

Proof of proposition 2.5.7. Let $\bar{s}$ be the sales of the representative firm in each industry. Then

$$\frac{\mathrm{d}\log(\bar{s})}{\mathrm{d}\alpha} = \frac{\mathrm{d}}{\mathrm{d}\alpha} \left(\log(\tilde{\beta}) + \log(\tilde{\alpha}) - \log(M) \right) + \frac{\mathrm{d}}{\mathrm{d}\alpha} \log(P_c^{\sigma} C).$$

By the chain rule,

$$\frac{\mathrm{d}\log(\bar{s})}{\mathrm{d}\alpha} = \left(\mathrm{diag}(\tilde{\beta})^{-1}\frac{\mathrm{d}\tilde{\beta}}{\mathrm{d}M} + \mathrm{diag}(\tilde{\alpha})^{-1}\frac{\mathrm{d}\tilde{\alpha}}{\mathrm{d}M} - \mathrm{diag}(M)^{-1}\right)\frac{\mathrm{d}M}{\mathrm{d}\alpha} + \mathrm{diag}(\tilde{\alpha})^{-1}\Psi_d e_k + \frac{\mathrm{d}}{\mathrm{d}\alpha}\log(P_c^\sigma C),$$

$$= \left(\mathrm{diag}(\tilde{\beta})^{-1}\Psi_1'\mathrm{diag}(M)^{-1} + \mathrm{diag}(\tilde{\alpha})^{-1}\Psi_2\mathrm{diag}(M)^{-1} - \mathrm{diag}(M)^{-1}\right)\frac{\mathrm{d}M}{\mathrm{d}\alpha}$$

$$+ \mathrm{diag}(\tilde{\alpha})^{-1}\Psi_d e_k + \frac{\mathrm{d}}{\mathrm{d}\alpha}\log(P_c^\sigma C),$$

$$= \left(\mathrm{diag}(\tilde{\beta})^{-1}\Psi_1' + \mathrm{diag}(\tilde{\alpha})^{-1}\Psi_2 - I\right)\mathrm{diag}(M)^{-1}\frac{\mathrm{d}M}{\mathrm{d}\alpha} + \mathrm{diag}(\tilde{\alpha})^{-1}\Psi_d e_k + \frac{\mathrm{d}}{\mathrm{d}\alpha}\log(P_c^\sigma C),$$

$$= \left[\left(\mathrm{diag}(\tilde{\beta})^{-1}\Psi_1' + \mathrm{diag}(\tilde{\alpha})^{-1}\Psi_2 - I\right)\mathrm{diag}(M)^{-1}\mathrm{diag}(M)\left(I - \mathrm{diag}(\tilde{\beta})^{-1}\Psi_1'\right.\right.$$

$$\left.\left. - \mathrm{diag}(\tilde{\alpha})^{-1}\Psi_2\right)^{-1} + I\right]$$

$$\mathrm{diag}(\tilde{\alpha})^{-1}\Psi_d e_k + \frac{\mathrm{d}}{\mathrm{d}\alpha}\log(P_c^\sigma C),$$

$$= 0,$$

where the second to last line uses proposition 2.5.4. ∎

Proof of proposition 2.5.8. Let s be the sales of the industries. Then

$$\frac{\mathrm{d}\log(s)}{\mathrm{d}\alpha} = \frac{\mathrm{d}}{\mathrm{d}\alpha}\left(\log(\tilde{\beta}) + \log(\tilde{\alpha}) + \log(P_c^\sigma C)\right).$$

By the chain rule,

$$\frac{\mathrm{d}\log(s)}{\mathrm{d}\alpha} = \left(\mathrm{diag}(\tilde{\beta})^{-1}\frac{\mathrm{d}\tilde{\beta}}{\mathrm{d}M} + \mathrm{diag}(\tilde{\alpha})^{-1}\frac{\mathrm{d}\tilde{\alpha}}{\mathrm{d}M}\right)\frac{\mathrm{d}M}{\mathrm{d}\alpha} + \mathrm{diag}(\tilde{\alpha})^{-1}\Psi_d e_k + \frac{\mathrm{d}}{\mathrm{d}\alpha}\log(P_c^\sigma C),$$

$$= \left(\mathrm{diag}(\tilde{\beta})^{-1}\Psi_1'\mathrm{diag}(M)^{-1} + \mathrm{diag}(\tilde{\alpha})^{-1}\Psi_2\mathrm{diag}(M)^{-1}\right)\frac{\mathrm{d}M}{\mathrm{d}\alpha} + \mathrm{diag}(\tilde{\alpha})^{-1}\Psi_d e_k + \frac{\mathrm{d}}{\mathrm{d}\alpha}\log(P_c^\sigma C),$$

$$= \left[\Lambda\mathrm{diag}(M)^{-1}\mathrm{diag}(M)(I - \Lambda)^{-1} + I\right]\left[\mathrm{diag}(\tilde{\alpha})^{-1}\Psi_d e_k + \frac{\mathrm{d}}{\mathrm{d}\alpha}\log(P_c^\sigma C)\right],$$

$$= (I - \Lambda)^{-1}\left[\mathrm{diag}(\tilde{\alpha})^{-1}\Psi_d e_k + \frac{\mathrm{d}}{\mathrm{d}\alpha}\log(P_c^\sigma C)\right],$$

which gives the desired result. Note that the second to last line of the derivation uses proposition 2.5.4, and the last line uses the fact that

$$\Lambda(I - \Lambda)^{-1} + I = (I - \Lambda)^{-1}.$$

∎

Proof of lemma 2.6.1. Element (u,v) of both Ψ_d and Ψ_s are a weighted sum of the number of directed walks from u to v. If u and v are not connected, then this is always equal to zero. Similarly, element (v,u) of both Ψ_s and Ψ_d must also be equal to zero.

Element (u,k) of both Ψ_d and Ψ_s are weighted sums of the number of directed walks from u to k. If u is not connected to v, then no walk from u to k goes through v, therefore the (u,k)th element of Ψ_d and Ψ_s do not depend on v. A similar argument implies that the (k,u)th element of Ψ_s and Ψ_d also do not depend on u.

Since $\tilde{\beta}_u$ and $\tilde{\alpha}_u$ are linear combinations of the uth column and row of Ψ_1 and Ψ_2, and since the (u,v) and (v,u)th elements of Ψ_d and Ψ_s are equal to zero, and the (k,u) and (u,k)th elements of Ψ_s and Ψ_d do not depend on v, then $\tilde{\beta}_u$ and $\tilde{\alpha}_u$ also do not depend on v. ∎

Proof of theorem 2.6.2. Consider a firm $i \in \mathcal{B}$. Suppose that the failure of v results in an equilibrium where firms in the set $\mathcal{C}$ fail. If i is not connected to any firms in $\mathcal{C}$, then by lemma 2.6.1, $\tilde{\alpha}_i$ and $\tilde{\beta}_i$ are the same regardless of whether or not v is rescued. Since the profits of firm i can be expressed as

$$\pi_i = \text{const } P_c^\varepsilon C \left(\tilde{\alpha}_i \times \tilde{\beta}_i \right),$$

if firm i prefers for v to be rescued, it must be that the exit of v causes $P_c^\varepsilon C$ to fall. Note that exits can only cause P_c^ε to increase, so if $P_c^\varepsilon C$ falls, it must be that C has fallen. Since C is the utility of the household, it follows that rescuing v must be Pareto-efficient.

Now suppose that two firms i and j in $\mathcal{B}$ disagree. In particular, i does not want v rescued but j does. Then it must be the case that $P_c^\varepsilon C$ is not falling, but either of $\tilde{\alpha}_j$ or $\tilde{\beta}_j$ has fallen – this will occur if and only if some firm in $\mathcal{C}$ is connected to firm j. ∎

Proof of proposition 2.5.11. By theorem 2.3.3, the profits of industry i are

$$\pi_i = \frac{\beta_i \alpha_i}{f_i \varepsilon_i} P_c^\sigma C,$$

when $\Omega = 0$. Therefore

$$\frac{d\pi_i}{dM_j} = \frac{\beta_i \tilde{\alpha}_i}{f_i \varepsilon_i} \frac{dP_c^\sigma C}{dM_j}.$$

We can write

$$P_c^\sigma C = \frac{1 - M'f}{\beta' \tilde{\alpha} - \mathrm{diag}(\varepsilon)^{-1} \beta' \tilde{\alpha}}.$$

Hence,

$$\frac{dP_c^\sigma C}{dM_j} = \frac{-f_j}{\beta' \tilde{\alpha} - \mathrm{diag}(\varepsilon)^{-1} \beta' \tilde{\alpha}} - \frac{1 - M'f}{\beta' \tilde{\alpha} - \mathrm{diag}(\varepsilon)^{-1} \beta' \tilde{\alpha}} \beta_i (1 - 1/\varepsilon_i) \frac{\sigma - 1}{\varepsilon_i - 1} M_i^{\frac{\sigma - \varepsilon_i}{\varepsilon_i - 1}} \mu_i^{\sigma - 1} \alpha_i.$$

Since $\sigma > 1$, both of these terms are negative and industries can only be enemies. ∎

Proof of proposition 2.5.13. Let $\pi(M)$ be the vector industrial profits with mass of entrants M. Note that an equilibrium of the continuous limit corresponds to $M(0)$ such that $\pi(M(0)) = 0$. Let $M(\Delta)$ correspond to the equilibrium mass of entrants for some $\Delta \gg 0$. Then $\pi(M(\Delta)) \geq 0$. By the mean-value inequality

$$\|\pi(M(\Delta))\| = \|\pi(M(\Delta)) - \pi(M(0))\| \leq \|D\pi(M^*)\| \|M(\Delta) - M(0)\|,$$

where M^* is in the convex hull of $M(0)$ and $M(\Delta)$. Rearrange this expression to get the desired result. ∎