NORTHWESTERN UNIVERSITY

Essays in Macroeconomics, Production Networks and Monetary History

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Economics

By

Pellet, Thomas Louis

EVANSTON, ILLINOIS

June 2023

Abstract

Essays in Macroeconomics, Production Networks and Monetary History

Pellet, Thomas Louis

Chapter 1 studies the impact of the Latin Monetary Union (1865-1927) on financial flows between European Countries. Using Machine Learning methods, we find that this monetary union fostered greater financial flows between members, especially before 1885 and except for France.

Chapter 2 studies a production network model with quantity rigidities and informational frictions, where firms may be restricted in how effectively they can adjust their intermediate input quantities in response to changes in the economic environment and they need to choose their quantities under incomplete information about the realizations of shocks. The characterization results show that these two frictions lead to a reduction in aggregate output, as firms may find it optimal to rely more heavily on less volatile suppliers, even if it comes at the cost of forgoing more efficient ones.

Chapter 3 builds on chapter 2 by relaxing the rational expectation assumption. It studies an agent-based model with quantity, price rigidities and decentralized labor and goods markets with search frictions. Firms make production decisions facing unquantifiable uncertainty about demand and unable to coordinate spontaneously on a Nash equilibrium. Inductive firms optimize over a set of forecasting strategies to best predict demand. Naive firms simply extrapolate from last period. The share of inductive firms

is a key parameter that dramatically affects business cycle characteristics. A larger share decreases the persistence and volatility of GDP and unemployment, making the process that firms try to forecast memory-less and harder to predict. This leads to the paradox that firms trying to forecast demand individually end up making larger forecast errors in the aggregate. It also qualitatively affects the passthrough from output shocks to prices and wages.

Acknowledgements

This dissertation would not have been possible without the support of my committee members, namely Martin Eichenbaum, Matt Rognlie and Alireza Tahbaz-Salehi. I would also like to thank my office mates Giovanni Sciacovelli, Michael Cai, José Diego Salas, Devis Decet, Alexandra Paluszynska, Federico Puglisi, Carl Hallmann and many others at Northwestern without which the entire process would have been a lot less stimulating. I thank Alessandro Pavan, George-Marios Angeletos, Nicolas Crouzet, Joel Mokyr, Carola Frydman, Laura Murphy, Joao Guerreiro, Leif Rasmussen and Andràs Borsos for great discussions that made this dissertation better. I would also like to thank Olivier Blanchard's mentorship along the way. I would not have started my Ph.D. without his support. I thank Jeanne Sorin for helping me during the Ph.D. application process. I thank Guido Lorenzoni for the elevator discussion that changed the second chapter of this dissertation for the better. I thank Uri Wilensky at Northwestern's computer science department for introducing me to Agent-Based-Modeling and Netlogo. I would like to thank my parents, François Pellet and Martine Tambuzzo, who gave me the education and curiosity to pursue a Ph.D. in the first place. Finally, I would like to thank Claudia Cummings for giving me the motivation to finish writing my dissertation in four years.

Dedication

À mes chers parents et Claudia.

Table of Contents

List of Tables

List of Figures

CHAPTER 1

Financial Flows in the Latin Monetary Union: A Machine Learning Approach

Written with Giovanni Sciacovelli[1]

[1]Prepared for the European Association of Banking History's conference on "Monetary Unions in History". We are grateful to Hugo Banziger, Alain Naef and conference participants for detailed comments and suggestions. We also thank Joel Mokyr, Carola Frydman and Martin Eichenbaum.

1.1. Introduction

If we could go back in time, we could generate all the data we need to answer the questions that haunt us today. But data collection cannot happen retrospectively. Economic historians are thus dependent on their predecessor's goodwill. How to access historical records of national accounts at times when the notion of national accounts did not exist? How to access records of bilateral financial flows across nations when nation states were still in their infancy? Historical records might not exist because their underlying economic concepts were yet to be discovered.

Accepting these intrinsic data limitations would greatly reduce the range of questions an economic historian can answer. The main danger is to fall for the "drunk and the lamp-post" fallacy, asking the questions one can answer instead of the questions one ought to ask.

One way forward is of course to keep searching for more data sources, discover new historical records. And there are still treasures in archives around the world to discover. It remains that this strategy is constrained by what contemporaries decided to record at the time they lived. Some variables of interest have simply never been recorded so that the precise information is lost forever. It is not possible to run a randomized control trial in the past tense, or introduce the concept of national accounting in antic Rome. And yet it might still be the information we need to answer important research questions.

A solution is to find clever ways of reinterpreting existing data in a new lights, to help us measure today what they missed then. The risk is that these natural proxies capture something else entirely. And it is not always possible to find natural proxies for the question one wants to answer.

Another solution, and the main focus of this paper, is to extract more information from the data we already have to generate synthetic proxies. In many historical applications, despite a missing variable of interest, many others variable are available. Building proxies given a set of observables is fundamentally a conditional prediction exercise. And this is exactly the type of settings where machine learning models perform well. The generalization of these methods in economic history could therefore relax the data availability constraint the same way that it did in other fields like finance (Jasova et al., 2021).

To illustrate the point, this paper considers the literature on the Latin Monetary Union (LMU), a currency union created in 1865 by France, Italy, Belgium and Switzerland to unify their monetary systems under a common bimetallic standard. Long forgotten with the global take-over of the gold standard at the end of the XIXth century, the literature on the LMU revived after the creation of the Euro area, its indirect descendant.

The LMU literature focused on establishing an extensive historical account of the events that led to its creation and later collapse (Einaudi, 2000; Willis, 1901; Einaudi, 2001) and few papers try to identify causal effects of the LMU (Flandreau, 2000; Timini, 2018). Despite being monetary and financial in nature, the literature has focused exclusively on trade in goods. The most likely explanation for this state of affairs is data availability: bilateral trade indicators are readily available, while dis-aggregated financial indicators are not.

This paper takes a different route. The LMU was effectively a common currency regime with fixed exchange rates, reducing foreign exchange risks and possibly enhancing financial market integration among its members. International financial flows rather than trade flows are for these reasons a more pertinent variable of interest. The problem is that the data does not exist at the bilateral level and only recently researchers have released measures of aggregated capital accounts for the period (Reinhart et al., 2016). Can we find a way to create a synthetic proxy for bilateral financial flows that would be good enough for causal inference applications?

This is where machine learning models can come to the rescue. By estimating the relationship between a large set of observables and our variable of interest in modern times, we are able to generate a proxy for our variable of interests in historical times, which can then be used for standard causal inference exercises.

To validate the methodology presented in this paper, we first estimate in post-WW2 data a model of trade flows for which we have 19^{th} century data. This exercise confirms that some machine learning models perform well out of sample, even decades before the estimation period. The best synthetic proxy has an out of sample R^2 of 0.53 in the 1861-1913 period and errors remain relatively homogeneous around 10-15% of the average true value in each given year.

With this new dataset, we are able to estimate the impact of the LMU on bilateral financial flows in a panel setting with country-year and country-pair fixed effects. This paper finds that the LMU had a significant impact on bilateral financial flows for its members, increasing them by 5% during the entire 1865-1913 period and by above 15% in the 1865-1885 period, when it was most active.

The structure of the paper is as follows. Section 1.2 presents the historical context. Section 1.3 present the available data. Section 1.4 describes the algorithm used to estimate the machine learning models. Section 1.5 discusses how we select the best performing model. Section 1.6 presents the main results of the paper. Section 1.7 concludes.

1.2. Historical Context

The Latin Monetary Union (LMU) was established in 1865 by France, Belgium, Switzerland and Italy[2]. The agreement revolved around the standardization of gold and silver

[2]Over time, additional countries joined the Union. Appendix A.1 provides additional details on the LMU chronology.

coinage among member countries, with the goal of reducing exchange rate uncertainties and strengthening the commercial and political relations of neighbouring nations. Both economic and political reasons led to the establishment of the Union. In the following sections, we will review both of these reasons and provide a historical recollection of the main events that characterized the life of the LMU.

Economic Reasons. From an economic point of view, Willis (1901)[3] emphasises the importance of French monetary history in the 19^{th} century to understand the reasons leading to the institution of the LMU. In 1803, France established a new law setting the ratio of exchange between gold and silver to 1:15.5. The rationale behind choosing this ratio was that, at the time, it was broadly consistent with the market value of the two metals. The consequence of setting such a fixed internal rate of exchange was that, in the years following the introduction of the law, changes in the relative market value of gold and silver led to rapid outflows of the undervalued metal. In particular, the adoption of the gold standard by England in 1816, together with the establishment of ratios equal to 1:15.873 and 1:16 in Holland and the United States, respectively, led to and increase in the world market value of gold short after the introduction of the French 1803 law. As a consequence, gold was massively exported out of France in the first half of the 19^{th} century, and the country's internal medium of exchange consisted predominantly of silver coins up until 1848. From this year thereafter there was a flow reversal, since the market value of gold relative to silver dropped below the 1:15.5 ratio: silver begun to outflow France, while gold started to be the most widely used medium of exchange within the country.

As a consequence of this rapid change in the nature of the prevailing stock of coin, the French public debate in the late 1850s was characterized by a growing interest in

[3]This work represents one of the most comprehensive reconstructions of the history of the Latin Monetary Union together with Einaudi (2001). These volumes are the main sources of the historical summary we provide in this section.

assuring a more convenient and stable medium of exchange. This interest culminated in the appointment, in 1858, of a commission[4] whose goal was to study how to solve the *monetary issue*. The commission highlighted the negative consequences that the current system had on commerce, and proposed policies aimed at stabilizing the internal medium of exchange by attacking money speculators. Despite the work of the commission, the recommended policies were not implemented by the French government.

In 1850, France, Belgium, Switzerland and Piedmont[5] unofficially agreed to have coins with the same nominal value. However, as the market values of gold and silver fluctuated, creating problems similar to the ones experienced by France, Switzerland (in 1860) and Italy (in 1862) decided to unilaterally reduce the fineness of their coins. Such unilateral practices led to a diverging currency fineness among neighbouring countries, so that arbitrage opportunity arose and the instability of the domestically used mediums of exchange was reinforced. The situation called for a collective response, which was invoked by Belgium in 1864 and that eventually took place with the monetary convention of 1865 involving France, Belgium, Switzerland and Italy, leading to the creation of the LMU.

Willis (1901) highlights that, unfortunately, the Union had the consequence of extending the *status quo* in France (conversion rate of 1:15.5 established by the 1803 law) to other smaller European countries. Importantly, while the LMU solved exchange rate problems among participating countries, it did not address the underlying issues of the French system. Although the Union was formally dissolved in 1927, Willis (1901) argues that, as a consequence of the structural instability of the French system, which was passed to the Union, it *de facto* ceased to exist already in 1885, when additional changes in the market prices of gold and silver[6] led member countries to substantially

[4] *Commission Chargeé d'Étudier la Situation monétaire*.

[5] Italy was unified in 1861.

[6] Mostly linked to the emerge of the gold standard as international monetary system (Meissner, 2015; Timini, 2018; Flandreau and Oosterlinck, 2012).

revise the original LMU agreement. In particular, in the years before 1885 there had been a reduction in the market value of silver and, similarly to the pre-LMU French experience, this had led to massive outflows of gold from LMU countries (especially France and Belgium) due to the official overvaluation of the metal imposed by the rules of the Union. As a consequence, countries reacted by reducing the possibility of silver conversion, undermining the LMU architecture.

Political Reasons

While the above reconstruction of the LMU history highlights the economic reasons that led to its creation, other authors have emphasised that political considerations also played an important role. Flandreau (2000), relying on notes by French senior officials from the Quai d'Orsay's archives, maintains that the Union represented "the starting point for an active French diplomatic campaign that aimed to introduce a franc-based international currency". According to his reconstruction, during the first half of the nineteenth century, French officials were concerned with the much greater prosperity of England relative to France, and tended to associate it with England's financial advancement and primary role as capital exporter. In particular, the rationale behind this belief was the idea that "investing abroad was spending at home" (Flandreau, 2000, p.34): by investing abroad, the investing country would stimulate a demand increase from the borrowing country, which would then buy goods from the lending nation. According to this view, then, the LMU, by imposing the French monetary system to its neighbouring countries and, therefore, easing financial exchanges, helped France in its goal of serving a more important role as lending nation in international markets. At the same time, as French capital exports to LMU members grew, borrowing countries had an incentive to denominate their liabilities in francs to reduce possible exchange rate risk premia, reinforcing the role of the French currency in capital markets.

From a political perspective, however, it is important to note that not only France, but also the other adhering countries had an incentive to join. According to Einaudi (2000), "By attempting to join the union, states with poor public finances wanted to facilitate their international trade, improve the standard of their internal currency, acquire monetary credibility, and gain access to international financial markets". Hence, Einaudi (2000) emphasises several benefits that smaller European states aimed at reaching by adhering to the Union: not only participation by these countries was seen as a way to solve monetary issues, but it was also a way to enhance participation in international trade and finance. In particular, many of these countries, such as Italy, wanted to acquire credibility as borrowers, and being part of the LMU was believed to be helpful in that regard.

The fact that adhering to the Union was also perceived as a way to access international financial markets helps explain why other countries decided not to join the Union. As a matter of fact, soon after the establishment of the LMU in 1865, the French government invited other countries, such as the United Kingdom and the German states, to join the Union. Einaudi (2000), using sources from diplomatic and banking archives, argues that, despite both Britain and Germany considered to join the Union, they lacked the incentives of Southern European countries of importing credibility or of entering international capital markets. Moreover, additional political considerations such as a potential subordinate position in the system to France, eventually led these countries to abandon the idea of adhering to the Union.

Connection to Empirical Analysis

Overall, the historical recollection of the LMU that we have provided highlights that countries that joined the Union expected to benefit from higher access to credit and international markets. Previous empirical work on the LMU has focused on identifying the effects that it had on trade flows across member countries (Flandreau, 2000; Timini, 2018) concluding that it had a very limited impact. But we believe there may be other important dimensions through which the Union may have played a role. In particular, the context surrounding the birth of the Union suggests that access to international financial markets was a critical goal. This observation provides the ground for our empirical analysis, to which we turn in the next sections.

1.3. Data

In order to implement our empirical exercise we aim to gather as much information as possible to accurately reconstruct a proxy for bilateral financial flows during the 19^{th} century. To achieve this goal, we rely on several data sources, which we describe in the next section. Afterwards, we describe how we merged these sources into the final dataset used for our exercise.

1.3.1. Data Sources

The first data source is Tradehist (Fouquin and Hugot, 2016), a dataset that has been recently developed for the empirical investigation of bilateral trade flows during the period 1827-2014. Five types of variables are included in the dataset: i) bilateral trade flows, ii) country-level aggregate exports and imports, iii) GDPs, iv) exchange rates, and v) additional bilateral factors that can favor or hamper trade[7]. Given the fact that Timini (2018), which represents the most up-to-date analysis of the impact of the LMU on trade flows, used a different dataset, it is worth emphasising why we believe Tradehist to be the appropriate data source for our analysis. Timini (2018)'s analysis relies on RICardo (Dedinger and Girard, 2017), a dataset containing bilateral trade flows during the 19^{th} century. Relative to this dataset, Tradehist has two major advantages. First, its

[7] Appendix A.2 provides a list of all variables included in this dataset that are used in our exercise.

coverage is larger than that of RICardo: combining primary sources with data with pre-existing datasets (inluding RICardo itself), Tradehist reports many more observations than those of RICardo. Second, Tradehist combines trade data with additional variables that are important to explain the observed trade flows. This is not the case for RICardo, whose focus is on providing only trade and exchange rate data. Because our forecasting exercise requires as much information as possible, having both more data points and variables represent makes Tradehist more advantageous.

The second dataset we use is the IMF's Coordinated Portoflio Investment Survey (CPIS) that measures bilateral financial asset positions and financial flows. The dataset provides detailed information on these flows, such as the sector of investment (governments, financial corporations, etc.) and the type of investment (equity, debt, etc.). In order to capture the entirety of financial flows, we download the variable measuring the overall investment of a country in assets of another country[8]. The variable is available for 15 years within the period 1997-2020, where the years 1998 and 1999 are not available. Table A.3 in Appendix A.3 provides summary statistics regarding our collected data.

Lastly, we supplement our dataset with a series on long-run interest rates. The rationale for including this series is that, since we are interested in financial flows, such a variable is expected to have an important informative power. In order to create this series, we collected information from different datasets, the most important ones being the Global Financial Dataset[9] and the Macrohistory Database[10]. Table A.5 in Appendix A.4 provides a detailed description of the data sources used to construct this series. Table A.4 provides summary statistics for our collected interest rate series.

---

[8]The variable we rely upon is "Total investment in foreign assets, Total Holdings", whose CPIS code is I_A_T_T_T_BP6_USD.T.T.

[9]Available at https://globalfinancialdata.com/insights.

[10]Available at https://www.macrohistory.net/database/.

1.3.2. Final Dataset

In our analysis, to be as close as possible to Timini (2018), we decide to focus on the sample of countries used in his analysis: Belgium, Denmark, Finland, France, Germany, Greece, Italy, Netherlands, Norway, Portugal, Spain, Sweden, Switzerland, United Kingdom[11]. Hence, we merge data from the three previously described sources, and restrict attention to these countries. As a consequence, our final dataset spans the period 1861-2014 (starting 4 years before the establishment of the LMU in 1865), includes 59 variables and has an overall size of 29,681 observations[12]. Starting from this dataset, we use the 1997-2014 sample to train our models in predicting bilateral financial flows, and use the 1945-2014 sample to train models in predicting trade flows for the model selection exercise (a more thorough description of these exercises is postponed to section 1.5).

1.4. Model Estimation

The goal is to design the best proxy for bilateral financial flows given the observables we have. This is a pure conditional prediction exercise that is well-suited for machine learning methods. The difficulty resides in preserving good out-of-sample performance despite the lack of historical data on financial flows to externally validate our predictions. From Kaggle data science competitions, XGBoost and LightGBM are supposed to perform best in a time series setting[13]. Yet, applications to economic history are slightly different from traditional time series forecasting exercises. It is possible that other models would actually perform better. The reason is economic historians are less interested in T steps ahead forecasts and more interested in predicting a variable over an entire historical period. Machine Learning models are complex objects and

[11]Timini (2018) includes Austria-Hungary in his sample. However, since we will be reconstructing financial flows data using post-WWII observations, and given that Austria-Hungary doesn't exist anymore, we don't have data for this country.
[12]Tables A.3 and Table A.4 report statistics of our newly assembled data. The remainder of the variables, coming from Tradehist, are thoroughly described in Fouquin and Hugot (2016).
[13]https://medium.com/analytics-vidhya/xgboost-lightgbm-and-other-kaggle-competition-favorites-6212e8b0e835.

it is therefore difficult to know a priori which one will do better. It is also essential that hyper-parameter tuning does not lead to over-fitting and preserves out-of-sample performance over long historical periods. The methodology developed in this paper and described in Algorithm 1 is grounded on two guiding principles to alleviate these concerns.

Algorithm 1 Cross-validation and model estimation

procedure ESTIMATION(N, X_o, X_n) ▷ X_o, X_n correspond to historical/modern data
  Split X_n sample in N period blocks
  for $F \in$ {set of ML models} **do** ▷ for Lasso, XGBoost, ...
    Create hyper parameter grid Δ_F
    for random draw $\delta \in \Delta_F$ **do**
      for $i \in N$ **do**
        Estimate model F_δ over $N \setminus \{i\}$ blocks ▷ Leave one out for cross-validation
        Compute cross-validation $R^2_{F_\delta(i)}$ over block i
      Compute average cross-validation score $R^2_{F_\delta(X_n)}$ over all blocks
    Select best hyper parameter $\delta^\star_F = \text{argmax}_\delta\, R^2_{F_\delta(X_n)}$
    Re-estimate model on full sample X_n with cross-validated hyperparameter $\delta^\star_F$
    Predict historical data using $F_{\delta^\star}(X_o)$
    Compute out of sample $R^2_{F_{\delta^\star}(X_o)}$ ▷ Possible only for a test variable
  Select best performing model out of sample $F^\star_{\delta^\star} = \text{argmax}_F\, R^2_{F_{\delta^\star}(X_o)}$

The first is to be agnostic regarding the "right" model and the "right" set of hyper-parameters to use in building our proxy variable. To account for this model uncertainty we benchmark 9 different models with potentially different strengths and weaknesses[14]. We also define a large hyper parameter grid space Δ_F. Using random grid search we explore a hundred hyper-parameter combinations for each model. This guarantees extensive grid search to find a hyper-parameter combination that is relatively close to the global optimum. Otherwise there would be a risk of false negatives, good models

[14] A description of each model and its characteristics is provided in the appendix A.5.

that are rejected by our algorithm because the right set of hyper-parameters has not been tried.

The second principle is to select our final model of choice to perform well even many decades prior to the available sample. The algorithm ensures that in two separate steps. First, we select hyper-parameters using KFold cross-validation. Practically, we split the sample of interest into 5 blocks. For each block, we compute a model prediction R^2 based on the estimation over the other 4 blocks. We average those into a cross-validation R^2 that measure how well the model can perform out-of-sample for a given set of hyperparameters. Hyper-parameters are thus selected so that the model has the highest average R^2 when predicting an out-of-sample block. This is the methodology that has been shown to perform best in the finance literature (Bryzgalova et al., 2019; Kaniel et al., 2021; Kozak et al., 2020). It is also better suited than time series split for our purpose given that we are less interested in step ahead forecasts. Figure 1.1 illustrates the difference between the two methods where year is the "sample index" of our sample[15].

Figure 1.1. Alternative Cross-validation Methods

A. KFold B. Time Series Split

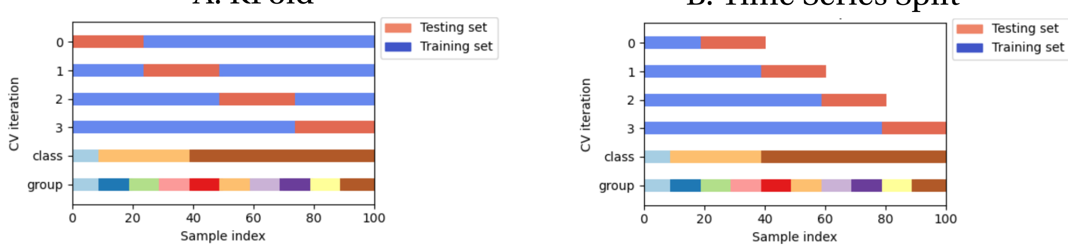

Source: https://scikit-learn.org/stable/modules/cross_validation.html

One possibility would be to simply pick the model with highest cross-validation R^2 and use it to build our proxy for financial flows. This is what is usually done for standard

[15]For a detailed discussion of the different cross-validation methods, the reader is referred to this article from scikit-learn developers https://scikit-learn.org/stable/modules/cross_validation.html.

time series exercises. Would that be enough to perform well with wide historical data? Simple KFold cross-validation guarantees that the model performs well out-of-sample, so long as the training set is not too far away in time from the evaluation set. When predicting historical data a century back, this methodology is likely to show its limits.

The second step is to select our final model of choice by comparing prediction performance far out-of-sample for a readily available historical variable. We choose a variable available for the entire 1861-2014 period and to be reconstructed for the 1861-1913 period. Since we want this exercise to be informative about the best performing model for bilateral financial flows, the test variable should be at the same disaggregated level and highly correlated with financial flows. As shown in figure A.5, bilateral trade flows is an important predictor of bilateral financial flows. We therefore train our models to predict bilateral trade flows on the 1945-2014 period. We use the same remaining observables and the same cross-validation procedure to predict the test variable and our variable of interest to make the comparison meaningful. Comparing our predictions with the actual data for the 1861-1913 period, we can obtain a measure of out-of-sample performance. Practically we select the model with highest out-of-sample R^2[16]. This guarantees that the model not only performs well a few years before the training sample, which is guaranteed by our Kfold cross-validation procedure, but also many decades before that. Doing so we pick the model that best captures long term trends and invariant economic relationships in the data, rather than a good forecasting model at shorter horizon but ill-suited to historical forecasting.

[16]This is equivalent to selecting the model based on the lowest RMSE criterion.

1.5. Model Selection

Starting from our nine forecasting models, we need to discriminate among them in order to evaluate which has the best forecasting power given the characteristics of our data.

Table 1.1. Performance on CPIS Financial Flows

	ET	RF	LGBM	NN	XGBoost	Ridge	Lasso	AdaBoost	SVM
R^2 (In-sample)	0.991	0.988	0.979	0.977	0.958	0.880	0.879	0.836	0.815
Folds	120	120	120	120	120	120	120	120	120
N	2483	2483	2483	2483	2483	2483	2483	2483	2483
Years	15	15	15	15	15	15	15	15	15

Notes: Regressors are ordered with decreasing in-sample R^2 values. "ET" stands for Extra Trees, "RF" stands for Random Forest, "NN" stands for Neural Network, "SVM" stands for Support Vector Machine. *Iterations* measures the number of iterations in our cross-validation exercise. N measures the number of folds available in the sample of our exercise. *Years* are the number of years we use to train our models (1997-2014, 1998 and 1999 are not available in the original IMF dataset).

Table 1.1 provides a summary of the performance of our models, which we ordered with decreasing R^2 values, while figure A.2 in Appendix A.6 provides a graphical representation of their performance. Two important points can be made looking at the table. First, all models perform fairly well in-sample, with R^2 values ranging between 0.815 for SVM, the worst performing model, to 0.991 for Extra Trees, the best performing model. Second, while the overall distance between the best and worst performing model is of 0.176, five of the nine models fall within a range of only 0.033 (ET, RF, LGBM, NN, XGBoost), so that their performance is almost identical. This table is informative about the capacity of the different models to fit the data in sample. And it is not surprisingly that most models do well given how flexible they are compared to a simple OLS. This is not however the way we select the "best" model.

Ideally we would like to rank our models based on their performance at predicting bilateral financial flows over the 1861-1913 period. While we cannot perform any out-of-sample exercise for the variable we are interested in forecasting due to the data limitations problem we are solving, we can evaluate our models on their performance at predicting bilateral trade flows over that same period. Based on these statistics, we choose which models to rely upon to estimate bilateral financial flows.

Table 1.2. Performance on Trade Flows

	Lasso	XGBoost	LGBM	AdaBoost	ET	RF	NN	Ridge	SVM
R^2 (In-sample)	0.963	0.989	0.989	0.932	0.994	0.988	0.989	0.966	0.778
R^2 (Out-sample)	0.531	0.529	0.313	0.296	0.260	0.213	0.205	-0.082	-2.566
Iterations	100	100	100	100	100	100	100	100	100
N	12381	12381	12381	12381	12381	12381	12381	12381	12381
Years	70	70	70	70	70	70	70	70	70

Notes: Regressors are ordered with decreasing out-of-sample R^2 values. "ET" stands for Extra Trees, "RF" stands for Random Forest, "NN" stands for Neural Network, "SVM" stands for Support Vector Machine. *Folds* measures the number of folds in our cross-validation exercise. N measures the number of observations available in the sample of our exercise. *Years* are the number of years we use to train our models (1945-2014).

Table 1.2 shows the in-sample and out-of-sample R^2 values of our models, while figures A.3 and A.4 in Appendix A.6 provide a graphical representation. The table, where models are ordered with decreasing out-of-sample R^2 values, shows the importance of relying on out-of-sample forecasts. Similarly to the statistics of Table 1.1, the in-sample performance of all models is very high, spanning from 0.994 for Extra Trees to 0.778 for SVM, a 0.216 difference. Yet, the picture that we get based on the out-of-sample R^2 is different: the ranking of the models changes, and the distance between their accuracy measures increases substantially. In particular, the two best performing models are Lasso and XGBoost, with R^2 values of 0.531 and 0.529, respectively. LGBM, the third-best model, has an R^2 that differs from that of XGBoost by 0.216, approximately the same difference that exists between the best and worst in-sample fit of all models. Extra Trees, the best in-sample performer, ranks fifth. The out-of-sample fit of some models (Ridge and SVM) is so mediocre that their R^2 values are negative.

Based on the results from table 1.2, we select Lasso and XGBoost as benchmark models to reconstruct bilateral financial flows: Lasso will be our preferred model, while

Figure 1.2. Out-of-Sample RMSE (Trade Flows)

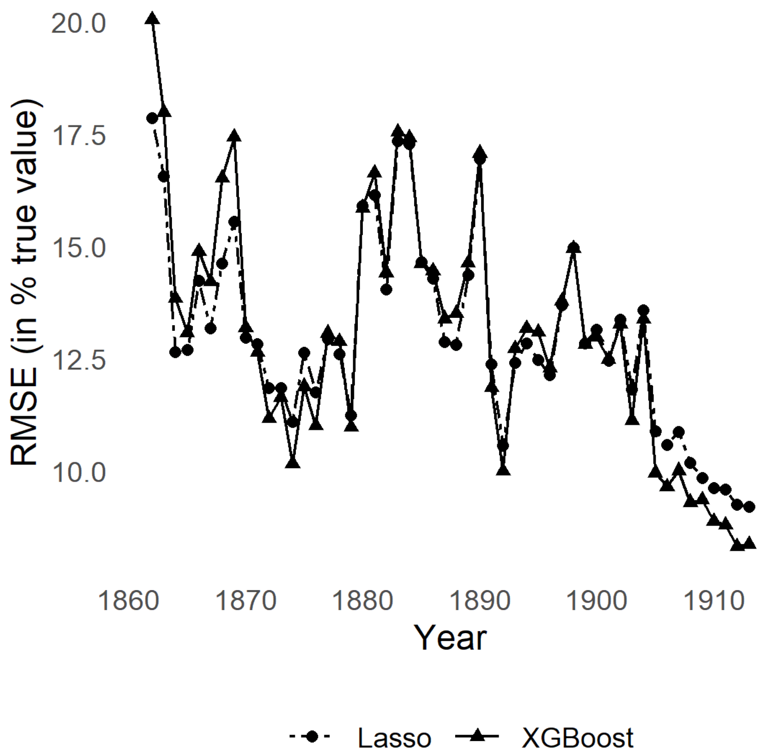

XGBoost will be used to implement a robustness exercise[17]. Even though our proxy variable cannot be a perfect measure, there are two reasons why we believe our two models will make reasonable predictions. First, their out-of-sample performance on trade flows, a structurally similar variable to financial flows, is high. This is shown not only by their out-of-sample R^2 values in table 1.2, but also by figure 1.2. The figure displays a measure of the average error in the yearly predictions of our models: the root of the mean squared error of trade flows predictions, expressed as a fraction of the average observed trade flows values. As we can see, with the exception of the very first year for XGBoost, the errors are always below 20% of the average yearly trade flows,

[17]For completeness, all of the results we show are also provided for the other models, and can be found in Appendix A.9.

and often below 15%[18]. Second, trade flows are an extremely important variable to forecast bilateral financial flows as shown in figure A.5 in Appendix A.7. Yet this important piece of information is dropped when predicting bilateral trade flows (to avoid autoregression), which suggests our models perform well even with a limited set of bilateral observables. This suggests our bilateral financial flow proxy benefits from an important extra variable, and possibly achieves higher prediction accuracy.

1.6. Consequences of the LMU for Financial Flows

After having reconstructed bilateral financial flows data using our Lasso and XGBoost models, we are ready to evaluate the effectiveness of the LMU on stimulating financial flows. As emphasized in the historical recollection of section 1.2, enhancing capital flows across members was an important reason for countries to join the Union. Unfortunately, data availability issues have not allowed researchers to investigate this dimension of the LMU so that, so far, the only focus has been on evaluating the impact that it had on trade flows. Thanks to our new methodology we can instead move on and address this question. In the following, we will first describe the empirical strategy we use to evaluate the impact of the LMU on bilateral flows. We will then show our results and, finally, implement a robustness exercise.

1.6.1. Empirical Strategy

In order to evaluate the impact of the LMU on bilateral financial flows, we rely on the best practice guidelines to implement structural gravity models compiled by the WTO Yotov et al. (2016). In particular, this implies that we will be using a Poisson regression, which is able to deal with zero flows values and is consistent with fixed-effects[19]; that we

[18]The figure provides an additional reason to prefer our Lasso model to XGBoost: as the chart shows, XGBoost tends to have higher RMSE relative to Lasso, especially in the first half of the sample. Since the LMU started in 1865, this is an important period for our analysis.

[19]All regressions are implemented using Stata's PPMLHDFE command Correia et al. (2020).

will include in our specification both directional time-varying fixed-effects and country-pair fixed-effects; and that we will use standard error clustered at the country-pair level. Accordingly, the main regression in our analysis is:

$$X_{i,j,t} = \beta_0 + \beta_1 LMU_{i,j,t} + \beta_2 GS_{i,j,t} + \beta_3 SMU_{i,j,t} + \gamma_{i,t} + \theta_{j,t} + \delta_{i,j} + \epsilon_{i,j,t} \tag{1.1}$$

where $X_{i,j,t}$ are our reconstructed bilateral financial flows, $LMU_{i,j,t}$ is a dummy variable equal to one when both country i and country j belong to the LMU at time t, $GS_{i,j,t}$ and $SMU_{i,j,t}$ are dummy variables equal to one when both countries belong to the Gold Standard and Scandinavian Monetary Union at time t, respectively (we include these two variables to be consistent with the specification for trade flows of Timini, 2018). $\gamma_{i,t}$, $\theta_{j,t}$, and $\delta_{i,j}$ capture importer time-varying, exporter time-varying, and country-pair fixed-effects.

Since the LMU was characterized by a country, France, that played a pivotal role, we follow Timini (2018) and additionally run the following regression:

$$\begin{aligned} X_{i,j,t} = &\beta_0 + \beta_1 LMU\_France_{i,j,t} + \beta_2 LMU\_Rest_{i,j,t} + \\ &\beta_3 GS_{i,j,t} + \beta_4 SMU_{i,j,t} + \gamma_{i,t} + \theta_{j,t} + \delta_{i,j} + \epsilon_{i,j,t} \end{aligned} \tag{1.2}$$

where $LMU\_France_{i,j,t}$ and $LMU\_Rest_{i,j,t}$ are dummy variables equal to one if flows among LMU countries involve France ($LMU\_France_{i,j,t}$) or not ($LMU\_Rest_{i,j,t}$). The idea of this regression is to test whether the LMU was particularly effective in stimulating flows between France and other members. Finally, we will run variations of these two main specifications including additional dummy variables to test whether the LMU was particularly effective during a sub-period of its entire existence. These will be the periods 1861-1885 and 1861-1874, both of which have been suggested by historians to be time frames during which the Union was particularly effective[20].

1.6.2. Results

Table 1.3 displays the results of our empirical exercise, where bilateral financial flows are estimated through Lasso, our preferred model. Since the 6 specifications reported in the table follow the main empirical exercises in Timini (2018) for trade flows, table A.7 in Appendix A.8 provides Timini (2018)'s results, the most recent on the effects of the LMU, for comparison.

The first two columns show the results of our main regressions, displaying the coefficients of equations 1.1 and 1.2, respectively. In column one, the LMU coefficient is positive and significant at the 5% level, with participation in the LMU being associated with an approximate 5% increase in bilateral financial flows. This represents the main result of this study on the effectiveness of the LMU of bilateral financial flows. Differently from the literature on the effectiveness of the LMU on trade flows (Flandreau, 2000; Timini, 2018), we find evidence in favor of a positive impact of the LMU on financial flows. Column 2 moves on to investigate whether the impact of the LMU was different when flows involved France. Both reported coefficients are positive, but only the one associated with flows not involving France is statistically significant. This result suggests that, during the entire 1865-1913 period, the Union was particularly effective in stimulating financial flows across these countries.

Columns 3 and 4 implement an exercise to evaluate whether the impact of the Union was greater during the 1865-1885 period since, as we discussed in section 1.2, historians have argued that the LMU ceased to *de facto* exist in 1885. To do so, we interact the

[20] As mentioned in section 1.2, Willis (1901) suggests that the LMU *de facto* ceased to exist in 1885. Moreover, Flandreau and Oosterlinck (2012) stress that in 1874 markets downgraded the possibility of bimetallism to last, so that 1874 can be seen as the earliest date in which the effectiveness of the Union started to decrease.

Table 1.3. Bilateral Financial Flows (Lasso)

	(1)	(2)	(3)	(4)	(5)	(6)
LMU	0.051*		-0.049***		-0.003	
	(0.021)		(0.014)		(0.008)	
LMU_France		0.047		-0.059*		-0.008
		(0.031)		(0.026)		(0.030)
LMU_Rest		0.087***		0.084***		0.097***
		(0.016)		(0.008)		(0.015)
LMU_1885			0.204***			
			(0.036)			
LMU_France_1885				0.222***		
				(0.045)		
LMU_Rest_1885				-0.033		
				(0.045)		
LMU_1874					0.147**	
					(0.048)	
LMU_France_1874						0.161**
						(0.055)
LMU_Rest_1874						-0.105
						(0.066)
GS	0.248***	0.247***	0.131**	0.124**	0.207***	0.203***
	(0.042)	(0.042)	(0.047)	(0.047)	(0.049)	(0.047)
SMU	-0.249***	-0.250***	-0.256***	-0.252***	-0.256***	-0.250***
	(0.048)	(0.045)	(0.015)	(0.036)	(0.031)	(0.035)
N	7169	7169	7169	7169	7169	7169

Notes: $^{***}p < 0.001, ^{**}p < 0.01, ^{*}p < 0.05$. Dependent variable: Estimated bilateral financial flows. All regressions include a constant, importer-year, exporter-year and importer-exporter fixed-effects. Clustered standard errors at the importer-exporter level.

LMU dummy with a dummy capturing the pre-1885 period. The coefficients in column 3 are both highly statistically significant and, similar to Timini (2018), of opposite sign. In particular, while the LMU coefficient is negative, the pre-1885 LMU coefficient is positive and of much larger magnitude, so that the overall LMU effect during this period is positive (LMU_1885 + LMU $\approx .155$). Importantly, this coefficient is larger than the one in column one, suggesting that the effectiveness of the LMU was indeed larger when we focus on the pre-1885 period. On the other hand, the negative LMU impact after 1885 may signal a deterioration of LMU members, which is in line with findings in Timini

(2018). Column 4 provides additional information regarding the results from column 3. Differently from the results in column 2, we can see that the overall impact of the LMU on stimulating flows involving France is positive (LMU_France_1885 + LMU_France $\approx$ $.163$) and statistically significant in the pre-1885 period, while it is negative after 1885. A different pattern is found when looking at flows among non-France LMU members: not much more affected in the pre-1885 period, but positive and statistically significant after throughout the entire LMU period. Overall, columns 3 and 4 give us additional insights on the effectiveness of the LMU: while columns 1 and 2 suggest an overall positive impact concentrated among non-France countries, columns 3 and 4 point to an even greater impact of the LMU and to a pivotal role of France in the pre-1885 period, and to a reduction in non-Franch flows from then onwards. It is important to point that such a pattern is similar to the one found in Timini (2018) for trade flows.

Finally, columns 5 and 6 report the estimates of an exercise similar to that of columns 3 and 4, but restricting attention to the 1865-1874 activity period of the LMU. The rationale for this further restriction is twofold. First, 1874 is the year in which markets started to downgrade the possibility of bimetallism to last Flandreau and Oosterlinck (2012), so that it could be considered the shortest possible period of actual existence of the LMU. Second, this represents the only period during which the LMU had an overall positive impact on trade flows according to Timini (2018).

Overall, despite minor changes in the magnitude of the coefficients, the story of these estimates is in line with that of columns 3 and 4: the effectiveness of the LMU was positive and larger during its first years; France was heavily responsible for these flows initially, while flows among other countries were constantly important in the century. However, it is important to point out that the magnitudes of the coefficients associated with the 1874 dummy are lower than those associated with the 1885 dummy in columns 3 and 4. Hence, restricting attention to the 1865-1874 period tends to decrease the importance of the LMU, implying that the Union kept being effective for an additional decade. This is important because signals a difference in financial flows pattern relative to trade flows: while the latter were stimulated only in the very first decade of existence of the Union, as shown in Timini (2018), the former were positively affected until 1885 (and thereafter among non-France members).

Finally, although this is not our focus of interest, we note that the coefficients on participation to the Gold Standard (GS) are always positive, statistically significant and fairly stable across specifications as we would expect. The coefficients on participation to the Scandinavian Monetary Union (SMU) are always negative, statistically significant and stable across specifications, similarly to the results of Timini (2018).

1.6.3. Robustness: XGBoost Results

In this section we evaluate the robustness of the main conclusions we reached in the previous section. In order to do so, we run our regressions using the bilateral financial flows as estimated through XGBoost, the second-best model according to our discussion in section 1.5.

Table 1.4 shows the results we obtain using these data. Confirming the results we obtained with Lasso data, columns 1 and 2 point to a positive and statistically significant impact of the LMU on financial flows during the entire 1861-1913 period, with an effect particularly pronounced for non-French flows. Comparing the magnitudes of these estimates, we can see that those of the statistically significant coefficients are very close to those of table 1.3.

Moving to columns 3 and 4, similarly to table 1.3, we see that the positive effects of the LMU tend to concentrate on the 1865-1885 period (column 3, LMU = 0.058) and that, during this time frame, the LMU was especially effective in stimulating financial flows with France as a counterpart (column 4). These results are therefore qualitatively

Table 1.4. Bilateral Financial Flows (XGBoost)

	(1)	(2)	(3)	(4)	(5)	(6)
LMU	0.046*		0.011		0.025	
	(0.019)		(0.019)		(0.019)	
LMU_France		0.005		-0.032***		-0.019
		(0.007)		(0.010)		(0.011)
LMU_Rest		0.082**		0.052		0.066*
		(0.031)		(0.036)		(0.031)
LMU_1885			0.058**			
			(0.019)			
LMU_France_1885				0.065***		
				(0.012)		
LMU_Rest_1885				0.051		
				(0.028)		
LMU_1874					0.055**	
					(0.021)	
LMU_France_1874						0.064***
						(0.015)
LMU_Rest_1874						0.042
						(0.036)
GS	-0.027*	-0.030**	-0.039**	-0.041***	-0.031*	-0.033**
	(0.011)	(0.011)	(0.012)	(0.012)	(0.013)	(0.012)
SMU	-0.026	-0.026	-0.020	-0.019	-0.021	-0.019
	(0.019)	(0.016)	(0.017)	(0.015)	(0.017)	(0.014)
N	7169	7169	7169	7169	7169	7169

Notes: $^{***}p < 0.001, ^{**}p < 0.01, ^{*}p < 0.05$. Dependent variable: Estimated bilateral financial flows. All regressions include a constant, importer-year, exporter-year and importer-exporter fixed-effects. Clustered standard errors at the importer-exporter level.

in line with those of the corresponding columns in table 1.3, but the magnitudes of these coefficients are lower. An additional difference is that, albeit positive, the coefficient on LMU_Rest is not significant using XGBoost data.

Finally, column 5 shows that, albeit less than before 1885 (column 3, LMU_1885 = 0.058), the LMU was effective during its early years (column 5, LMU_1874 = 0.055), while column 6 shows that it led to increased flows involving France up until 1974, and to increased flows involving other LMU members thereafter (column 6). Both of these results are qualitatively in line with the results in table 1.3, but their magnitudes are smaller.

Overall, although the estimates of bilateral financial flows coming from our XGBoost model are a second-best option, the results in columns 1 and 2 of table 1.4 tell us that we would have reached very similar conclusions regarding the effectiveness of the LMU during the 1865-1913 period if we had used this data. Moreover, columns 3 to 6 show that the qualitative conclusions we would have reached regarding the effectiveness of the LMU during different sub-periods would have been in line with those of our Lasso model.

Yet, it is important to stress two issues. First, as we mentioned in previous paragraphs, the magnitudes of the estimated effects are lower once we rely on this alternative model. Second, although this is not the focus of our exercise, the coefficients of participation to the Gold Standard and to the Scandinavian Monetary Union are different from those of table 1.4. In particular, the GS coefficients, although displaying a very low statistical significance, are negative[21]. Differently, the SMU coefficients lose their statistical significance.

1.7. Conclusion

This paper emphasizes that a lot more information and correlation patterns can be extracted from existing historical data. Machine learning models can extract that information in a systematic, comprehensive and replicable way, creating synthetic proxies for a wide range of variables that cannot be measured otherwise. Accordingly, bringing these methods into the economic history literature, similarly to what has been done in

[21]Importantly, the data we are using for our exercise on the LMU, excluding many non-European countries, such as the United States, are not well-suited to evaluated the overall effectiveness of the Gold Standard on financial flows. Accordingly, this variable is only introduced to control for potential omitted variables biases in our regressions. Nonetheless, this exercise points to an incongruency of our results depending on which proxy we use (Lasso vs. XGBoost).

other fields, could allow to tackle important research questions that tend to be neglected because of data availability issues.

One such example is the literature on the Latin Monetary Union, which has been concerned with trade flows precisely because of data availability issues. From both a theoretical perspective and the historical accounts at the time, the LMU was monetary and financial in nature. A natural exercise would have been to study the effect of the LMU on financial flows absent existing data limitations.

Relying on machine learning techniques, we were able to circumvent that data limitation by reconstructing a proxy for financial flows across 14 countries between 1861 and 1913. It makes possible the measurement of the impact of the Latin Monetary Union on the pattern of European financial flows through standard causal inference methods.

Our main finding is that, differently from what has been found for trade flows, the Latin Monetary Union did favor financial flows among its members, increasing bilateral financial flows by 5% between 1865 and 1913 and by approximately 15% when restricting attention to the 1865-1885 period, during which the Union was most active according to historical accounts. Moreover, we find that while flows heavily involved France until 1885, this stopped being the case thereafter, when flows began to concentrate among other member countries.

Overall, these results provide new insights about the history of the Latin Monetary Union, showing that it did help member countries achieve some of the goals that had pushed them to join the Union in the first place.

CHAPTER 2

Rigid Production Networks

Written with Alireza Tahbaz-Salehi[1]

[1]Prepared for Carnegie-Rochester-NYU Conference on Public Policy. We are grateful to Luigi Iovino, Ariel Zetlin-Jones, and conference participants for detailed comments and suggestions. We also thank Marios Angeletos, Vasco Carvalho, Stephen Morris, Alessandro Pavan, Ali Shourideh, and Venky Venkateswaran.

2.1. Introduction

Firms in any modern economy rely on a wide range of goods and services for production, while simultaneously serving as input suppliers of other producers in the economy. For example, as documented by Bernard et al. (2022), the median firm in Belgium uses inputs from 53 suppliers and sells to 26 different customers. The extent of such firm-to-firm linkages can become significantly more skewed for larger firms, as a firm like Airbus works with roughly 12,000 different suppliers that provide products and services for flying and non-flying parts.

While clearly an integral part of the production process, supply chain linkages between different firms can function as a source of macroeconomic risk, as disruptions due to natural disasters (Carvalho et al., 2021), wars (Korovkin and Makarin, 2022), or foreign trade shocks (Dhyne et al., 2022) can escalate from localized events into broader disruptions. Not surprisingly then, an extensive (and growing) body of work in macroeconomics studies how the economy's production network can serve as a mechanism for propagation and amplification of shocks.[2]

However, this literature, for the most part, abstracts from two realistic frictions. First, the benchmark models of production networks assume that firms can adjust their input and output quantities frictionlessly in response to shocks. This is despite the fact that, in reality, many production processes, especially in manufacturing, require advance planning and setting up production lines that cannot be ramped up or down instantaneously. Similarly, while firms may be able to produce using different mixes of inputs, switching from one mix to another may require reorganizing the production process, a potentially costly and time-consuming endeavor. Yet another factor that may reduce

[2]For example, see Long and Plosser (1983) and Acemoglu et al. (2012a). More recent examples include production network models with fairly general production functions (Baqaee and Farhi, 2019), endogenous technologies (Acemoglu and Azar, 2020), nominal rigidities (La'O and Tahbaz-Salehi, 2022; Rubbo, 2022), and extensive margin adjustments (Baqaee and Farhi, 2021; Acemoglu and Tahbaz-Salehi, 2020).

firms' short-term ability to drastically adjust their production process in response to shocks is the presence of significant lead times in acquiring inputs: according to the Institute for Supply Management, in January 2023, the average lead time for obtaining production materials for manufacturing firms in the United States was 87 days. In the same month, the average lead time for acquiring capital inputs—such as machinery, plant equipment, software and the like—was 166 days, with 23% of firms facing lead times exceeding one year (Institute for Supply Management, 2023a).

Of course, frictions in quantity adjustments may not be of first-order importance if firms can perfectly anticipate future shocks and can set up contingency plans accordingly. This brings us to the second friction missing from the benchmark models, which assume that firms make their production decisions under perfect information about all the shocks in the economy. But, this is not necessarily the case either, as it would require firms to not only have perfect foresight, but also obtain and process a large volume of information about a wide range of shocks.

The relevance of the above-mentioned frictions became more visible during the COVID-19 pandemic, as various industries experienced significant disruptions in their supply chains. For example, at the onset of the pandemic, U.S. automakers "underestimated demand for their products," and "expecting weak demand, they cancelled orders of semiconductors, an item with a long lead time and with a secular increase in demand from other industries" (Helper and Soltas, 2021). This led to significant disruptions and price increases in the motor vehicle sector in the later stages of the pandemic.

In this paper, we develop a model that incorporates quantity rigidities and informational frictions into an otherwise standard model of production networks, where (i) firms may be restricted in how effectively they can adjust their input quantities in response to changes in the economic environment, and (ii) they may have to choose those quantities with only incomplete information about the realizations of shocks. Specifically, we consider a multisector general equilibrium economy à la Long and Plosser (1983) and Acemoglu et al. (2012a) in which firms are linked to one another via input-output linkages and are subject to industry-level productivity and aggregate demand shocks. Firms in each industry use Cobb-Douglas production technologies with constant returns to transform labor and intermediate inputs into output. Additionally, as in Baqaee and Farhi (2022), we allow for downward nominal wage rigidities. However, in a departure from the rest of the literature, we assume that firms make (some or all of) their intermediate input decisions in the presence of incomplete information about the realization of supply and demand shocks. This modeling approach has two key implications. First, it ensures that by the time the shocks are realized (or observed), certain input decisions made by the firms are sunk and hence cannot be adjusted. Second, while firms in our model are subject to quantity adjustment frictions, they nonetheless optimally plan their production in anticipation of future shocks (subject to their information sets).

As our main theoretical result, we provide a system of equations that (implicitly) characterizes equilibrium prices and quantities in terms of model primitives, namely, the realized supply and demand shocks, the economy's production network structure, each industry's set of rigid and flexible inputs, and the information sets of all firms in the economy. Despite being implicit, this characterization result captures the key economic forces that are active in the model. Specifically, it shows that, when deciding on their quantities under incomplete information, firms need to make forecasts about prices charged by their suppliers as well as the quantity demanded by their customers. As such, our characterization result highlights that equilibrium prices and quantities depend on firms' expectations of shocks both upstream and downstream their supply chains.

We then apply our implicit characterization result to three specific environments that lend themselves to closed-form solutions. As our first case study, we consider an economy consisting of only a single rigid industry subject to informational frictions, with the remaining industries capable of adjusting their quantities with no frictions. Focusing on such an environment allows us to identify the role played by quantity rigidities and informational frictions at each industry separately, and in particular, identify the industries that can act as production bottlenecks for the rest of the economy. We find that these two frictions result in a reduction in aggregate output. This is because when firms make their intermediate input decisions under uncertainty about the realizations of productivity shocks, they find it optimal to rely more heavily on less volatile suppliers, even if this comes at the cost of forgoing more efficient ones. Additionally, our result indicates that, all else equal, a rigid industry functions as a tighter "production bottleneck" for the entire economy if it is simultaneously (i) an overall large supplier in the economy and (ii) an important direct or indirect customer of other firms in the economy.

We then apply our results to study how incomplete information and the frictions in quantity adjustments—which we refer to as *real rigidities*—change the mapping from supply and demand shocks to aggregate output and inflation. We find that, in the presence of real rigidities and informational frictions, the first-order impact of productivity shocks is dampened compared to the fully flexible benchmark; that this dampening effect is stronger the higher the degree of real rigidities; and that the extent to which shocks to an industry propagate to aggregate outcomes depends on the size of the rigid industry and its exposure to the shock. As for shocks to aggregate demand, we find that, compared to the benchmark without informational frictions, the real effect of aggregate demand shocks is dampened, a positive demand shock is inflationary, and that the magnitudes of both effects depend on the exact position of the rigid industry in the production network.

We follow up these results by focusing on two other information structures: one in which all firms in the economy observe the same public signal, and a more general case with an arbitrary information structure (albeit for a simplified production network structure). Focusing on these environments, we show how the effect of real rigidities and informational frictions can build up over the production network.

We conclude the paper with a simple quantitative assessment of our model's implications.

Related Literature. Our paper belongs to the literature on production networks, which explores the implications of the disaggregated structure of the economy for aggregate, macroeconomic outcomes. In addition to the papers already mentioned, some of the more recent works in this literature include Bigio and La'O (2020), Liu and Tsyvinski (2023), and Baqaee and Farhi (2022). See Carvalho and Tahbaz-Salehi (2019) and Baqaee and Rubbo (2022) for recent surveys.[3] We contribute to this literature by relaxing two of the key standing assumptions in most production network models: that firms can make decisions under complete information and can frictionlessly adjust their intermediate input decisions in response to changes in the economic environment. We study how incomplete information together with frictions in quantity adjustments change the mapping from supply and demand shocks to aggregate output and inflation.

Our approach in using incomplete information in modeling frictions builds on earlier works, such as Mankiw and Reis (2002), Woodford (2003), and Maćkowiak and Wiederholt (2009), among others. For the most part, this literature relies on incomplete information as a source of *nominal* rigidities, with the assumption that firms set their nominal prices without complete information about the economy's fundamentals. In contrast to the bulk of this literature, firms in our framework are subject to *real* rigidities and choose their input quantities in the presence of incomplete information. As such,

[3]Also see Barrot and Sauvagnat (2016), Boehm et al. (2019), and Carvalho et al. (2021) for empirical studies of the role of production networks in propagation and amplification of shocks.

our paper is more closely related to Angeletos et al. (2016) and Angeletos and La'O (2020), who consider models in which firms' incomplete information about the shocks is a source of both real and nominal rigidities. Our point of departure from these two papers is our focus on how real rigidities arising from firms' incomplete information interact with the economy's production network structure.

More closely related to our work are two recent papers that also explore the role of incomplete information in the context of supply chains. Kopytov et al. (2022a) develop a model of endogenous network formation to investigate how uncertainty about the productivity of suppliers impacts firms' choice of technology. As in our paper, one of the key tradeoffs faced by firms in their framework is that supply chain uncertainty induces firms to rely more heavily on less volatile inputs, even if this comes at the cost of forgoing more efficient one.[4] They key distinction, however, is in the two papers' modeling approach: Kopytov et al. (2022a) assume that firms choose their production technology under incomplete information, but can flexibly adjust their quantity demand in response to shocks. As such, and given the assumption of constant returns, firms in their framework only need to form forecasts about their marginal costs. In contrast, firms in our framework are forced to make (some or all of) their quantity decisions prior to observing the shocks. As a result, they not only need to form forecasts about their upstream prices, but also about their downstream demand.

The second related paper is the contemporaneous work of Bui et al. (2023), who introduce informational frictions into a multi-country, multi-sector model with global value chains. As in our paper, firms do not observe the shocks and only have access to imperfect signals about productivities of various industries and countries. However, whereas Bui et al. (2023) assume that firms choose their primary inputs (e.g., labor) under incomplete information, firms in our model need to choose their intermediate

[4] Also see Grossman et al. (2023) for a related mechanism in the context of global supply chains and in the presence of relationship-specific risk and country-wide supply disturbances.

inputs before learning the realizations of the shocks. This distinction is consequential: since in Bui et al. (2023) firms face no frictions in choosing their intermediate inputs, the equilibrium only depends on firms' expectations of their suppliers' decisions. In contrast, in our model, the equilibrium depends not only on firms' forecasts of their suppliers' forecasts, but also on their forecasts of their customers' forecasts.

Finally, our paper is related to the growing body of works that studies network interactions in the presence of incomplete information. Examples include Calvó-Armengol et al. (2015), de Martí and Zenou (2015), Bergemann et al. (2017), and Golub and Morris (2018). We complement this literature, which is mostly focused on reduced-form games over networks, by studying a micro-founded, general equilibrium macro model where firms' decisions and outcomes are interlinked with one another as a result of the economy's disaggregated production network structure.

Outline. The rest of the paper is organized as follows. Section 2.2 sets up the environment and defines the equilibrium concept. Section 2.3 contains the main characterization result of the paper, where we show how informational frictions and real rigidities shape equilibrium prices and quantities. In Section 2.4, we consider a few special cases of the general model that lend themselves to explicit characterizations. We present a quantitative analysis of the model in Section 2.5. All proofs and some additional technical details are presented in the Online Appendix.

2.2. Model

In this section, we present a multisector model that forms the basis of our analysis. The model, which is in the spirit of general equilibrium models of Long and Plosser (1983) and Acemoglu et al. (2012a), closely follows the framework in La'O and Tahbaz-Salehi (2022). As our main point of departure from the prior literature, we assume that firms may have to make some of their intermediate input quantity decisions under incomplete information about the realizations of shocks.

2.2.1. Firms and Production

Consider an economy consisting of n industries indexed by $i \in \mathcal{N} = \{1, 2, \ldots, n\}$. Each industry consists of two types of firms: (i) a unit mass of monopolistically-competitive firms, indexed by $k \in [0, 1]$, producing differentiated goods and (ii) a competitive producer whose sole purpose is to aggregate the industry's differentiated goods into a single sectoral output. The output of each industry can be either consumed by the households or used as an intermediate input for production by firms in other industries.

The monopolistically-competitive firms within each industry use a common constant-returns-to-scale technology to transform labor and intermediate inputs into their differentiated products. More specifically, the production function of firm $k \in [0, 1]$ in industry i is given by

$$y_{ik} = z_i \zeta_i l_{ik}^{\alpha_i} \prod_{j=1}^{n} x_{ij,k}^{a_{ij}}, \tag{2.1}$$

where y_{ik} is the firm's output, l_{ik} is the firm's labor input, $x_{ij,k}$ is the quantity of sectoral commodity j purchased by the firm, and z_i is an industry-specific productivity shock. The constant $\alpha_i > 0$ denotes the share of labor in industry i's production technology, $a_{ij} \geq 0$ parameterizes the importance of good j in the production technology of firms in industry i, and $\zeta_i = \alpha_i^{-\alpha_i} \prod_{j=1}^{n} a_{ij}^{-a_{ij}}$ is a normalization constant. As is standard in this literature, we summarize input-output linkages in this economy by matrix $\mathbf{A} = [a_{ij}]$, which with some abuse of terminology, we refer to as the economy's *input-output matrix*. We also define the economy's *Leontief inverse* as $\mathbf{L} = (\mathbf{I} - \mathbf{A})^{-1}$, whose (i, j) element captures the role of industry j as a direct or indirect intermediate input supplier to industry i. Throughout the paper, we normalize the steady-state value of all (log) productivity shocks to 0, i.e., $\log z_i^{\text{ss}} = 0$ for all i.

Given the production technology (2.1), the nominal profits of firm k in industry i are given by

$$\pi_{ik} = (1-\tau_i)p_{ik}y_{ik} - wl_{ik} - \sum_{j=1}^{n} p_j x_{ij,k}, \tag{2.2}$$

where p_{ik} is the nominal price charged by the firm, p_j is the nominal price of industry j's sectoral output, w denotes the nominal wage, and τ_i is an industry-specific revenue tax or subsidy levied by the government.

As already mentioned, each industry also contains a competitive producer, which transforms the differentiated products produced by the unit mass of firms in that industry into a sectoral good using a constant-elasticity-of-substitution (CES) production technology, with elasticity of substitution $\theta_i > 1$:

$$y_i = \left(\int_0^1 y_{ik}^{(\theta_i - 1)/\theta_i} \mathrm{d}k \right)^{\theta_i/(\theta_i - 1)}.$$

The sole purpose of this producer is to ensure that each industry produces a single sectoral good, while at the same time allowing for monopolistic competition among firms within the same industry. Throughout the paper, we assume that the industry-specific tax in (2.2) is set to $\tau_i = 1/(1-\theta_i)$. As is well-known, this choice undoes the effect of monopolistic markups and ensures that the distortions in the economy are not due to firms' market power.

2.2.2. Households

In addition to the firms, the economy consists of a representative household, with preferences

$$U(C,L)=\log C-\chi\frac{L^{1+1/\eta}}{1+1/\eta},$$

where C and L denote the household's consumption and total labor supply, respectively, η is the Frisch elasticity of labor supply, and $\chi>0$ is a constant that parameterizes the representative household's disutility of labor supply. The representative household's final consumption basket is a Cobb-Douglas aggregator of the sectoral goods produced in the economy,

$$C=\prod_{i=1}^{n}(c_i/\beta_i)^{\beta_i},$$

where c_i is the amount of good i consumed and $(\beta_1,\dots,\beta_n)$ are nonnegative constants that measure various goods' shares in the household's consumption basket, normalized such that $\sum_{i=1}^{n}\beta_i=1$. The representative household's budget constraint is therefore given by

$$PC=\sum_{i=1}^{n}p_ic_i=w\sum_{i=1}^{n}\int_0^1 l_{ik}\mathrm{d}k+\sum_{i=1}^{n}\int_0^1\pi_{ik}\mathrm{d}k+T,$$

where $P=\prod_{i=1}^{n}p_i^{\beta_i}$ is the nominal price of the household's consumption bundle, w denotes the nominal wage, π_{ik} is given by (2.2), and T denotes lump-sum transfers from the government. To ensure that the government's budget constraint is satisfied, we assume that $T=\sum_{i=1}^{n}\tau_i\int_0^1 p_{ik}y_{ik}\mathrm{d}k$.

Finally, we assume that the representative household is subject to the following cash-in-advance constraint:

$$PC = m,$$

where m denotes the nominal aggregate demand in the economy. In what follows, we interpret a decrease in m as a negative aggregate demand shock. While the most natural source of such a shock is a monetary policy shock (say, due to a reduction in money supply), as shown by Baqaee and Farhi (2022) in a simple dynamic extension of the model, a decrease in expected future output, an increase in nominal interest rate, an increase in the household's discount factor, or a decrease in future prices—which can be thought of as a proxy for forward guidance—can all generate effects that are isomorphic to a decrease in m.

2.2.3. Real Rigidities and Informational Frictions

While the benchmark models of production networks assume that firms can adjust their input and output quantities in response to supply and demand shocks, in reality, many firms may have limited ability to do so, at least in the short run. For example, many production processes, especially in manufacturing, require advance planning, setting up production lines that cannot be ramped up or down instantaneously, or acquiring inputs with significant lead times.[5] Similarly, while firms may be able to produce using different mixes of inputs, switching from one mix to another may require reorganizing

[5]As pointed out in the Introduction, such lead times can be particularly significant for capital inputs, such as machinery, plant equipment, and software. According to the Institute of Supply Management (ISM), in January 2023, only 15% of firms in the U.S. manufacturing sector faced lead times that were less than 30 days, whereas 59% of firms faced lead times that exceeded 6 months (Institute for Supply Management, 2023a). The ISM defines lead time as "the time that elapses from placement of an order until receipt of an order, including time for order transmittal, processing, preparation and shipping."

the production process, a potentially costly and time-consuming endeavor. Additionally, firms may not even have access to all the relevant information that would be necessary for adopting their production plans in response to economic disturbances.

We model the presence of such frictions by following Angeletos et al. (2016) and Angeletos and La'O (2020) and assuming that firms make (some or all of) their intermediate input decisions under incomplete information about supply and demand shocks. Specifically, we assume that the economy lasts for two periods, $t \in \{0, 1\}$. At $t = 0$, firms in industry i receive a common signal $\omega_i \in \Omega_i$ about the realizations of the supply and demand shocks (z, m), where $z = (z_1, \ldots, z_n)$. Given ω_i, each firm k in industry i chooses the intermediate input quantities $x_{ij,k}$ for any input $j \in \mathcal{R}_i$ at $t = 0$, where $\mathcal{R}_i \subseteq \mathcal{N}$ denotes the set of *rigid* inputs of industry i. The productivity and demand shocks are then observed by all firms at $t = 1$, which is when firms set prices, choose their labor input l_{ik}, and choose the remainder of their intermediate inputs, $\{x_{ij,k}\}_{j \in \mathcal{F}_i}$, where $\mathcal{F}_i = \mathcal{N} \setminus \mathcal{R}_i$ denotes the set of *flexible* inputs of industry i. Production and consumption also take place at $t = 1$.

A few remarks are in order. First, note that since firms choose their rigid intermediate inputs before the realization of shocks, these input choices are subject to a measurability constraint: $\{x_{ij,k}\}_{j \in \mathcal{R}_i}$ can be contingent on ω_i, but not on (z, m). While related, this measurability constraint on quantities is distinct from measurability constraints on nominal prices, which are a source of *nominal* rigidities as opposed to real ones.[6] Second, it is immediate to see that if either (i) all firms observe perfectly informative signals about the realizations of the shocks, i.e., $\omega_i = (z, m)$ for all i; or (ii) all inputs of all firms are flexible, i.e., $\mathcal{F}_i = \mathcal{N}$ for all i, then the above framework reduces to a standard production network model, such as Acemoglu et al. (2012a). Third, note that firms face no frictions in adjusting their labor input, as we assume they choose l_{ik} at $t = 1$. This

[6]See Mankiw and Reis (2002), Maćkowiak and Wiederholt (2009), and La'O and Tahbaz-Salehi (2022) for examples of models with informational frictions as a source of nominal rigidities.

is to ensure that at least one input is free to adjust in response to realized demand, as otherwise markets may fail to clear.[7]

We conclude this discussion by introducing a measure for firms' uncertainty about the shocks' realizations. For any given pair of industries i and j, define

$$\kappa_{ij} = \frac{\mathbb{E}[\text{var}_i(\log z_j)]}{\text{var}(\log z_j)}, \tag{2.3}$$

where $\log z_j$ is the log productivity shock to industry j, $\text{var}_i(\cdot)$ denotes the variance conditional on the information set of firms in industry i, and $\mathbb{E}[\cdot]$ and $\text{var}(\cdot)$ denote the unconditional expectation and variance operators, respectively. The interpretation of κ_{ij} as a measure of uncertainty is fairly natural: it captures the (ex ante) volatility of $\log z_j$ conditional on i's information set as a fraction of its unconditional volatility. By the law of total variance, κ_{ij} is always in the unit interval, $[0, 1]$, and obtains its maximum value of 1 if firms in industry i receive no informative signals about the realization of $\log z_j$ (in which case, $\text{var}(\mathbb{E}_i[\log z_j]) = 0$). At the other end of the spectrum, $\kappa_{ij} = 0$ if firms in industry i face no uncertainty about the shock to industry j (i.e., $\text{var}_i(\log z_j) = 0$).[8]

We can define a similar object to measure firms' uncertainty about the realization of the demand shock:

$$\mu_i = \frac{\mathbb{E}[\text{var}_i(\log m)]}{\text{var}(\log m)}, \tag{2.4}$$

[7]This is also the key modeling distinction between our framework and that of Bui et al. (2023), who assume that labor input decisions are made under incomplete information, while all intermediate input quantities can adjust freely in response to shocks.

[8]In the special case that all (log) shocks and signals are normally distributed, κ_{ij} takes a familiar form in terms of the signal-to-noise ratio. In particular, suppose firm i observes a single signal given by $\omega_i = \log z_j + \epsilon_i$, where $\log z_j$ and ϵ_i are independent and normally distributed with variances σ_z^2 and σ_ϵ^2, respectively. In that case, $\kappa_{ij} = \sigma_\epsilon^2/(\sigma_z^2 + \sigma_\epsilon^2)$. The expression in (2.3) generalizes this concept to any arbitrary joint distribution of shocks and signals.

where m is the nominal aggregate demand. Once again, $\mu_i = 0$ if firms in industry i face no uncertainty about aggregate demand shocks, whereas $\mu_i = 1$ if they receive no informative signals about the realization of m.

2.2.4. Downward Nominal Wage Rigidities

While firms can set their prices p_{ik} flexibly at $t = 1$ after observing productivity and demand shocks, we allow for downward nominal wage rigidities by assuming that the nominal wage w cannot fall below an exogenously-specified value $\bar{w}$. This restriction on nominal prices means that there are two possibilities. One possibility is that $w > \bar{w}$ and the labor market clears. The other possibility is that the constraint on the nominal wage binds (so that $w = \bar{w}$), in which case the labor market is slack and does not clear, in the sense that the total demand for labor falls short of what the representative household is willing to supply at that wage. Taken together, we say the labor market is in equilibrium if

$$(w - \bar{w})\left(L - \sum_{i=1}^{n} \int_0^1 l_{ik} \mathrm{d}k\right) = 0, \qquad w \geq \bar{w}, \qquad L \geq \sum_{i=1}^{n} \int_0^1 l_{ik} \mathrm{d}k, \tag{2.5}$$

where L denotes the household's labor supply and l_{ik} is the labor demand of firm k in industry i. Clearly, the special case that $\bar{w} = 0$ corresponds to an economy with no nominal rigidities.

2.2.5. Equilibrium

With the various model ingredients in hand, we are now ready to define our solution concept.

Definition 1. An *equilibrium* is a collection of nominal prices, nominal wage, and quantities such that

(i) at $t = 0$, monopolistically-competitive firms in each industry choose their rigid intermediate input quantities to maximize expected real value of their profits given their information;

(ii) at $t = 1$, firms set their nominal prices and choose their labor and flexible intermediate inputs to maximize profits, taking the realized demand and their rigid input quantities as given;

(iii) the competitive producer in each industry chooses inputs to maximize its profits given prices;

(iv) the representative household chooses consumption and labor supply to maximize utility subject to its budget constraint;

(v) the labor market is in equilibrium, i.e., condition (2.5) is satisfied;

(vi) all sectoral good markets clear, i.e.,

$$y_i = c_i + \sum_{j=1}^{n} \int_0^1 x_{ji,k} \mathrm{d}k \qquad \text{for all } i \in \mathcal{N}. \tag{2.6}$$

Equilibrium conditions (ii)–(vi) are all standard—capturing firm and household optimizing behavior and consistency restrictions on quantities—with the measurability constraints on the rigid quantities captured by condition (i). Note that while firms in our model are subject to quantity adjustment frictions, they nonetheless optimally plan their production process in anticipation of future shocks subject to their information sets.

2.3. Equilibrium Characterization

In this section, we characterize the equilibrium in terms of model primitives, namely, the economy's production network structure, the set of rigid and flexible intermediate inputs, and the information sets of firms in each industry. We do so via backward induction.

Starting with decisions at $t = 1$, recall that firms optimally choose their labor and flexible intermediate input quantities to meet the realized demand. This means that firm k in industry i faces the following cost-minimization problem:

$$(2.7)\qquad \begin{aligned} (l_{ik}, \{x_{ij,k}\}_{j\in\mathcal{F}_i}) \in \arg\min \quad & wl_{ik} + \sum_{j\in\mathcal{F}_i} p_j x_{ij,k} \\ \text{s.t.} \quad & y_{ik} = z_i \zeta_i l_{ik}^{\alpha_i} \prod_{j=1}^{n} x_{ij,k}^{a_{ij}}, \end{aligned}$$

while taking prices, realized demand, and rigid intermediate input quantities as given. Next, turning to the firm's price-setting decision, the firm sets its nominal price optimally to maximize profits while taking its rigid intermediate inputs and all other nominal prices as given, that is,

$$(2.8)\qquad p_{ik} \in \arg\max \quad (1-\tau_i)p_{ik}y_{ik} - wl_{ik} - \sum_{j\in\mathcal{F}_i} p_j x_{ij,k}$$

subject to the demand curve $y_{ik} = (p_{ik}/p_i)^{-\theta_i} y_i$ and the labor and flexible intermediate input optimality condition (2.7). Finally, at $t = 0$, firm k in industry i chooses its rigid intermediate input quantities to maximize expected real value of its profits given its information, that is,

$$(2.9)\qquad (\{x_{ij,k}\}_{j\in\mathcal{R}_i}) \in \arg\max \quad \mathbb{E}_i\left[\frac{U'(C)}{P}\left((1-\tau_i)p_{ik}y_{ik} - wl_{ik} - \sum_{j=1}^{n} p_j x_{ij,k}\right)\right],$$

where $\mathbb{E}_i[\cdot]$ denotes the expectation operator with respect to the information set of firms in industry i, $U'(C) = 1/C$ is the household's marginal utility, and P is the nominal price of the consumption good bundle.

Given the above, we can characterize the equilibrium by solving for the optimization problems (2.7)–(2.9) recursively and imposing the market clearing condition (2.6) for all industries i. To present the equilibrium characterization result, let $\lambda_i = p_i y_i / PC$ denote industry i's *Domar weight*, defined as its sales as a fraction of GDP. Given that all firms in the same industry have identical technologies and information sets, we can drop the firm index k in our characterization. We have the following result:

Proposition 1. *Equilibrium nominal prices and Domar weights solve the system of equations:*

$$p_i = \frac{1}{z_i} w^{\alpha_i} \prod_{j \in \mathcal{F}_i} p_j^{a_{ij}} \prod_{j \in \mathcal{R}_i} \left(m \frac{\mathbb{E}_i[p_j/m]}{\mathbb{E}_i[\lambda_i]/\lambda_i} \right)^{a_{ij}} \tag{2.10}$$

and

$$\lambda_i = \beta_i + \sum_{j:i \in \mathcal{F}_j} a_{ji} \lambda_j + \sum_{j:i \in \mathcal{R}_j} a_{ji} \mathbb{E}_j[\lambda_j] \frac{p_i/m}{\mathbb{E}_j[p_i/m]} \tag{2.11}$$

for all $i \in \mathcal{N}$, where m is the nominal aggregate demand and w is the nominal wage.

Proposition 1 provides a system of $2n$ equations and $2n$ unknowns that expresses sectoral Domar weights $(\lambda_1, \dots, \lambda_n)$ and nominal prices $(p_1, \dots, p_n)$ in terms of the firms' information sets, the realized productivity shocks, the nominal wage, and nominal aggregate demand, m.

Focusing on equation (2.10), it is easy to verify that if firms in industry i are not subject to real rigidities—either because all their inputs are flexible or because they receive completely informative signals—then (2.10) reduces to $p_i = \frac{1}{z_i} w^{\alpha_i} \prod_{j=1}^{n} p_j^{a_{ij}}$. In other words, the nominal price of industry i is simply equal to its nominal marginal cost, as anticipated. More generally, however, as equation (2.10) indicates, the nominal price of industry i depends on industry i's expectation of the prices of its rigid intermediate inputs relative to nominal aggregate demand, $\mathbb{E}_i[p_j/m]$, as well as its expectation of its own equilibrium Domar weight, $\mathbb{E}_i[\lambda_i]$. To see the intuition for these dependencies, note that either an increase in $\mathbb{E}_i[p_j/m]$ or a decrease in $\mathbb{E}_i[\lambda_i]$ result in a reduction in the quantity x_{ij} of good j that firms in industry i demand at $t = 0$. Given that this quantity is sunk by the time firms set their prices at $t = 1$, a lower x_{ij} is akin to a lower productivity from the point of view of firms at $t = 1$, thus inducing firms in industry i to set a higher nominal price.

The intuition underlying equation (2.11) is similar. Recall that the Domar weight of industry i—which represents that industry's size in equilibrium—would be larger when it faces higher demand from its downstream customers (larger x_{ji}'s). Also recall that the demand from a customer j that is subject to real rigidities is increasing in $\mathbb{E}_j[\lambda_j]$ and is decreasing in $\mathbb{E}_j[p_i/m]$. Therefore, as is evident from (2.11), λ_i increases in its customers' expectations of their size and decreases in their expectations of i's price. Finally, note that, if none of i's customers are subject to quantity adjustment or informational frictions, then equation (2.11) implies that $\lambda_i = \beta_i + \sum_{j=1}^{n} a_{ji}\lambda_j$, as would be the case in the benchmark models of production networks (Carvalho and Tahbaz-Salehi, 2019; Baqaee and Rubbo, 2022).

As already mentioned, Proposition 1 characterizes equilibrium nominal prices and Domar weights in terms of the nominal wage and nominal aggregate demand. Given prices and Domar weights, one can then characterize the entire allocation in terms of m and w. In particular, from the definition of λ_i, it follows immediately that the output of industry i is given by $y_i = \lambda_i m/p_i$. Also, as we show in the proof of Proposition 1, industry i's demand for its flexible and rigid intermediate inputs are given by

$$x_{ij} = \begin{cases} a_{ij} m \lambda_i / p_j & \text{if } j \in \mathcal{F}_i \\ a_{ij} \mathbb{E}_i[\lambda_i] / \mathbb{E}_i[p_j/m] & \text{if } j \in \mathcal{R}_i, \end{cases} \tag{2.12}$$

respectively. Finally, the household's first-order conditions imply that $c_i = \beta_i m/p_i$.

To see how quantity adjustment and informational frictions shape firms' sourcing decisions, it is instructive to consider the log-quadratic approximation to equation (2.12) for inputs $j \in \mathcal{R}_i$:

$$\log x_{ij} = \log a_{ij} + \mathbb{E}_i[\log s_i] + \frac{1}{2}\text{var}_i(\log s_i) + \text{cov}_i(\log s_i, \text{sdf})$$
$$- \mathbb{E}_i[\log \hat{p}_j] - \frac{1}{2}\text{var}_i(\log \hat{p}_j) - \text{cov}_i(\log \hat{p}_j, \text{sdf}),$$

where $s_i = p_i y_i / P$ is the real sales of industry i, $\hat{p}_j = p_j / P$ is the real price of good j, $\text{sdf} = \log U'(C)$ is the log stochastic discount factor (SDF), and $\text{var}_i(\cdot)$ and $\text{cov}_i(\cdot, \cdot)$ denote, respectively, the variance and covariance operators conditional on the information sets of firms in industry i. A few observations follow. First, note that $\log x_{ij}$ is not only decreasing in the expected value of the (log) real price of industry j, but also in its variance. This reflects the fact that the two frictions induce firms in industry i to rely more heavily on suppliers with less volatile prices, even if this comes at the cost of forgoing inputs that are cheaper in expectation. Second, observe that $\log x_{ij}$ is decreasing in the covariance of the log price of good j with the log SDF. This term reflects the fact that firms in industry i reduce their demand for input j if it tends to be more expensive in the states of the world with a high marginal utility of consumption. Finally, the term $\text{cov}_i(\log s_i, \text{sdf})$ indicates that firms in industry i increase their demand for all inputs if they have higher sales in the states of the world with high marginal utility.

As a final remark on Proposition 1, we note that while nominal aggregate demand m is a model primitive (and a proxy for demand shocks), the nominal wage w is an endogenous object that is determined in equilibrium. The following simple lemma completes the characterization of equilibrium by providing an additional equation that expresses the nominal wage in terms of nominal aggregate demand, the minimum nominal wage, and sectoral Domar weights:

Lemma 1. *The nominal wage is given by*

$$w = \max \left\{ m\chi^{\frac{\eta}{1+\eta}} \left(\sum_{i=1}^{n} \alpha_i \lambda_i \right)^{\frac{1}{1+\eta}}, \bar{w} \right\}. \tag{2.13}$$

Furthermore, in the special case that labor supply is fully elastic, $w = \max\{\chi m, \bar{w}\}$.

Taken together, Proposition 1 and Lemma 1 provide a complete (albeit implicit) characterization of equilibrium in the presence of informational frictions and real rigidities. Unfortunately, in general—and unless one imposes some discipline on the economy's information structure—system of equations (2.10)–(2.13) does not lend itself to a closed-form solution. This is due to two sources of complexity in the model. First, as in Golub and Morris (2018), La'O and Tahbaz-Salehi (2022), and Bui et al. (2023), the presence of network interactions in the economy means that the equilibrium depends not only on the firms' first-order expectations, but also on their expectations of higher order. This is because firms need to forecast their Domar weights and input prices, which depend not only on the firms' own forecasts of the realized shocks, but also on their forecasts of other firms' forecasts, and so on. Second, and in contrast to the prior literature, the relevant higher-order expectations do not have a simple iterative representation in terms of cross-sectional (weighted) averages of firms' lower-order expectations. To see this, note that, according to (2.10), the nominal price set by firms in industry i depends on i's expectation of its input prices—an object that depends on the actions of its *upstream* suppliers—as well as on i's expectation of its own Domar weight—an object that is determined by the demand from its *downstream* customers. This means that the equilibrium depends on iterations of expectations both upstream and downstream over the network, significantly complicating how informational frictions interact with the production network structure.[9]

[9] See the economy in Subsection 2.4.3 for a more detailed discussion.

To explore the implications of Proposition 1 in a transparent manner, in the next section we study various special cases of the general setting in Section 2.2 by focusing on particular information structures. However, before doing so, we conclude this section by studying the equilibrium's efficiency properties.

Recall that if either (i) all firms observe perfectly informative signals about the realizations of the shocks or (ii) all inputs of all firms are flexible, then the model in Section 2.2 reduces to a standard production network model, where all firms price at marginal cost.[10] This means that in the absence of informational frictions and real rigidities—and as long as the downward nominal wage rigidity constraint does not bind—the equilibrium is (first-best) efficient. However, such a strong efficiency result no longer holds if frictions limit firms' ability to adjust their input quantities in response to shocks, as a planner could improve welfare by changing firms' input and output quantities contingent on the shocks' realizations. Nonetheless, our next result establishes that, despite the frictions, the equilibrium remains *constrained* efficient in the sense of Angeletos and Pavan (2007): a planner who is subject to the same informational and quantity rigidity frictions cannot improve upon the equilibrium welfare.

Proposition 2. *If the downward nominal wage rigidity constraint does not bind, then the equilibrium is constrained efficient.*

This proposition extends the constrained efficiency result of Angeletos et al. (2016) to an economy with non-trivial input-output linkages and an arbitrary specification of rigid and flexible inputs. It establishes that, in the absence of nominal rigidities, the equilibrium remains constrained efficient irrespective of the economy's production network structure, the set of rigid and flexible inputs of each industry, and the particular

[10] Even though firms in each industry are monopolistically competitive, setting the industry-specific taxes in (2.2) to $\tau_i = 1/(1-\theta_i)$ undoes the effect of monopolistic markups and ensures that there are no distortions due to market power.

information structure. Proposition 2 also establishes that, despite the complex nature of interactions between firms' expectations over the production network captured by equations (2.10) and (2.11), more precise information unambiguously improves equilibrium welfare: since more precise information improves welfare in the planner's solution, it would also do so in equilibrium.

2.4. Closed-Form Results

As discussed in the previous section, unless one imposes some discipline on either the economy's information structure or its production network architecture, equilibrium conditions (2.10)–(2.13) do not lend themselves to closed-form solutions. Therefore, to explore the implications of Proposition 1 and Lemma 1 in a transparent manner, we next focus on various special cases of the general setting in Section 2.2 that allow us to explicitly characterize the equilibrium in terms of model primitives.

2.4.1. Frictions in a Single Industry

As a first special case of the general setting in Section 2.2, we assume that only one single industry is subject to the quantity adjustment and informational frictions. Specifically, we assume that firms in industry r are the only firms in the economy with incomplete information about the realizations of productivity shocks $(z_1, \dots, z_n)$ and the aggregate demand shock, m. Firms in all other industries are not subject to real rigidities and make all their decisions at $t = 1$, that is, $\mathcal{R}_i = \varnothing$ for all $i \neq r$. Focusing on this special case allows us to identify the role played by real rigidities and informational frictions at each industry separately.

We also impose the following assumption on the economy's production network structure:

Assumption 1. $\ell_{ir}\ell_{ri} = 0$ for all industries $i \neq r$.

To interpret the above assumption, recall that the (i,j) element of the economy's Leontief inverse, ℓ_{ij}, captures the extent to which industry i relies on industry j as a (direct or indirect) input supplier. Therefore, under Assumption 1, there is no industry i in the economy that is simultaneously upstream and downstream to r. We impose this assumption to tease out the role of upstream and downstream relationships vis-à-vis industry r in the most transparent manner. Finally, to investigate the impact of supply shocks separately from that of demand shocks, we first assume that the downward nominal wage rigidity constraint does not bind, making shocks to nominal aggregate demand neutral for real outcomes. We have the following result.

Proposition 3. *Suppose r is the only rigid industry, Assumption 1 is satisfied, and the downward nominal wage rigidity constraint does not bind.*

(a) Then, to a first-order approximation,

$$\log C = \log C^* - \frac{\lambda_r^{\text{ss}}}{1 + 1/\eta\left(1 - \lambda_r^{\text{ss}} \sum_{i \in \mathcal{R}_r} a_{ri}\right)} \sum_{j \in \mathcal{R}_r} \sum_{i=1}^{n} a_{rj} \ell_{ji} (\log z_i - \mathbb{E}_r[\log z_i]), \tag{2.14}$$

where $\log C^* = \sum_{i=1}^{n} \lambda_i^{\text{ss}} \log z_i - \frac{1}{1+1/\eta} \log \chi$ *is the log output in the absence of frictions.*

(b) Additionally, if labor supply is fully elastic ($\eta \to \infty$), then

$$\log C = \log C^* - \lambda_r^{\text{ss}} \sum_{j \in \mathcal{R}_r} a_{rj} \mathbb{K}_r \left(- \sum_{i=1}^{n} \ell_{ji} \log z_i \right), \tag{2.15}$$

where $\mathbb{K}_r[x] = \log \mathbb{E}_r[e^x] - x$.

Proposition 3 characterizes the impact of productivity shocks on aggregate output when firms in industry r face informational frictions in setting up their production processes. Statement (a) characterizes aggregate output in terms of model primitives to a first-order approximation, whereas statement (b) provides an exact characterization for the special case that labor supply is fully elastic. It is immediate to see that the last terms on the right-hand sides of (2.14) and (2.15) vanish when firms in industry

r have complete information about the realizations of productivity shocks—and hence make all their price and quantity decisions with full anticipation of all shocks. In such a case, the expressions in (2.14) and (2.15) reduce to the standard result in the literature, according to which the impact of a productivity shock to industry i on aggregate output is equal to its pre-shock Domar weight, λ_i^{ss}. However, (2.14) and (2.15) also show that such a simple relationship no longer holds if firms in industry r face uncertainty about the productivity of their upstream supply chains. Specifically, the result in Proposition 3 leads to the following observations.

First, the expression in (2.15) illustrates that an increase in the uncertainty faced by firms in the industry with rigid inputs translates into unambiguously lower (expected) aggregate output. In particular, taking unconditional expectations from both sides of (2.15) and using the observation that $\mathbb{E}[\log \mathbb{E}_r[e^x] - x] > 0$ for any non-degenerate random variable x implies that $\mathbb{E}[\log C] < \mathbb{E}[\log C^*]$. This is, of course, intuitive: the fact that firms in industry r make (some or all of) their intermediate input decisions under incomplete information about productivity shocks induces them to rely more heavily on less volatile suppliers, even if this comes at the cost of forgoing more efficient ones. This role of uncertainty can be seen more clearly if one considers a second-order approximation to (2.15):

$$\mathbb{E}[\log C] = \mathbb{E}[\log C^*] - \frac{1}{2}\lambda_r^{\text{ss}} \sum_{j \in \mathcal{R}_r} a_{rj} \mathbb{E}\left[\text{var}_r\left(\sum_{i=1}^{n} \ell_{ji} \log z_i\right)\right], \tag{2.16}$$

which captures how higher uncertainty—as measured by the conditional variance of productivities—reduces expected aggregate output.

Second, the last term on the right-hand side of (2.15) (or its second-order approximation in (2.16)) indicates that the negative impact of informational frictions on aggregate output is increasing in industry r's Domar weight, as well as in r's uncertainty about the supply chains of each of its rigid intermediate inputs separately, $\text{var}_r\left(\sum_{i=1}^{n} \ell_{ji} \log z_i\right)$.

Not surprisingly, this uncertainty is weighted by the importance of industry j in r's production technology (as captured by the expenditure share a_{rj}). This means that, all else equal, a rigid industry r functions as a tighter "production bottleneck" for the entire economy if it is simultaneously (i) a larger supplier in the economy (as captured by the larger Domar weight, λ_r^{ss}) and (ii) a more important direct or indirect customer of other firms in the economy (proxied by greater a_{rj} and ℓ_{ji}).

Third, one can use Proposition 3 to obtain an expression for how microeconomic shocks translate into macroeconomic outcomes. Under the assumption that industry-level productivity shocks are independent, equation (2.14) implies that the slope coefficient of regression

$$\log C = \gamma_0 + \gamma_i \log z_i + \varepsilon_i \tag{2.17}$$

—which captures the (average) first-order impact of shocks to industry i on aggregate output—is given by

$$\gamma_i = \lambda_i^{\text{ss}} - \left(\frac{\lambda_r^{\text{ss}}}{1 + 1/\eta \left(1 - \lambda_r^{\text{ss}} \sum_{i \in \mathcal{R}_r} a_{ri}\right)} \sum_{j \in \mathcal{R}_r} a_{rj} \ell_{ji} \right) \kappa_{ri}, \tag{2.18}$$

where $\kappa_{ri} = \mathbb{E}[\text{var}_r(\log z_i)] / \text{var}(\log z_i)$ parameterizes the uncertainty of firms in industry r about shocks to industry i (as defined in (2.3)).[11] Equation (2.18) establishes that, in the presence of real rigidities and informational frictions, (i) the first-order impact of productivity shocks is dampened compared to the predictions of Hulten's theorem for the fully flexible benchmark, for which the first-order impact of a shock to industry i is equal to λ_i^{ss} irrespective of the value of η; (ii) this dampening effect is stronger the more uncertain firms in industry r are about the shock's realization and the more elastic the labor supply is; and (iii) the extent to which shocks to industry i shape aggregate

[11]Regression coefficient γ_i serves as the counterpart to $\mathrm{d}\log C/\mathrm{d}\log z_i$ in our incomplete-information economy.

outcomes depends on the size of the rigid industry r and the extent to which r is exposed (directly or directly) to those shocks (as captured by $\sum_{j\in\mathcal{R}_r} a_{rj}\ell_{ji}$).

Finally, it is worth pointing out that whereas aggregate labor supply in the frictionless economy is independent of the shocks (because of the household's logarithmic utility), that is no longer the case in the presence of informational frictions. This is because of the fact that rigid input choices of firms in industry r are not indexed to the realization of shocks, and yet, all good markets have to clear irrespective of the shocks' realizations. The only way this can happen is for aggregate labor supply to adjust in response to the shocks. For example, aggregate labor supply has to rise if realized shocks turn out to be smaller than the firm's expectations.

We next turn to the implications of demand shocks for output and inflation. To ensure that such shocks have a non-trivial impact on real variables, we focus on the case in which the downward nominal wage rigidity constraint bind. We have the following counterpart to Proposition 3.

Proposition 4. *Suppose r is the only rigid industry and Assumption 1 is satisfied. If the downward nominal wage rigidity constraint binds, and in the absence of productivity shocks,*

$$\log C = \log m - \log \bar{w} - \left(\lambda_r^{\text{ss}} \sum_{j\in\mathcal{R}_r} a_{rj}\right) \mathbb{K}_r\left(-\log m\right) \tag{2.19}$$

$$\log P = \log \bar{w} + \left(\lambda_r^{\text{ss}} \sum_{j\in\mathcal{R}_r} a_{rj}\right) \mathbb{K}_r\left(-\log m\right), \tag{2.20}$$

where $\mathbb{K}_r[x] = \log \mathbb{E}_r[e^x] - x$.

This result shows that the real rigidities and firms' uncertainty about the realization of demand shocks reduce aggregate output, while increasing the price level (which serves as a proxy for inflation in our model). For example, it is immediate to see that the expression for $\log C$ in (2.19) can be approximated by $\log C = \log m - \log \bar{w} - (\log m - \mathbb{E}_r[\log m] + \frac{1}{2}\operatorname{var}_r(\log m))\lambda_r^{\text{ss}} \sum_{j\in\mathcal{R}_r} a_{rj}$ to a second order, thus indicating that an increase in $\operatorname{var}_r(\log m)$ would reduce aggregate output.

As with supply shocks, one can use Proposition 4 to characterize the (average) first-order impact of demand shocks on output and inflation by calculating the slope coefficients of the following regressions:

(2.21) $$\log C = \gamma_0 + \gamma_m \log m + \varepsilon_m$$

(2.22) $$\log P = \delta_0 + \phi_m \log m + \varepsilon_m.$$

Using equations (2.19) and (2.20), it is easy to verify that

(2.23) $$\gamma_m = 1 - \phi_m \qquad \text{and} \qquad \phi_m = \left(\lambda_r^{\text{ss}} \sum_{j\in\mathcal{R}_r} a_{rj}\right)\mu_r,$$

where $\mu_r = \mathbb{E}[\operatorname{var}_r(\log m)]/\operatorname{var}(\log m)$ parameterizes the uncertainty of firms in industry r about the realization of $\log m$, as defined in (2.4). Given that $\phi_m > 0$ whenever $\mu_r(\log m) > 0$, it follows immediately that real rigidities and informational frictions dampen the real effect of positive aggregate demand shocks, while increasing their inflationary effects. Not surprisingly then, both of these effects are determined by the size of the rigid industry, λ_r^{ss}, and the extent to which it relies on rigid intermediate inputs, $\sum_{j\in\mathcal{R}_r} a_{rj}$.

It is also instructive to compare the coefficients γ_m and ϕ_m in (2.23) with the results of Baqaee and Farhi (2022), who find that, in a Cobb-Douglas economy with a single factor of production subject to nominal wage rigidity, aggregate demand shocks translate one-for-one to aggregate output, with no impact on inflation. They also show that, holding sectoral Domar weights constant, the details of the economy's production network structure are irrelevant for how aggregate demand shocks impact aggregate output. As is evident from (2.23), neither statement is no longer true in the presence of informational frictions: the real effect of aggregate demand shocks is dampened, a positive demand shock leads to an increase in the price level, and both effects depend on the production network structure.

2.4.2. Public Information

By focusing on an economy with a single rigid industry, the results in Subsection 2.4.1 abstract from the possibility that firms may be subject to multiple rigid suppliers and customers. In this subsection, we apply the general result in Proposition 1 to an economy in which firms in all industries have access to the same public information about the shocks (that is, $\omega_i = \omega$ for all industries i). Focusing on such an economy allows us to explore how the impact of real rigidities can build up over production chains.

To express our results in this more general case, it is convenient to define the following objects. Let $\mathbf{A}_f$ denote the matrix whose (i, j) element is equal to the corresponding element of matrix $\mathbf{A}$ if $j \in \mathcal{F}_i$ and is equal to zero otherwise. Therefore, $\mathbf{A}_f$ captures input-output relationships that are flexible and are not subject to quantity adjustment frictions. Similarly, we define matrix $\mathbf{A}_r = \mathbf{A} - \mathbf{A}_f$ to capture input-output relationships that are subject to frictions. Finally, let $\mathbf{L}_f = (\mathbf{I} - \mathbf{A}_f)^{-1}$ denote the Leontief inverse corresponding to the flexible inputs in the economy. We have the following result.

Proposition 5. *Suppose all firms share a common information set and the downward nominal wage rigidity constraint does not bind. If labor supply is fully elastic, then*

$$\mathbb{E}[\log C] = \mathbb{E}[\log C^*] - \frac{1}{2}\lambda^{\text{ss}\prime}\mathbf{A}_r \operatorname{diag}(\mathbf{Q}) + \frac{1}{2}\lambda^{\text{ss}\prime} \operatorname{diag}(\mathbf{A}_r\mathbf{1}) \operatorname{diag}(\mathbf{H}'\mathbf{Q}\mathbf{H}), \tag{2.24}$$

to a second-order approximation, where $\mathbf{H} = \mathrm{diag}(\mathbf{A}_r'\lambda^{\mathrm{ss}})\mathbf{L}_f\mathrm{diag}^{-1}(\lambda^{\mathrm{ss}})$*,*

$$\mathbf{Q} = \sum_{k=0}^{\infty} \left(\mathbf{L}_f \,\mathrm{diag}(\mathbf{A}_r\mathbf{1})\mathbf{H}'\right)^k \mathbb{E}[\mathrm{var}_\omega(\mathbf{L}_f \log z)] \left(\mathbf{H}\,\mathrm{diag}(\mathbf{A}_r\mathbf{1})\mathbf{L}_f'\right)^k,$$

and $\mathrm{var}_\omega(\cdot)$ *denotes the conditional variance-covariance matrix with respect to the firms' common information set.*

This result generalizes Proposition 3 by allowing for multiple rigid industries and an arbitrary pattern of rigid and flexible inputs, while also relaxing Assumption 1. In fact, it is easy to verify that if there is a single industry with rigid inputs r and if Assumption 1 is satisfied, then $\mathbf{H}\,\mathrm{diag}(\mathbf{A}_r\mathbf{1}) = 0$, in which case $\mathbf{Q} = \mathbb{E}[\mathrm{var}_\omega(\mathbf{L}_f \log z)]$ and as a result, (2.24) reduces to the expression in (2.16). Given its generality, we will use the characterization in Proposition 5 as the basis of our quantification exercise in Section 2.5.

While significantly more involved than Proposition 3, the above result captures the same economic force, according to which an increase in firms' uncertainty about the shocks reduces aggregate output. The intuition is also similar: the fact that intermediate input decisions are sunk by the time firms observe the realized productivities means that they shift their demand from more uncertain suppliers to more reliable ones, even if it comes at the cost of lower (expected) productivity.

Next, turning to demand shocks in this setting, we can establish the following result:

Proposition 6. *Suppose all firms share a common information set, the downward nominal wage rigidity constraint binds, and labor supply is fully elastic. Then, in the* *absence of productivity shocks,*

(2.25)

$$\mathbb{E}[\log C] = \mathbb{E}[\log m] - \log \bar{w} - \frac{1}{2}\lambda^{\mathrm{ss}\prime}\left(\mathbf{A}_r \operatorname{diag}(\mathbf{G}) - \operatorname{diag}(\mathbf{A}_r\mathbf{1}) \operatorname{diag}(\mathbf{H}'\mathbf{G}\mathbf{H})\right) \mathbb{E}[\operatorname{var}_\omega(\log m)]$$

(2.26)

$$\mathbb{E}[\log P] = \log \bar{w} + \frac{1}{2}\lambda^{\mathrm{ss}\prime}\left(\mathbf{A}_r \operatorname{diag}(\mathbf{G}) - \operatorname{diag}(\mathbf{A}_r\mathbf{1}) \operatorname{diag}(\mathbf{H}'\mathbf{G}\mathbf{H})\right) \mathbb{E}[\operatorname{var}_\omega(\log m)]$$

to a second-order approximation, where

$$\mathbf{G} = \sum_{k=0}^{\infty} (\mathbf{L}_f \operatorname{diag}(\mathbf{A}_r\mathbf{1})\mathbf{H}')^k \mathbf{L}_f \alpha\alpha' \mathbf{L}_f' (\mathbf{H} \operatorname{diag}(\mathbf{A}_r\mathbf{1})\mathbf{L}_f')^k.$$

2.4.3. Dispersed Information and Higher-Order Expectations

While the information structure in Subsection 2.4.2 exhibits incomplete information throughout the economy, it abstracts from the possibility that information may be dispersed among different firms, as all signals are assumed to be publicly observable. This allows for a significant degree of coordination between different firms in the economy: firms can use these public signals to coordinate not only with their direct suppliers and customers, but also with potentially distant firms on their supply chains. We therefore conclude this section by studying the implications of heterogeneity in the firms' information sets. In order to keep the analysis tractable, we focus on a simple production network structure.

Consider the economy depicted in Figure 2.1 consisting of three industries organized on a vertical production chain, where industry 3 is the sole input supplier to industry 2 (with expenditure share $a_{23} = a_2$) and industry 2 is the sole input supplier to industry 1 (with expenditure share $a_{12} = a_1$). Furthermore, assume that industry 3 only uses labor for production ($\alpha_3 = 1$) and that industry 1 is the only industry that sells to the households ($\beta_1 = 1$). We have the following result:

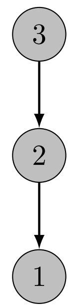

Figure 2.1. Vertical Production Network

Note: Each vertex corresponds to an industry, with a directed edge present from one vertex to another if the former is an input-supplier to the latter. In addition to intermediate inputs, all firms use labor as an input for production (not depicted in the figure).

Proposition 7. *Suppose labor supply is fully elastic and that* $m < \bar{w}/\chi$*. In the absence of productivity shocks,*

$$\log C = \log m - \log \bar{w} - a_1(1-a_2)\sum_{s=0}^{\infty} a_2^s\Big(\log m - (\mathbb{E}_1\mathbb{E}_2)^s\mathbb{E}_1[\log m]\Big) \tag{2.27}$$

$$\log P = \log \bar{w} + a_1(1-a_2)\sum_{s=0}^{\infty} a_2^s\Big(\log m - (\mathbb{E}_1\mathbb{E}_2)^s\mathbb{E}_1[\log m]\Big) \tag{2.28}$$

to a first-order approximation, where $(\mathbb{E}_1\mathbb{E}_2)^{s+1}[\cdot] = \mathbb{E}_1\mathbb{E}_2[(\mathbb{E}_1\mathbb{E}_2)^s[\cdot]]$.

Proposition 7 illustrates that, when information is dispersed, the impact of aggregate demand shocks on aggregate output and the price level depends on not just firms' first-order expectations but also on all expectations of higher order. As is well-known, higher-order beliefs adjust more sluggishly than first-order beliefs (Angeletos and Huo, 2021). Therefore, equations (2.27) and (2.28) underscore how rigidities build up over the production chain, dampening the real effects of positive monetary shocks, while amplifying their inflationary effects.

That firms' higher-order expectations can matter for aggregate economic outcomes is not, in and of itself, novel, and it is line with prior work such as Golub and Morris

(2018), Angeletos and Lian (2016, 2018), and La'O and Tahbaz-Salehi (2022), among others. What distinguishes the expressions in Proposition 7 from the results in prior work is how these higher-order expectations matter for aggregate output and price level. In particular, even though goods flow only in one direction in this economy—from upstream suppliers to their downstream customers—the macroeconomic variables in (2.27) and (2.28) depend on the iterated expectations in both directions: the supplier's expectations of the customer's expectation and vice versa. This is a consequence of the fact that, when choosing their input quantities, firms need to form forecasts not only about their suppliers' input prices, but also about demand from their customers.[12] But of course the same logic applies to those customers and suppliers as well, resulting in an infinite regress of expectations between firms in industries 1 and 2.

To further explore the implications of dispersed information in the vertical economy in Figure 2.1 and the expressions in Propositions 7, it is instructive to focus on a parametric information structure with normally distributed shocks and signals. Specifically, suppose all firms in the economy have a common prior about the aggregate demand shock: $\log m \sim N(0,1)$. Additionally, suppose that firms in industry $i \in \{1,2\}$ receive a public and a private signal $\omega_i = (\tilde{s}, s_i)$, where

$$\begin{aligned} \tilde{s} &= \log m + \tilde{\epsilon} \qquad && \tilde{\epsilon} \sim N(0, \sigma^2/(1-\delta)) \\ s_i &= \log m + \epsilon_i \qquad && \epsilon_i \sim N(0, \sigma^2/\delta), \end{aligned} \tag{2.29}$$

noise terms $(\epsilon_1, \epsilon_2, \tilde{\epsilon})$ are independent, and $\delta \in (0,1)$ and $\sigma^2 > 0$ are parameters. Note that the parametrization in (2.29) implies that firms' uncertainty about $\log m$ is the same irrespective of the value of δ. In particular, $\text{var}_i(\log m) = \sigma^2/(1+\sigma^2)$ for all $\delta \in (0,1)$. This means that δ parameterizes the strength of the private signal vis-à-vis the public signal:

[12] Importantly, as one can see from equations (2.10) and (2.11), the suppliers' and customers' expectations do not appear symmetrically.

as $\delta \to 1$ all information is private, whereas at the other end of the spectrum as $\delta \to 0$ all information is public.

Corollary 1. *Under the information structure in* (2.29)*, the first-order effects of an aggregate demand shock on aggregate output and the price level are given by*

$$\gamma_m = 1 - \phi_m \qquad \textit{and} \qquad \phi_m = a_1\sigma^2 \frac{(1+\sigma^2) + a_2\delta\sigma^2}{(1+\sigma^2)^2 - a_2\delta^2\sigma^4},$$

where γ_m *and* ϕ_m *are the coefficients of regressions* (2.21) *and* (2.22)*, respectively.*

Corollary 1 illustrates how δ, which parameterizes the extent of information dispersion, shapes the shock's aggregate impact. Specifically, it is easy to verify that a more dispersed information (i.e., a larger δ) results in greater inflation in response to positive aggregate demand shocks. This reflects the fact that with dispersed information, it is harder for firms in the supply chain to coordinate their production decisions with another. This diminishes the real effect of positive aggregate demand shocks and implies that such shocks manifest themselves as higher inflation.

2.5. Quantitative Analysis

In this section, we use our theoretical results to quantify the effect of informational frictions and real rigidities on aggregate output in a calibrated version of the model. We calibrate our model at a quarterly frequency. This amounts to assuming that firms are unable to adjust their rigid inputs within a quarter after the realization of the shocks. For this quantitative analysis, we also ignore the role of downward nominal wage rigidities by setting the lower bound on the nominal wage to $\bar{w} = 0$.

As is typical in the literature on production networks, we calibrate the model to the U.S. data using input-output tables constructed by the Bureau of Economic Analysis (BEA). These tables provide intermediate input expenditures by various industries, as well as each industry's contribution to final uses. However, to calibrate the model, we also need to specify (i) the sets of rigid and flexible inputs for each of the industries in the sample, as well as (ii) each industry's corresponding information set. We therefore start by discussing how we specify each of these.

Rigid and Flexible Inputs. To designate the sets of flexible and rigid inputs, we distinguish between intermediate inputs and investment goods each industry purchases from other industries. As mentioned in the Introduction, whereas the average lead time for obtaining production materials for manufacturing firms in the United States in January 2023 was 87 days, the average lead time for acquiring capital inputs—such as machinery, plant equipment, software and the like—was roughly twice as large (166 days). Similarly, while only 6% of firms faced lead times of over one year for their intermediate inputs, the same number was 23% for capital inputs (Institute for Supply Management, 2023a). Given this disparity in lead times, for the purpose of the calibration, we designate investment goods used by each industry as that industry's rigid inputs, while treating intermediates as flexible.

To calibrate the model according to the above criterion, we rely on the "investment network" constructed by vom Lehn and Winberry (2022). Focusing on a 37-sector disaggregation of the entire private nonfarm economy, vom Lehn and Winberry (2022) construct an annual network of flows that measures the share of the total investment expenditure of a given sector i that is purchased from another sector j for each pair of sectors (i, j) in the economy. We treat these shares as the corresponding expenditure shares on rigid inputs in our model, while setting expenditure shares on intermediates equal to the expenditure shares on flexible inputs. To be more specific, we consider the following static variant of Horvath's (2000) model, where each industry i produces inputs according to the following constant returns production technology:

$$y_i = z_i \zeta_i l_i^{\alpha_i} k_i^{\rho_i} \prod_{j=1}^{n} x_{ij}^{\psi_{ij}} \qquad , \qquad \alpha_i + \rho_i + \sum_{j=1}^{n} \psi_{ij} = 1, \tag{2.30}$$

where ζ_i is a normalization constant, x_{ij} is the intermediate input purchased from industry j, ψ_{ij} captures the expenditure share on intermediate input j, ρ_i is the share of capital, and k_i is the capital input, which itself is produced from other industries' output:

$$k_i = \prod_{j=1}^{n} g_{ij}^{\gamma_{ij}} \qquad , \qquad \sum_{j=1}^{n} \gamma_{ij} = 1, \tag{2.31}$$

where g_{ij} is the amount of industry j's output used by industry i as an input into i's capital bundle. The market-clearing condition for the good produced by industry i is thus given by

$$y_i = c_i + \sum_{j=1}^{n} x_{ji} + \sum_{j=1}^{n} g_{ji}, \tag{2.32}$$

thus accounting for the fact that the output of industry i can either be consumed by the households, used as an intermediate input by other industries, or serve as an input in other industries' capital bundle.[13] Therefore, in this model—and unlike the model in Section 2.2—the output of industry j takes a dual role in the production technology of industry i: once as an intermediate input and once as an input in i's investment bundle. Nonetheless, the model in equations (2.30)–(2.32) can be mapped to the model in Section 2.2 in a straightforward manner. By designating the capital bundle k_i as a separate industry whose inputs are all rigid and using the output of this industry as a (flexible) input in the production technology of industry i (with share ρ_i), equations

[13] This model coincides with the model in vom Lehn and Winberry (2022) when the depreciation rate of capital in their model is set equal to 100%, in which case the stock of capital in each industry becomes equal to the investment good.

(2.30)–(2.32) reduce to the model in Section 2.2, with the following (expanded) input-output matrix:

$$\mathbf{A} = \begin{bmatrix} \mathbf{\Psi} & \mathrm{diag}(\rho) \\ \\ \mathbf{\Gamma} & \mathbf{0} \end{bmatrix},$$

where $\mathbf{\Psi} = [\psi_{ij}]$ and $\mathbf{\Gamma} = [\gamma_{ij}]$ denote the matrices corresponding to the intermediate input and investment networks in equations (2.30) and (2.31), respectively, and $\mathrm{diag}(\rho)$ is a diagonal matrix with entries equal to the capital shares in each industry. We can thus calibrate matrix $\mathbf{A}$ using the network data of vom Lehn and Winberry (2022).

Information Sets. Next, we turn to specifying the information sets. To this end, we construct the variance-covariance matrix of $(\log z_1, \ldots, \log z_n)$ by first detrending the TFP process for each industry and then setting the variance-covariance matrix of the log-productivity shocks equal to the empirical variance-covariance matrix of the detrended processes. This amounts to assuming that while firms observe all past productivity shocks and are aware of their corresponding trends, they do not observe productivity innovations at the beginning of each quarter. Finally, note that this specification of information structure means that all firms share a common information set (or equivalently all signals are public).

Quantitative Analysis. With the economy's information structure and the sets of flexible and rigid inputs specified, we now turn to the quantitative assessment of our model's implications.

As a first exercise, we quantify the role of informational and quantity adjustment frictions for aggregate output. Recall that we assume that (i) all firms share the same common information structure, (ii) all intermediate inputs purchased from other industries are fully flexible, and (iii) all inputs purchased to construct the capital bundle are rigid. We can therefore use equation (2.24) in Proposition 5 to measure the drop in expected output due to the presence of frictions, where the matrices corresponding to the flexible and rigid inputs are given by

$$\mathbf{A}_f = \begin{bmatrix} \mathbf{\Psi} & \mathrm{diag}(\rho) \\ \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \qquad \text{and} \qquad \mathbf{A}_r = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \\ \mathbf{\Gamma} & \mathbf{0} \end{bmatrix},$$

respectively.

Using the calibrated model to calculate the expression on the right-hand side of (2.24), we find that informational frictions and real rigidities result in roughly a 1% drop in aggregate output (measured as a percentage of steady-state consumption). This means that, according to our calibration, the interaction of the two frictions with the economy's production network can generate significant macroeconomic effects.

To get a more granular picture of the sources of this output loss, in a second exercise, we calculate the expected drop in aggregate output while assuming that only capital inputs purchased from a single industry r are subject to informational and quantity adjustment frictions. Specifically, we assume that whereas firms in any industry i need to commit up front to the quantity g_{ir} used for the production of their capital input, they can flexibly adjust all other capital and intermediate inputs—including the quantity of intermediate inputs x_{ir} purchased from industry r—in response to the realized productivity shocks. To calculate the corresponding drop in aggregate output, we once again use equation (2.24), but this time specifying matrices $\mathbf{A}_f$ and $\mathbf{A}_r$ to reflect the fact that only capital inputs from a single industry are rigid. Repeating this exercise for each of the 37 industries in our sample allows us to rank industries based on their role as "supplier bottlenecks" in the economy's production network.

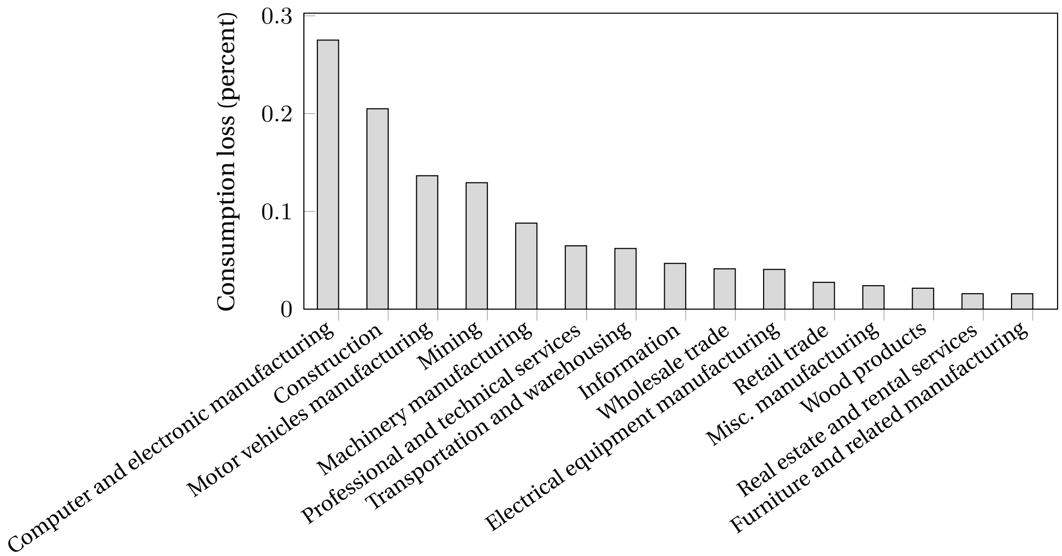

Figure 2.2. Supplier Bottlenecks

Notes: The figure plots output loss as a percentage of steady-state consumption when capital inputs purchased from a single industry are subject to frictions, while all other intermediate and capital inputs can be adjusted flexibly. The figure only reports for the 15 industries with the largest corresponding output loss.

Figure 2.2 reports the results for the 15 industries that result in the largest drop in expected output. At the top of the list is 'Computer and electronic manufacturing,' which generates a drop in GDP equal to 0.05% of steady-state output, followed by 'Construction' and 'Motor vehicles manufacturing'. Notably the resulting output loss diminishes rapidly, indicating that a large majority of the 37 industries do not play a significant role as supplier bottlenecks. It is also worth pointing out that four out of top six industries on this list—construction, machinery manufacturing, motor vehicles manufacturing, and professional/technical services—coincide with the industries vom Lehn and Winberry (2022) identify as "investment hubs" that are responsible for producing nearly 70% of total investment. This is to be anticipated: since these capital inputs are used widely by many customer industries, any friction in adjusting those input quantities in response to shocks would result in more significant aggregate effects.

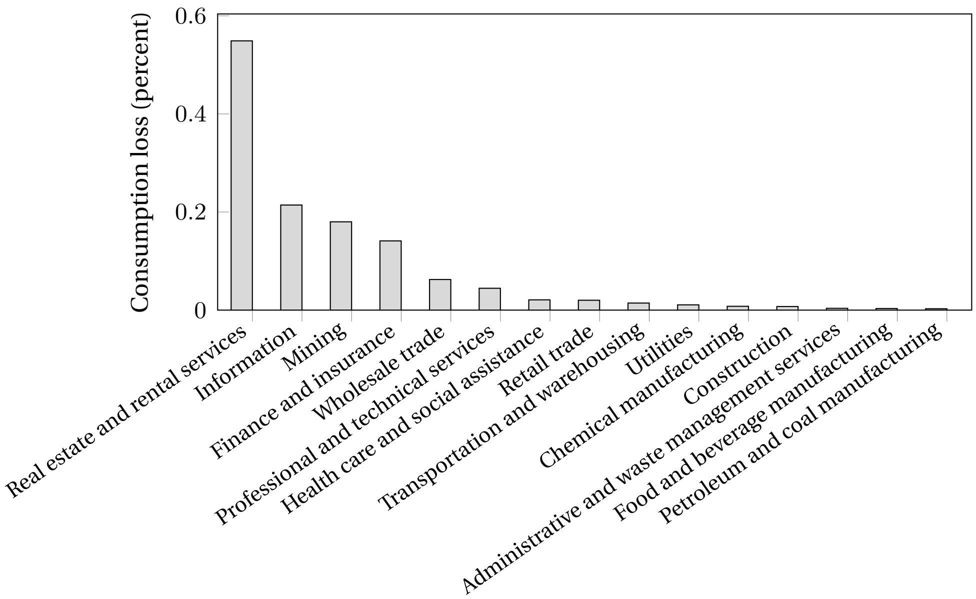

Figure 2.3. Customer Bottlenecks

Notes: The figure plots output loss as a percentage of steady-state consumption when capital inputs purchased by a single industry are subject to frictions, while all other intermediate and capital inputs can be adjusted flexibly. The figure only reports for the 15 industries with the largest corresponding output loss.

As our final exercise, we repeat the previous analysis but this time assuming that only a single industry r faces frictions when acquiring all its capital inputs, while all other industries $i \neq r$ face no frictions in acquiring either their intermediate or capital inputs. This exercise, which is in the spirit of the results in Subsection 2.4.1, allows us to rank different industries based on their roles as "customer bottlenecks:" what fraction of output is lost because any given industry cannot adjust its investment bundle in response to shocks. Figure 2.3 reports the results for the 15 industries that result in the largest drop in expected output. The picture that emerges is considerably different from the one in Figure 2.2. This time industries with larger Domar weights—and not necessarily larger shares in the investment network—tend to appear towards the top of the list.

2.6. Conclusions

In this paper, we develop a production network model, in which firms are subject to informational frictions and real rigidities. The presence of such frictions means that (i) firms may be restricted in how effectively they can adjust their intermediate input quantities in response to changes in the economic environment and (ii) firms may have to choose their quantities only with incomplete information about the realizations of shocks. Our main theoretical result provides an implicit characterization of equilibrium nominal prices and Domar weights in terms of model primitives, namely, the economy's production network structure, the firms' information sets, and the set of rigid intermediate inputs. While only implicit, this result illustrates that equilibrium prices and quantities are determined by the firms' expectations of their upstream input prices as well as their expectations' of their downstream demand. We then consider various special cases of this economy to obtain closed-form solutions for how supply and demand shocks impact aggregate output and inflation.

A few insights emerge from the model. First, the presence of the real rigidities and informational frictions results in an unambiguous drop in aggregate output, as firms decide to shift demand from more efficient suppliers towards those that are less volatile. Second, these frictions in turn reduce the passthrough of productivity shocks to aggregate output and dampen the real aggregate effect of positive demand shock, while at the same time, increasing their inflationary effect compared to the frictionless benchmark.

While aimed at incorporating two realistic frictions into an otherwise standard model of production networks, the model developed in this paper is nonetheless still very stylized. First, one implicit assumption in the model is that firms face the same degree of rigidity irrespective of whether they decide to increase various quantities or decrease them. In reality, such frictions are most likely asymmetric, as firms can more easily reduce input and output quantities than increase them. Second, we abstracted from adjustment costs by assuming that once the intermediate input quantity decisions are made, they are completely sunk and cannot be changed. But it is easy to imagine various scenarios in which firms can adjust their production processes (even if only partially) by incurring some adjustment costs. Third, our static model ignores the role of inventory management as one of the key tools available to firms for responding to unanticipated supply and demand shocks.[14] We leave exploring the implications of these realistic features for future research.

[14] See Ferrari (2023) for a tractable model of production networks with inventories.

CHAPTER 3

Production without FIRE

[0] I am grateful to Alireza Tahbaz-Salehi, George-Marios Angeletos, Matt Rognlie, Uri Wilensky, Laura Murphy, Michael Cai, Giovanni Sciacovelli and Martin Eichenbaum for comments and suggestions. I am also grateful to seminar participants at WEHIA 2021. This research was supported in part through the computational resources and staff contributions provided for the Quest high performance computing facility at Northwestern University which is jointly supported by the Office of the Provost, the Office for Research, and Northwestern University Information Technology.

3.1. Introduction

Firms take time to produce goods. Only 8% of firms are "hand-to-mouth" and risk ordering intermediates less than 30 days before using them in production. One reason is that average lead time for intermediates is roughly 3 months, with more than 15% of firms ordering at least 6 months in advance to accommodate long lead times (Institute for Supply Management, 2023b). As a result, firms have to make some of their production decisions in advance and the firm problem is inherently dynamic. In practice, it often means hiring workers, setting production targets and posting prices long before goods are sold and profits are hypothetically generated.

If firms had perfect foresight about the economy in all future periods, these production frictions would be harmless. Firms would simply adjust production in advance given their future demand curve to maximize their future profits. Under incomplete information, the situation is murkier. If a firm starts producing today, what will be the demand when goods are finally produced and available? What will be the prevailing price then? Firms face idiosyncratic risk at the micro level and this risk is existential because it involves spending upfront large amount of money on intermediates, labor, capital and overhead before revenues are generated. Indeed, post-production marketing costs only represent 10% of firms revenues on average (Gartner, Inc., 2022). Firms frontload large and certain payments and backload uncertain revenues, as long as production of the good they sell is not instantaneous.

A large literature on granularity has emphasized how micro risks like demand uncertainty can aggregate to macroeconomic fluctuations. More specifically (Pellet and Tahbaz-Salehi, 2023; Bui et al., 2022; Kopytov et al., 2022b) have stressed the importance of production rigidities (due to inputs, labor or technology) associated with informational frictions over the production networks to explain supply bottlenecks and inflation surprises. To account for the decentralized nature of supply chains, a key innovation was to introduce heterogenous expectations and imperfect information sharing between industries. Indeed, supply chains remain imperfectly integrated and cannot be well approximated as a representative and centralized multi-sector firm.

However this literature still endows firms with a large amount of information about the structure of the economy and the nature of the risk they face in the production process. These papers only deviate slightly from the benchmark Full-Information Rational Expectation (FIRE) framework. First, firms face a kind of uncertainty that is relatively limited. They know the statistical distribution and the structure of the model they operate in. In that sense, uncertainty is fully quantifiable and actionable. This is the realm of known unknowns with known probabilities and state space, as opposed to a world of unknown unknowns or "Knightian uncertainty" (Knight, 1921). Second, by the nature of the rational expectation assumption, firms know the structural equations of the model and reason through these equations. For example, monopolistic competitors know the demand curve equations and use them in their maximization problem. This fact is also common knowledge across firms, so that they can clearly reason through the reasoning and behaviors of other firms as well. This gives rise to model structures that often resemble that of well-defined strategic games. It is then natural to study game theoretical solution concepts like Nash equilibrium and their predictions. This posits a high degree of information sharing and processing across firms, especially so when the problem they face is dynamic and spans months if not years. These assumptions are often made for tractability reason, and it is true that they considerably simplify model derivations and allow for analytical solutions otherwise unaccessible.

What if firms do not have access to that structural knowledge about the model itself? When firms face this second and more radical type of uncertainty, many building blocks of traditional FIRE models fall apart. Because firms cannot reason through structural equations, profit maximization alone does not suffice to guide action, since every action leads to an unknown distribution of potential outcomes (Alchian, 1950). Firms cannot reason through other agents' behavior either because they lack the structural understanding that would allow them to predict others' behavior. In this fundamentally uncertain environment expectations are hence necessarily mis-specified. Acknowledging their limitations and using second best reasoning, Firms use the only information they have access to: past observations generated by the model. That is why the non-FIRE firm is necessarily a statistician, formulating expectations conditional on past observations alone. This also implies that agents' policy rules cannot be a function of the model itself. Using the terminology somewhat loosely, this is akin to a measurability constraint. The same way that firms cannot make decisions based on exogenous shocks they have not observed yet, they cannot make decisions based on a model they do not know.

There is strong evidence that this latter approach is closer to how the economy actually behaves, and how firms formulate expectations. Empirical evidence strongly suggests persistence of price dispersion for homogeneous goods (Federal Reserve Bank of Richmond and Trachter, 2017), failure of the law of one price, frequent backlogs (Institute for Supply Management, 2023b) or rationing. Involuntary inventories and unemployment are the norm rather than the exception. Yes, these facts could be captured by adding multiple frictions to a FIRE framework[1]. It would come at the cost of tractability and high computational complexity, which was the reason for making these assumptions in the first place. The fact is that FIRE-like models often have the structure of Brouwer fixed point problems, which are hard computational problems for modelers and firms alike (Roughgarden, 2010). Adding non-linear frictions to match certain characteristics of the data only makes the problem harder computationally. And

[1] Search frictions (Burdett and Judd, 1983), exogenous price or wage rigidities are often used to model these deviations

yet this hard computational problem the modeler is solving is precisely the problem firms should implicitly be solving under FIRE solution concepts.

This paper takes the view that these empirical facts are more than mere deviations from FIRE and should be modeled accordingly. Firms operate on decentralized markets surrounded by a fog of war for at least three reasons. First, they don't necessarily have the incentives to coordinate with competitors, and could be hoarding information strategically. Second, there is strong empirical evidence that sales growth and firm size both follow fat-tailed processes, making inference harder even under complete information (Bottazzi and Secchi, 2006). Third, there is also strong evidence that firms have bounded rationality and limited capacity to process available information effectively (Larsen-Hallock et al., 2022). Instead of using all public information and a well specified model to forecast sales, firms use simple linear statistical rules, i.e. mis-specified models, to extrapolate future sales from past observations. This makes it all the more difficult to define an optimal coordination strategy across firms. According to this view, the fact that the aggregate economy is relatively well-behaved is the anomaly to explain. A divine coincidence that should not be assumed ex ante by the modeler, but generated endogenously.

In these environments fast and frugal reasoning like simple inductive or adaptive rules become necessary and can paradoxically be more performant (Gigerenzer and Brighton, 2009). This is what this paper focuses on, with an emphasis on the consequences of frugal reasoning on macroeconomic outcomes. How sensitive is the aggregate economy to the type of strategies firms are using to predict demand at the micro level? Are business cycle characteristics a function of firms behavior and their ability to predict what demand will be? These are the main questions this paper is trying to answer. To capture the relevant dimensions detailed above, the model extended from Delli Gatti et al. (2011) will include quantity rigidities, incomplete information and bounded rationality on the firm side à la Arthur (1994), search frictions on decentralized non-Walrasian markets, price rigidities and financial frictions. Price rigidities arise naturally given the lack of coordination around the price vector. Firms need to post prices somewhat independently from each other before demand is realized. They effectively face a search problem chasing the customers to sell to. Similarly, firms hire workers for fixed-term contracts by posting vacancies and wages in advance. In that sense both labor and goods markets are decentralized search markets. Because firms need to produce before demand is realized, production cannot be internally financed and requires financing. And external financing begets financial constraints. As firms try to forecast demand the best they can, these frictions could interact in unexpected ways and lead to interesting emergent properties at the aggregate level.

Modeling all these frictions together in a standard equation based DSGE would be a daunting task for several reasons. By their nature, the traditional solution concepts assumes some sort of consistency at the aggregate level. The usual practice is to use a Nash equilibrium solution concept coordinated on prices for RBC models, or on labor hours for New Keynesian models with nominal rigidities. Existence of a flexible margin guarantees that the model "closes". This assumption would be inconsistent with the view that firms do not coordinate among each other. Heterogeneity is also particularly costly in these types of models, and the curse of dimensionality hits relatively soon, with very few state variables. This is particularly problematic when modeling firms with multiple rigidities as the number of state variables increases dramatically. For instance, decentralized labor markets with fixed term contracts make the history of posted wages firm-level state variable. Wages history will indeed determine the average cost of goods sold and therefore influence future production decisions.

To circumvent these difficulties, this paper uses a different computational approach called Agent-based Modeling (ABM). The principle is to start from the bottom-up, endowing agents with policy functions at the micro level, defining precise and consistent rules for transactions between agents and observing how this artificial economy behaves in the aggregate. In effect, it is taking the logic of the micro-foundation revolution to its natural limit. The macroeconomy is effectively the sum of all its agents, the same ways that humans are the sum of their cells. But because individual agents and cells interact in complex ways, the aggregated result is non-trivial and cannot be easily predicted from its parts. One agent is not representative of other agents, nor is the aggregate economy well approximated by a representative agent. Heterogeneity is a core element of this type of model. It can endogenously build up over time due to repeated interactions and path dependencies. That is because complex interactions can be an source of heterogeneity by themselves. One key advantage of this modeling approach is that the marginal cost of modeling additional heterogeneity is also much lower Because behaviors are simulated, rather than solved for under Nash equilibrium conditions.

The literature on Agent-based models is extensive and the interested reader can profitably read Axtell and Farmer (2023) for a recent and comprehensive account. What I want to emphasize here is the ability of this class of models to generate emergent behaviors at the aggregate level because aggregation is not trivial in complex systems. A canonical example is Schelling (1971)'s segregation model which illustrates how segregation can emerge in the aggregate against the preference of all agents involved. Similarly, the model presented here generates business cycle fluctuations in absence of aggregate shocks and despite the fact that agents have no taste for fluctuations. Business cycles simply emerge by the aggregation of accumulated granular errors of coordination in the system. Indeed, one of the great advantages of ABMs is to combine granularity with large scale interactions and realistic levels of heterogeneity to generate the type of firm size distributions observed in the real world (Axtell, 2006).

The main result of this paper is a sufficiency result. In the stylized economy presented here, firm forecasting behavior is critical in determining the structure of the business cycle. Inductive reasoning among firms breaks down the relationship between quantities and prices in the simulated economies. The correlation structure of the economy is also greatly affected, with the first principal component explaining a smaller share of the variance of the business cycle. Finally, the sales process is much more fat-tailed in presence of inductive firms. The paper is organized as follows: Section 3.2 reviews the literature. Section 3.3 describes the model. Section 3.4 discusses the calibration strategy. Section 3.5 presents some key results. Section 3.6 concludes.

3.2. Literature Review

This paper belongs to the literature on the origins and causes of business cycles, which studies the mechanisms that generate fluctuations in aggregate macroeconomic variables. The study of business cycles has a long history in macroeconomics and a proper account is beyond the scope of this paper. Contrary to the traditional literature on Real Business Cycles (Lucas and Prescott, 1971; Kydland and Prescott, 1982; Long and Plosser, 1983) where exogenous TFP shocks are the necessary disturbances generating output fluctuations, the model presented in this paper allows for fluctuations in absence of TFP shocks thanks to micro level mismatches between supply and demand and the finite number of firms and households.

From that perspective, this paper is similar in spirit to the literature exploring how micro-level fluctuations can generate aggregate fluctuations with granularity (Gabaix, 2011) and propagate through the network structure of the economy (Horvath, 2000; Acemoglu et al., 2012b; Carvalho, 2014; Carvalho et al., 2021; Baqaee and Farhi, 2019, 2022). In particular, it is closer to Pellet and Tahbaz-Salehi (2023) where firms cannot adjust quantities in response to unexpected demand fluctuations. The key departure from this paper is to allow for quantity, price and wage rigidities, and therefore disequilibrium on the labor and goods market. The flexible margins of adjustment to close the model are involuntary inventories and unemployment, while labor and goods rationing guarantee that demand never exceeds supply. The presence of non-Walrasian markets in the model and the random matching process between suppliers and customers or employers and employees will naturally introduce a random network structure in the economy. Prices and wages are not sufficient statistics in these markets to know who will get what. The conjonction of micro uncertainty, granularity and disequilibrium dynamics will propagate along the customer-supplier network and the worker-employer network, causing feedback mechanisms and the endogenous aggregate fluctuations at business cycle frequencies.

The endogeneity of business cycles in the model allows for a complex dynamic system with periodicity in steady state. The possibility of limit cycles echoes a recent literature Beaudry et al. (2015, 2020a,b) reviving old ideas about endogenous cycles dating back to Kalecki (1935); Kaldor (1940); Goodwin (1951); Guesnerie and Woodford (1993), suggesting that the interactions across agents could, *by themselves*, generate fluctuations in the economy. In that view, the economy naturally exhibits cyclicality and shocks to production and demand only affect the amplitude and frequency of these cycles.

This paper also relates to the learning and bounded rationality literature (Woodford, 2013; Townsend, 1983). Faced with endogenous uncertainty about demand and lacking a proper definition of the full problem, no unique model can be used to deduct the optimal solution. Many possible expectation models of demand are reasonable in face of ambiguity and heterogeneity in other firms' beliefs. Firms would need an inaccessible amount of information to know the true statistical process for demand. Absent a readily available "rational" expectation model of the world, firms can only learn from past observations using inductive reasoning. They will adapt their production plans according to a parametrized process developed in Arthur (1994).

This paper is also related to a long literature on inventories. Explaining inventory management as a result of intertemporal optimization with rational expectation (Blanchard, 1983) has long been a focus of study. Several motives for optimal holding of inventories have been proposed: the production-cost smoothing motives posit that firms use inventories to smooth out convex costs (Eichenbaum, 1989; Kryvtsov and Midrigan, 2013). The stockout avoidance motive suggests that firms mostly hold inventories under risk-aversion. Fixed delivery costs for intermediate inputs is an additionally source of inventory buildup that has been proposed (Khan and Thomas, 2007). The focus of this paper is the reverse, studying how unexpected inventory management affects firm decisions and ultimately output. Firms in the model are lean and would ideally not hold inventories. Because of bounded rationality and inductive reasoning, they are forced to hold some, which will affect their future production decisions and possibly aggregate fluctuations.

The structure of the model presented in this paper is inspired by a parallel literature on ABMs. This literature has stressed the importance of bounded rationality with simpler adaptive behaviors at the micro level but more complex interactions across agents using micro-simulation methods imported from physics. One of fundamental characteristics of micro-simulation methods is the ability to generate emergent phenomena at the aggregated level despite having simple behaviors at the micro level[2].Canonical examples in natural sciences include traffic jams generation, flocking birds, waves in fluid mechanics. Listing all contributions is outside the scope of this paper. In the

[2]Interestingly, the exact opposite is true of standard macroeconomic modeling strategies. Macro models tend to have complex microeconomic problems (micro- foundations) and implausibly simple interactions between these microeconomic units or even complete homogeneity (representative agent)

field of economics, original contributions were made in explaining segregation patterns (Schelling, 1971) and the Pareto distribution of firm size (Axtell, 2006) as emergent properties of simple individual policy rules. After the financial crisis and a renewed interest in macro-finance topics, these methods have been exploited to study systemic risks (Geanakoplos et al., 2012). The model is a direct extension from previous work by Delli Gatti et al. (2011) and Gualdi et al. (2015) with non-Walrasian labor and goods markets. The main contribution of this paper is to allow for inductive behaviors on the firm side to forecast demand.

3.3. Model

The model presented in this section is an agent-based model extended from Delli Gatti et al. (2011); Gualdi et al. (2015); Delli Gatti et al. (2019). As such it is a finite set of actions taken sequentially by agents every period. Unlike DSGE models there is no need to solve for a fixed point and equilibrium policy functions. Firms and consumers' policy functions are the primitives of the model. In practice simulating an ABM is like simulating a DSGE model after policy functions have been solved for, except that there is no flexible margin guaranteeing that supply equals demand[3].

As no Walrasian auctioneer magically sets prices to equate supply and demand, prices will be set individually by each firm according to their pricing policy function. Labor is not perfectly flexible and workers won't supply more of it than they want to. Mismatches on the labor and goods market will emerge as policy functions are aggregated up, sometimes leading to rationing or involuntary unemployment and inventories.

A key timing assumption is that supply, hiring and financing decisions are made before consumption decisions, so that contemporaneous demand is not part of the

[3]One advantage of this approach is that it does not require the counterfactual assumption that goods or labor supply adjusts perfectly to demand as in most New Keynesian models. See Smets and Wouters (2007); Christiano et al. (2005) for canonical examples.

information set of firms. This non-Walrasian economy can see supply equate demand on average, as firms are approximately accurate in forecasting demand over time as they learn the statistical properties of the demand they face.

3.3.1. The agents

There are F firms and N consumers. Consumers want to consume all their cash on hand every period. Because they get matched to a finite number of firms on the goods market and supply decisions are sunk, they might be rationed and forced to save cash for the next period. Consumers also own assets in the form of a fixed equity share $\frac{1}{N}$ of all firms and receive the corresponding dividends every period. Consumers are also endowed with one unit of labor every period that they can give to any firm in exchange for a wage. They only take two actions in this economy: consume and supply labor.

Firms are the main focus of this model and have multiple policy rules. They have a demand forecasting rule, a production planning rule, a pricing rule, a wage setting rule and a hiring rule that are partly dependent on each other. Because it is not an equilibrium model, some firms will make mistakes like overproducing or overpaying for labor factors and generate negative profits and possibly negative equity. A Firm that reaches negative equity at the end of the period is declared bankrupt and fires its workers. A new firm with mean equity is created in the next period so that the number of firms remains constant over time. At this stage it can help illustrate how the model works by listing the sequences of actions taken by consumers and firms in one iteration of the model.

One Iteration of the Model. The sequence of events can be summarized as follows:

(1) Firms forecast end of period sales $s_{i,t}$ conditional on information set $I_{i,t}$

(2) Firms formulate a production plan $(p_{i,t}, Q_{i,t})$ conditional on expected sales $\mathbb{E}_i s_{i,t}$ and past listed price relative to competition

(3) Firms finance production plan with external borrowing $B_{i,t}$ at exogenous rate R conditional on borrowing constraint

(4) Firms post vacancies $v_{i,t}$ and corresponding offered wage $w^o_{i,t}$ consistent with the production plan and financing restrictions

(5) Unemployed workers randomly draw a fixed number of vacancies, and apply to the highest paying wage offer

(6) Firms fill vacancies in a random order from the pool of unemployed applicants

(7) Firms produce goods conditional on production plan and financial/labor constraints

(8) Consumers set consumption and savings target given their income, wealth and employment status

(9) Consumers buy goods from a random set of firms that includes the biggest one, with a preference cheaper goods

(10) Consumers are forced to save when no goods are available to buy in the set of firms they visited

(11) Firms pay dividends if ex post profits are positive

(12) Firms go bankrupt if ex post equity is negative

(13) Back to step 1 and repeat

3.3.2. Firm Policy Rules

In this sections I will describe the firms' policy rules in more detail, starting with firms and their forecasting, production, price adjustment and financing rules. In a second part I will discuss what these rules imply in terms of inventory management and objective function.

Adaptive Sales Forecast. Before making a production decision, a firm needs to forecast sales. Naive firms have a simple extrapolative expectation rule[4]:

$$\mathrm{E}s_{i,t} = s_{i,t-1} + \xi$$

where ξ is random noise that could capture news shocks[5].

inductive firms have an adaptive expectation rule based on Arthur (1994). Firms keep a history of the last $2M$ observations of past sales $\{s_{i,t-2M}, \ldots, s_{i,t-1}\}$ and past aggregate sales $\{S_{t-2M}, \ldots, S_{t-1}\}$ each periods. Using this history, they can backtest possible forecasting strategies with memory M by assessing the forecast error of these strategies over the last M periods and pick the most accurate one. To give a simple example, if a firm has memory 2 and is trying to forecast demand in period 5. A possible strategy is to weight equally $\{s_{i,4}, s_{i,3}, S_4, S_3\}$. The firm can then look at the historical validity of that strategy by looking at its previous prediction for period 3 and period 4.

Formally, each new firm draws J strategies corresponding to two vectors of M scalar weighting past firm level and aggregate level sales. Every period, these J strategies are ranked by their ability to forecast the last M observations of individual sales. The one that performs best is selected to predict sales in the current period.

In the limit when $J \to \infty$, it is as if a firm with risk neutral preferences was forecasting expected sales by estimating an $AR(M)$ process with cross-validated parameters over the last M observations. In other words, this limit case would correspond to natural expectation formation with a zero weight on the rational expectation solution Fuster et al. (2010)

[4]Note that even in rational expectation models, expectation formation is always backward looking. Agents can only predict stochastic processes conditional on realized random variables, which by definition are past observables. So the deviation here is that the forecasting rule is simple, not that it is backward-looking.

[5]In practice small enough ξ shocks will tend to make pessimistic depression with zero forecasted sales, zero output and zero income an absorbing state.

This approach has several advantages. First, it a plausible way of formulating firms' expectation process when faced with Knightian uncertainty(Knight, 1921). Because of the decentralized nature of the model, neither the researcher nor the firm know the model well enough to have a well specified model of the demand process. In that way, it also corresponds to the type of stastistical rules that firms have been using in practice Larsen-Hallock et al. (2022). Therefore, firms cannot formulate their decision problem as an optimization problem and make their production plan conditional on equilibrium demand and equilibrium prices as they would in a competitive equilibrium model. Real and nominal rigidities combined together beget model ambiguity so that model consistent expectations are not well-defined. Firms can only rely on imperfect statistical forecasting rules, as opposed to exogenous knowledge about the model itself, to make decisions.

Second, this specific adaptive approach is versatile among statistical forecasting rules and accommodates complex non-linear dynamics while limiting memory and computational costs. For example, exponential smoothing and auto-regressive processes are both nested in the possible set of strategies, as long as the set of strategies is large enough. Moving averages, trends and up to M-period cycles are other possible strategies allowed by this adaptive approach

Third, it offers a strategy selection process that satisfies the condition that firms try to "optimize" in a loose sense and would not pick an arbitrarily bad forecasting rule like naive firms would.

Finally, this approach is computationally tractable from the modeler's perspective, which is essential given that the optimization step is repeated F times every period.

Why not use OLS instead? A key specificity of this algorithm is to allow for model uncertainty through the backtesting of the strategies M times. If firms were simply running OLS on past observation to predict future sales, they would be using a different but still naive form of extrapolation based on past observations. A possible alternative would be for firms to run a cross-validation algorithm with OLS estimation. With that alternative strategy, firms would run OLS M times over the M different subsamples and perform some kind of model averaging. This alternative approach is left for future research.

Production plan. After making its demand forecast, the firm decides how much to produce and at what price. Let us denote $q_{i,t-1}, s_{i,t-1}$ the quantities produced and quantities sold last period and $\bar{p}_{t-1}$ the average price for the consumer good as listed by all firms on the market. For simplicity, firms are only able to adjust prices or quantities every period. The action space can therefore be divided into four cases:

(1) **overproduction:** $s_{i,t-1} < q_{i,t-1}$ and $p_{i,t-1} \leq \bar{p}_{t-1}$: firm decreases production relative to expected demand up to a fixed factor η_q

(2) **uncompetitive:** $s_{i,t-1} < q_{i,t-1}$ and $p_{i,t-1} > \bar{p}_{t-1}$: firm reduces listed price by up to fixed factor η_p

(3) **Increase margins:** $s_{i,t-1} = q_{i,t-1}$ and $p_{i,t-1} \leq \bar{p}_{t-1}$: firm increases listed price up to fixed factor η_p

(4) **Gain market share:** $s_{i,t-1} = q_{i,t-1}$ and $p_{i,t-1} > \bar{p}_{t-1}$: firm increases production relative to expected demand up to a fixed percentage η_q

Hence, one can write the price and sales target rules as highly non-linear functions of multiple state variables:

$$s_{i,t}^T = f(\mathbb{E}s_{i,t}, s_{i,t-1} - q_{i,t-1}, p_{i,t-1}, \bar{p}_{t-1}) \tag{3.1}$$

$$p_{i,t} = g(\mathbb{E}s_{i,t}, s_{i,t-1} - q_{i,t-1}, p_{i,t-1}, \bar{p}_{t-1}) \tag{3.2}$$

Note that firms won't be allowed to price below their average cost of goods sold to avoid firms selling at a loss. Since the only factor of production is labor, the average cost of goods sold is equal to the average wage paid by the firm. When work duration is one period, this will correspond to past posted wages. When work duration is longer, cost of goods sold will be as weighted average of wages posted in previous periods.

For firms that want to minimize inventories while maximizing sales (say for example because the cost of an unsold good is equal to the cost of a missed sale in equilibrium), the quantity effectively supplied will be the sales target net of existing inventories:

$$q_{i,t} = s^T_{i,t} - I_{i,t-1}$$

If the forecast is correct for all firms in the economy, the economy will be in equilibrium. In all other cases, there will be either rationing or involuntary inventories. In that sense, the economy is partially supply determined, as production plans collectively set a ceiling on how much can be consumed each period. The economy remains partially demand determined as well given that overproduction remains a possibility.

Financing needs. Firms have a simple linear production function in labor with constant total factor productivity for simplicity. The production target is therefore equal to the number of workers needed for production, and production costs can be computed as:

$$payrolls_{i,t} = \sum_{c=1}^{q_{i,t}} W_{c,t}$$

Firms will preferably finance payrolls through internal funds $E_{i,t}$. If these funds are insufficient, firms will borrow up to a leverage constraint:

$$B_{i,t} = \min\left(\max\left(payrolls_{i,t} - E_{i,t}, 0\right), \ell(E_{t,i}, s_{i,t-1} \times p_{i,t-1}, I_{i,t-1}\right)$$

The leverage constraint ℓ is a function of current equity level $E_{i,t}$, past revenues and pledgeable inventories $I_{i,t-1}$. This is consistent with the recent corporate finance literature on corporate borrowing constraints that demonstrates how corporate financing relies heavily on cash flow constraints, as opposed to asset constraints (Lian and Ma, 2021). This is an additional deviation from Delli Gatti et al. (2011) which uses capital constraints.

If the firm is financially constrained, it will not be able to hire as many workers as expected and will therefore reduce its production level to $q^c_{i,t} \leq q_{i,t}$. If the level of collateral is too low, it is possible that not even existing workers can be paid. In that case the firm is forced to fire workers until the borrowing limit allows financing of the wage bill.

For aggregate accounting consistency in a closed economy, aggregate firm borrowing B_t will show up as consumer wealth W. The financial system is essentially transparent here, with consumers lending to firms directly.

3.3.3. The labor market

All work contracts on this labor market have a fixed duration and can only be terminated early when firms face bankruptcy or tight financing constraints. The contract duration is a constant parameter τ identical across firms. It can be interpreted as the number of periods workers are expected to work with the same firm. Each period firms will have an existing legacy payroll as a consequence of past contracts period $L^l_{i,t}$ corresponding to contracts that have not yet expired in period t. Firms will post vacancies if they need new workers to reach their production target. Vacancies will therefore be equal to the difference between production target and legacy payroll. Where legacy payroll is last period payrolls minus worker exit at the beginning of the period $L_{i,t-1} - E_{i,t}$. Exits could be due to contract expiration, or financially constrained terminations. Posted vacancies are capped by zero since they cannot fire workers for non-financial reasons.:

$$V_{i,t} = \max\left(Q^c_{i,t} - L^l_{i,t-1}, 0\right)$$

The assumption here is that the firm does not post negative vacancies, i.e. fire workers, simply because production plan is small than existing workforce. In effect, layoffs only happen because of financial constraints and bankruptcy in this economy. The firm also needs to choose the posted wage at which new workers are hired:

$$w^p_{i,t} = \max\left(\bar{w}, w_{i,t-1}\left(1-\eta_w\right)\right) \text{ if } V_{i,t-1} = 0$$

$$w^p_{i,t} = \max\left(\bar{w}, w_{i,t-1}\left(1+\eta_w\right)\right) \text{ if } V_{i,t-1} > 0$$

where $\bar{w}$ is a parameter setting the minimum wage. Firms will increase posted wages if they struggled to fill vacancies in the previous period, and lower wages if they filled all vacancies, up to a minimum wage limit $\bar{w}$.

Once firms have publicly posted vacancies, unemployed workers will randomly select S_w of them, and rank them in descending order of wage offered. Workers apply to the top S_j. Firms that move first in accepting the application gets the worker, even if their posted wage was not the at the maximum of the distribution. S_j can be thought as the parameter governing how many applications unemployed workers send in the period. Once a worker is matched to a firm, he will work for that firm for a fixed duration L_d, unless the firm goes bankrupt or faces financial constraints. Workers that have expiring contract will prioritize their former employer in the matching process, regardless of posted wages.

Firms then select the needed number of workers from the pool of applicants. Notice that if F_w is larger, many workers will apply to the same high paying job so that very few will effectively be matched. Many workers will be left unemployed and firms will have unfilled vacancies. Hence this parameter directly drives the efficiency of the matching process on the labor market.

After all vacancies have been filled or the pool of unemployed workers has been exhausted, firms update their production plan based on their available labor force. Firms end up producing

$$q_{i,t}^{realized} \leq q_{i,t}^{c} \leq q_{i,t} = s_{i,t}^{T} - I_{i,t-1}$$

Hence the effective production function is an indirect non-linear function of the labor supply, expected demand, existing inventories, financial constraints and labor constraints. Even though the production technology is apparently trivial, the sequencing of firm decisions generates complex non-linearities ex post. Aggregate fluctuations will affect the availability of unemployed workers, corporate financing, sales targets, inventories and feedback into the firm production plan. This mechanism will also generate non-trivial path dependencies. The presence of non-linearities will also make revenues per factors of production endogenous, as realized sales will fluctuate for a given workforce.

3.3.4. The goods market

Because the focus of this paper is on the production side of the economy, the goods market is made as simple as possible. All consumers are hand-to-mouth. They want to consume everything they have every period. Yet, demand forecasting remains non-trivial for firms because firms do not know how many workers other firms will hire and at what wage, so that aggregate income is highly uncertain ex ante. Supply might be too small to fulfill demand, if for example wages have increased due to tight labor markets, in which case consumers will be rationed and forced to save.

In practice, consumers randomly select S_g firms as a pool of supplying firm to visit. They also add to this list the largest firm they have interacted with last period, which captures household preference for reliability in inventories availability. They rank these firms in ascending order of posted price. So long as their purchasing power and the available inventories are non-zero, consumers spend their money in the same store. Once inventories reach zero, a they visit the next cheapest store and keep on spending as much as they can until their cash on hand is exhausted or the store is empty, in which case they visit the next store in their list. This iterative process stops if consumer purchasing power goes to zero, or if it is inferior to the price of one unit of good sold (units are indivisible here), or if the consumer has searched through all the S_g firms.

Consumers enter this market in a random order, and some might be left out from the most affordable stores. Unsold inventories are stored for next period, net of depreciation:

$$I_{i,t} = (1-\delta)(I_{i,t-1} + q_{i,t}^{realized} - s_{i,t})$$

3.3.5. Firms entry and exit

After trading occurs on the goods market, firms' revenues are captured by turnovers $p_{i,t}s_{i,t}$. Net profits are computed as:

$$\pi_{i,t} = p_{i,t}s_{i,t} - payrolls_{i,t}^{realized} - (R-1)B_{i,t}$$

where $payrolls_{i,t}^{realized}$ corresponds to labor costs after considering financial and labor constraints. The distribution of profits to consumers follows the simple rule:

$$D_{i,t} = \phi \max(\pi_{i,t}, 0)$$

where ϕ is the share of profits distributed to households.

Hence the law of motion for equity is the following:

$$E_{i,t} = E_{i,t-1} + \pi_{i,t} - D_{i,t} - B_{i,t} = E_{i,t-1} + (1-\phi)\pi_{i,t} - B_{i,t}$$

If equity turns negative, the firm is considered bankrupt. It exits the pool of firms, with a corresponding tax on household cash on hand to capture the equity loss. Bankrupt firms are replaced by a new firm with fresh equity.

New firms start with an equity level that corresponds to the average equity in the population of firms truncated at the 5th and 95th equity percentile. The posted wage and listed price they start with correspond to the average values in the truncated pool of firms. One thing to note is that the process of firm exit and entry can have dramatic effects on model results. If entering firms have too low a posted wage and price, they will capture the entire market quickly and make prices converge to the minimum wage. If entering firms' prices and wages are too high, they can fuel large inflation dynamics.

3.4. Calibration

With 22 parameters including 3 parameters governing initial condition, the parameter space is large. It is not computationally feasible to explore the entire parameter space. The goal of the calibration is to make the model a relatively accurate representation of the economy so that experiments and counterfactual varying one parameter at a time in the model can have external validity. To restrict degrees of freedom, I will calibrate externally parameters that have a reasonable interpretation outside of the model. It is one of the advantages of ABMs that some parameters have relatively transparent empirical interpretation. Other deeper parameters capture multiple mechanisms at once and cannot be extracted from the data so easily. For these 12 parameters, I will use a simulated methods of moments approach described in 3.4.2 to calibrate them.

3.4.1. External Calibration

Table 3.1 shows the values of each parameter externally calibrated using multiple sources. The number of workers in the economy I is set to match the labor force in the US at a scale of 1 per 200 000 due to computational constraints. The number of firms is calibrated to match a ratio of workers to firms of 4.8% as in the data.

The quarterly depreciation rate of inventories is set so that in 4 periods the annual rate depreciation rate matches a ratio of corporate depreciations and amortizations over GDP 5.54%. Similarly the qcompounded quarterly interest rate matches the average rate on Aaa corporate bonds in 2023 according to Moody's. Dividends share of profits is calibrated to match the ratio of net dividends to net income from the NIPA tables in Q4 2022. The work contract duration matches the median work duration of 4.1 years as reported by the Bureau of Labor Statistics. The minimum wage is calibrated to match a ratio of minimum wage to average wage of 33 %. Given our wage numeraire it roughly corresponds to what an employee at full time would earn over the course of a quarter.

All the other parameters will be estimated according to the procedure explained below. One thing to note is that TFP is inversely related to the average price of goods sold and can be used to pin down the average cost of a transaction.

In the initial period, firms are uniformly distributed along two dimensions: how much equity they have and their markup over labor costs. Firm wages are distributed $\mathcal{U}(\bar{w}, w_{max})$ so that the average wage is a function of $\bar{w}, w_{max}$. The maximum markup μ_{max} is calibrated to 134% to match an average markup of 67.5%, in line with recent estimates by De Loecker et al. (2020). Prices are therefore distributed $\mathcal{U}(\bar{w}, w_{max}(1+\frac{\mu_{max}}{100}))$ as a markup over offered wages. Calibration of wages is done starting from nominal GDP in the current year and taking a labor share of 59.7%

	Parameter	Value	Units	Source
I	Number of workers	830	1 per 200 000	WB
F	Number of firms	40	1 per 200 000	SUSB - WB
$\frac{I}{F}$	Consumer/firm ratio	4.8	$\frac{I}{F}$	SUSB - WB
δ	inventories depreciation rate (Q)	1.36	%	IRS - FRED
r	interest rate (Q)	1.11	%	Moody's - FRED
ϕ	Dividends share of profits	66.23	%	BEA
L_d	Work contract duration	16	Quarters	BLS
$\bar{w}$	Minimum wage	6	\$1000	FRED
μ_{max}	Maximum markup of firms	134	%	De Loecker et al. (2020)
E_{max}	Maximum equity of firms	1781	\$1000	FRB
w_{max}	Maximum posted wage	41	\$1000	BEA

Figure 3.1. Externally Calibrated Parameters

Equity is distributed $\mathcal{U}(0, E_{max})$. Since households own firms and they only hold equity, E_{max} is used to match the ratio of US net worth over GDP of 566% calculated from the financial accounts of the United States using the Z.1 table of the Federal Reserve Board of Governors and FRED.

Initial expected demand $\mathbb{E}_i s_{i,0}$ pins down how much firms are producing in the first period and does not have a clear empirical interpretation. Given TFP, the number of firms and workers, a target unemployment rate in the first period will define it implicitly. It will be set to match an unemployment rate of 5% in line with US long-term average.

Results do not tend to be affected by μ_{max} as the distribution of markups quickly converges to its ergodic distribution. E_{max} does affect the likelihood of observing bankruptcies early on in the simulation, which could make the system converge to an absorbing state with no employment and no production. When equity is too low relative to minimum wage, firms simply cannot ramp up production, so that income never picks up and demand remains depressed. If $\mathrm{E}_i s_{i,t}$ is large, the economy can momentarily see a boom in production with rising prices and wages, but will usually revert back to its ergodic distribution after sufficiently many iterations.

3.4.2. Internal Calibration: Simulated Method of Moments

The goal of the calibration exercise is to make the model sufficiently representative of the US economy so to have a credible sandbox on which to run parameter experiments and see how firms behaviors affect aggregate outcomes. The remaining 12 parameters (table 3.1) are estimated by simulated method of moments (SMM). 20, 000 parameter vectors are drawn from a 12-D grid space described in table 3.2. For each parameter vector, the model is simulated for 400-periods simulation. To reduce the impact of initial conditions, the first 100 periods are dropped from the simulated sample. The targeted moments are the covariance matrix elements between unemployment, real GDP, CPI and average hourly earnings. The average unemployment rate is also added as additional moment 3.3. The preferred set of parameters listed in table 3.1 is the one that minimizes the mean squared error relative to their empirical counterpart in the US economy, as estimated from post-WW2 data (1950-2023).

Several elements make the calibration plausible. First of all, all estimated parameters are interior to the grid space (table 3.2), suggesting that bounds are not driving

	Parameter	Value	Units
θ	Share of smart firms	21	%
M	Memory size of inductive firms	27	periods
J	Number of strategies of inductive firms	17	# of weight vectors drawn
η_p	Maximum prices adjustment rate	24.01	%
η_q	Maximum quantities adjustment rate	9.26	%
η_w	Maximum wages adjustment rate	2.54	%
S_g	Goods search friction	37	# store visits per period
S_l	Labor market friction	5	# job screenings per period
S_j	Job applications limit	2	# applications per period
ℓ	Maximum leverage ratio	1.54	$\frac{B_{i,t}}{p_{i,t}s_{i,t}+Inv+E_t}$
ξ	Maximum size of expected demand shock	6	goods per period
TFP	Real output per worker	2	# goods produced

Table 3.1. Calibrated parameters using Simulated Method of Moments

estimation outcomes. Second, the calibration implies that wages are more rigid than prices, with quantities in between the two. This is consistent with the macroeconomic literature stressing the importance of wage rigidities over price rigidities in explaining output fluctuations (Christiano et al., 2005). Third, the estimated maximum leverage ratio is in the range of values used by banks when providing funding to firms.

	Parameter	Lower Bound	Upper bound
θ	Share of smart firms	0	100
M	Memory size of inductive firms	1	100
J	Number of strategies of inductive firms	10	300
η_p	Maximum prices adjustment rate	1	30
η_q	Maximum quantities adjustment rate	0.5	15
η_w	Maximum wages adjustment rate	0.5	30
S_g	Goods search friction	1	40
S_l	Labor market friction	1	40
S_j	Job applications limit	1	S_l
ℓ	Maximum leverage ratio	1	5
ξ	Maximum size of expected demand shock	0	10
TFP	Real output per worker	1	15

Table 3.2. Sampling Space for the SMM

An interesting finding is that the estimated model suggests firms have limited forecasting ability. Only 21% of them are estimated to be using inductive reasoning. And when they do, their memory is limited to roughly 7 years. They also use very few strategies drawing only 17 vectors of weights to forecast from past observations.

3.5. Quantitative Results

In this section I will present some of the key findings of the paper. First, I will analyze the behavior of the artificial economy under SMM calibration. Then I will make experiments changing the share of inductive firms to see how it impacts the business cycle. Two strategies are adopted to limit the influence of initial conditions. For the same calibration and share of inductive firms, 40 parallel economies are simulated to control for possible initial condition effects. The first 300 observations are also discarded from all samples to limit further the influence of initial conditions. Ultimately, there are 40 economies per parameter set with 1001 periods for each.

3.5.1. Baseline Calibration

The baseline calibration matches first and second order moments of the US data fairly well. Unemployment is well-behaved around its average of 5.8%[6]. An interesting feature of the model under this calibration is the emergence of very long "super-cycles" of price/wage spirals. Prices and wages would increase progressively as inventories are low and firms struggle to keep up with demand. In the next phase of the cycle firms would be forced to cut prices and wages as inventories are high and demand is low. This would lead to a recession and the cycle would start again. These cycle drive nominal GDP up or down in the long-run despite constant potential output. The dynamics are similar to a model of conflict inflation, where firms fight over the price of their goods relative to wages and the price of other goods.

Most distributions generated by the model tend to be right skewed, which is consistent with the data as well. For example, the distribution of prices and wages are right skewed with a long right tail. The distribution of markups is also right skewed with a

[6]All summary statistics are presented in appendixes C.1 and C.2

long right tail. The distribution of sales follows a similar pattern due to the emergence of super-star firms that receive most of consumer demand.

Under this parametrization, the financial constraint is not really binding. Firms do not use much private debt to finance production because they have plenty of equity available. When needed firms are able to borrow as much as they need to finance their operations. Few firms go bankrupt as well. This is probably due to the fact that by matching the equity to GDP ratio without having volatility in asset prices, firms have a large cushion of equity to absorb demand shocks in the model.

The model also generates volatility in the vacancy rate, oscillating around 25% of job postings on average.

Regarding the forecasting process, it appears that inductive firms end up making more forecasting mistakes than naive firms in at least some simulation runs. As a result, naive firms end up having a higher share of sales than inductive firms. One way to explain this apparent paradox is to realize that with only 17 strategies to choose from in a space of $([-1,1]^21, [-1,1]^21)$, inductive firms are actually quite limited in their ability to forecast. They are also limited in their memory, which is only slightly more than 5 years. As a result, they end up making more mistakes than naive firms. Another factor that could play a role is that given the relatively simple nature of that artificial economy, a simple $AR(1)$ process might be a good model.

As described in section 3.3, the microeconomic policy rules are chosen to be simple, without heterogeneity in productivity across households and firms. Yet, the model endogenously generates relatively deep heterogeneity in most calibrations.

First of all, the heterogeneity in firm size as measured by sales or wealth accumulates over time. Some firms do consistently well at predicting demand, pricing competitively and producing the right amount so as to capture a large share of the goods market. The economy is relatively quickly concentrated between few superstar firms. At the

Table 3.3. Comparison of Empirically Estimated and Simulated Moments

	US	Model
$\bar{u}$	5.7477	5.8876
σ_u	0.3185	0.6065
σ_p	0.0199	0.0220
σ_w	0.0110	0.0082
σ_g	0.0247	0.0949

same time, the distribution of markups is wide and skewed to the right as previously illustrated in Axtell (2006), with a mode slightly below zero and an average above zero. This is consistent with superstar behaviors with some firms capturing most of the profits while more than 50% of all firms remain unprofitable.

Another interesting finding is the emergent property that firm size is positively correlated with higher wages. This is true despite not having any difference of productivity across firms or workers, which goes against standard microeconomic reasoning on marginal costs of production factors. Putting the "superstar firms" literature in perspective (Autor et al., 2020), it illustrates how firms with superior understanding of the demand process can afford to pay higher wages, even if workers are identical, and even if the superstar firms have no specific productivity advantage.

At an aggregate level, the economy can be relatively efficient at equalizing supply and demand. As detailed in the next section, the economy will be closer to equilibrium on average with more inductive firms. The economy also shows strong path dependency. Series of positive demand surprises can trigger firms to ramp up production in expectation of future positive surprises, leading to higher payrolls and aggregate income for some time. An extreme case of path dependency is the existence of absorbing states. Low output, low income, and self-fulfilling pessimistic expectations tend to be an absorbing state as firms with low revenues, equity have tight financial constraints.

3.5.2. Sensitivity analysis

On the production side. In the macroeconomic literature there is large debate about the role of sticky prices and wages in generating business cycle amplification. This model makes it possible to test whether economic activity would be more volatile under different pricing regimes. Setting η_p to zero, we have fully sticky prices. Firms can only adjust their production plan through the quantity margin. It appears that under this calibration economic fluctuations in real variables are amplified, with higher volatility in output and unemployment. When η_w is set to zero however, there is no induced increased volatility. It is a puzzle as to why the two margins of adjustment would affect volatility differently. I suspect that wage differentiation makes the labor market less liquid and tends to amplify fluctuations. Uniform wages therefore tend to dampen this channel.

Perhaps surprisingly, having very flexible prices and wages also tend to increase volatility. Firms react too much to current economic conditions through price changes and amplify local shocks. It is when prices are reasonably flexible, but not too much, that the economy is most stable. This is consistent with stories of debt-deflation Fisher (1933) where flexible prices can in fact amplify shocks rather than dampen them.

On the demand side. One key parameter on the demand side is the number of firms that households are visiting before buying goods. When this number is small, competition across markets is limited. One firm tends to capture the entire market over time. As it grows larger and larger, more households keep it in their list of firms. Prices go down at first. Then prices go up as the number of bankruptcies increases and competition diminishes in the economy. Private debt is endogenously higher in this economy than before. This is because small firms have limited equity and need external financing to survive. Reintroducing competition increases prices at first, since small firms with limited capacities need to cope with the new demand coming in. After some time, the number of bankruptcies and the firm distribution becomes less skewed. Output also increases and the level of unemployment goes down with higher competition.

3.5.3. Varying the Share of Inductive Firms

What happens when the share of Inductive firms in the economy changes relative to the baseline calibration of 21%, holding everything else constant? In this section, I will examine the consequences of varying the share of inductive firms from 0% to 100% in increments of 25%.

Non-monotonous Impact on Forecast Errors. The first question one can ask is whether firms collectively make less forecasting mistake when more firms use inductive reasoning. As shown in figure 3.2, the mean squared forecast error is increasing in the share of inductive firms. It is highest when the share of inductive firms is 100%, and lowest when the share of inductive firms is 0%.

Figure 3.2. Distribution of Forecast Errors across Simulations for varying θ

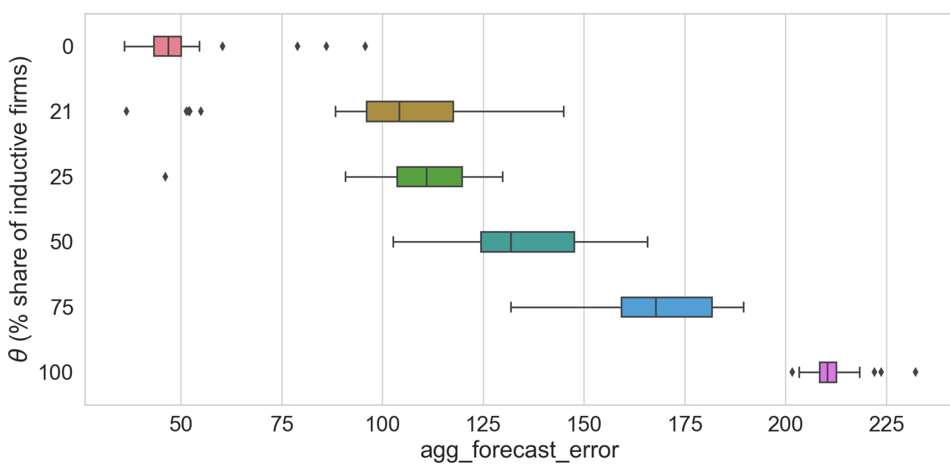

Notes: Aggregate forecast errors in a given period are computed as the mean squared error of individual forecast errors, where a firm's forecast error is simply the difference between forecasted demand and realized demand. Aggregate forecast errors are then averaged over the 1000 periods in a given simulation. For each value of θ, the share of inductive firms in the economy, a boxplot depicts the distribution of averaged aggregate forecast errors across the 40 simulation runs.

Figure 3.3. Inductive Firms - Distribution of Forecast Errors across Simulations for Inductive Firms

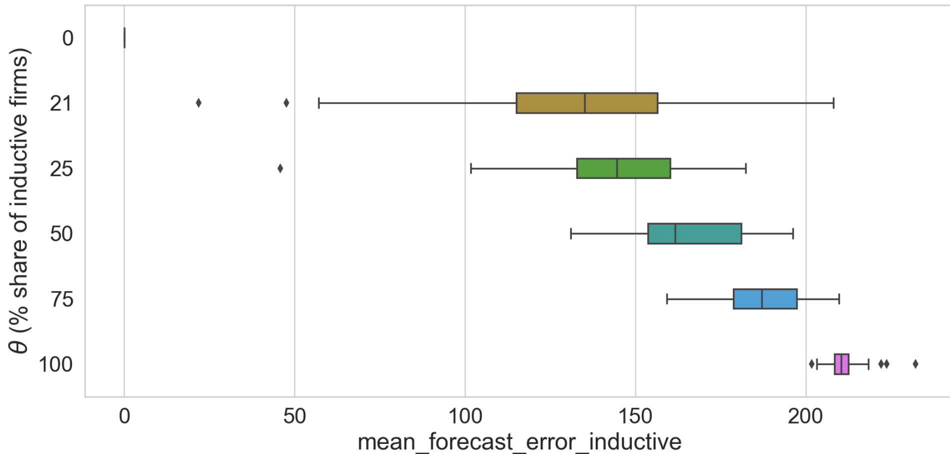

Notes: Aggregate forecast errors in a given period are computed as the mean squared error of individual forecast errors, where a firm's forecast error is simply the difference between forecasted demand and realized demand. Aggregate forecast errors are then averaged over the 1000 periods in a given simulation. For each value of θ, the share of inductive firms in the economy, a boxplot depicts the distribution of averaged aggregate forecast errors across the 40 simulation runs.

Figure 3.4. Naive Firms - Distribution of Forecast Errors across Simulations

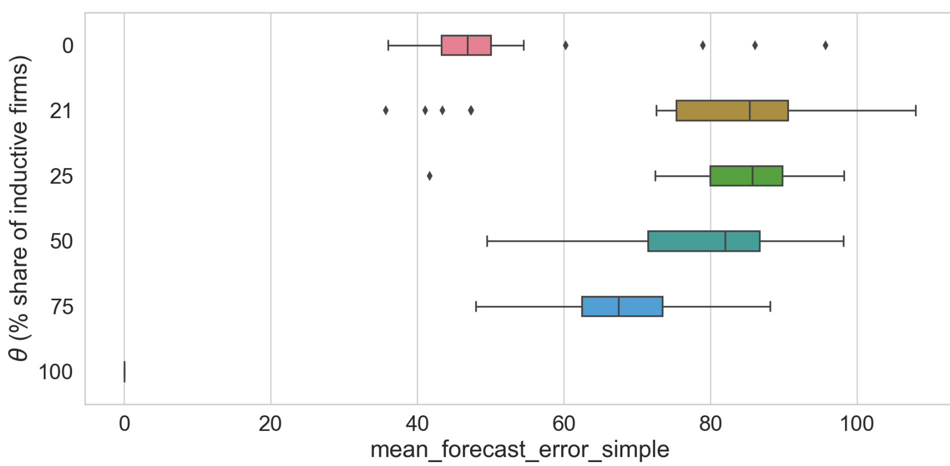

Notes: Aggregate forecast errors in a given period are computed as the mean squared error of individual forecast errors, where a firm's forecast error is simply the difference between forecasted demand and realized demand. Aggregate forecast errors are then averaged over the 1000 periods in a given simulation. For each value of θ, the share of inductive firms in the economy, a boxplot depicts the distribution of averaged aggregate forecast errors across the 40 simulation runs.

To better understand this somewhat counterintuitive result, figure 3.3 and figure 3.4 show the distribution of forecast errors for inductive and simple firms respectively. The picture that emerges is that naive firms are relative consistent in their forecast error regardless of the share of inductive firms, although they tend to do somewhat worse once inductive firms are introduced. On the contrary, inductive firms perform a lot better when few use inductive reasoning. It appears that $\theta = 100\%$ is a case where they do particularly badly. One element to keep in mind is that given the baseline calibration, firms are endowed with a limited number of 17 strategies to choose from. As a result, there are not necessarily better off than naive firms. In fact, their sales performance is on average worse than naive firms. This is illustrated in figure 3.5 which shows that in most simulation runs, inductive firms sell less than naive firms on average.

This suggests a paradox in that the presence of inductive firms might make inductive reasoning less performant as aggregation of behaviors change the characteristics of the economy. One possible intuition for this is result is that when all firms are using inductive reasoning, the commonality of the expectation process across firms shapes the business cycle so as to make expectations less likely to be self-fulfilling, as if the information was already incorporated on the market already. In section 3.5.3, I will explore how business cycle characteristics are indeed affected by the share of inductive firms.

Figure 3.5. Distribution of Relative performance of Inductive Firms across Simulations

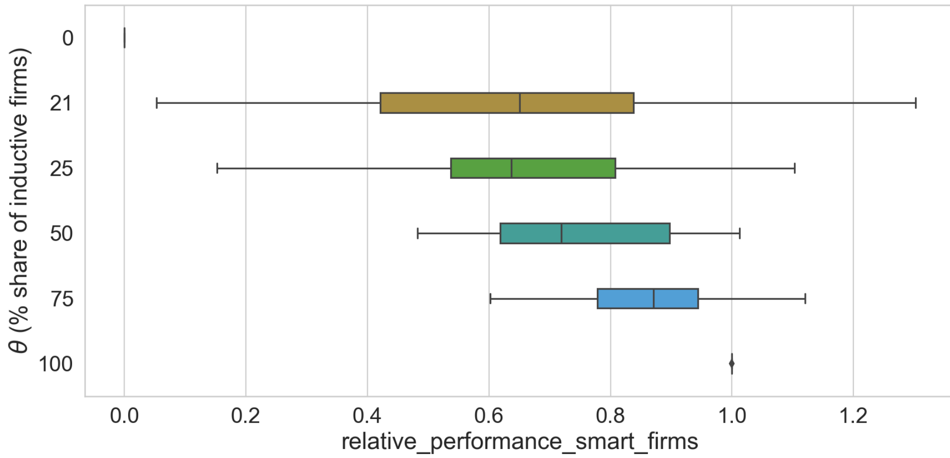

Notes: Aggregate relative performance of inductive firms is the ratio of sales per firm for inductive firms to sales per firm for naive firms. A ratio above 1 indicates that inductive firms sell more goods than naive firms on average. This relative performance ratio is averaged over the 1000 periods in a given simulation. For each value of θ, the boxplot depicts the distribution of averaged relative performance ratio across the 40 simulation runs.

Changes in Business Cycle Characteristics. As suggested by the previous section, the share of inductive firms in the economy has an impact on the predictability of sales. This suggests that business cycle characteristics could be strongly affected.

To illustrate this, one can look at the variation of certain moments of the business cycle across simulations. Figure 3.6 shows the distribution of the standard deviation of aggregate sales for different values of θ. The picture that emerges is that the standard deviation of aggregate sales is negatively correlated with θ. At the same time, aggregate sales tend to be significantly less auto-correlated, going from 0.9 in absence of inductive firms to 0.5 (figure 3.7). It is therefore not surprising that firms struggle to forecast a process that is now less persistent.

Figure 3.6. Standard Deviation of Aggregate Sales

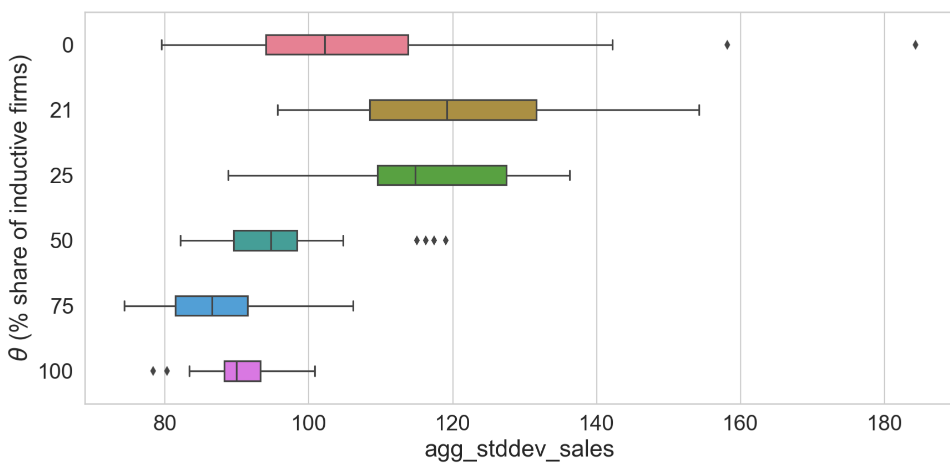

Notes: The statistic is computed over the 1000 periods in a given simulation. For each value of θ, the boxplot depicts the distribution of that statistic across the 40 simulation runs.

The same patterns emerge with the unemployment rate. Again, the standard deviation decreases with the share of inductive firms, and the autocorrelation is dramatically reduced as shown in Figure 3.8 and 3.9. The unemployment rate process goes from very persistent with an autocorrelation above 0.9 with naive firms to a process that has no persistence when $\theta = 100\%$.

Figure 3.7. Auto-correlation of Aggregate Sales

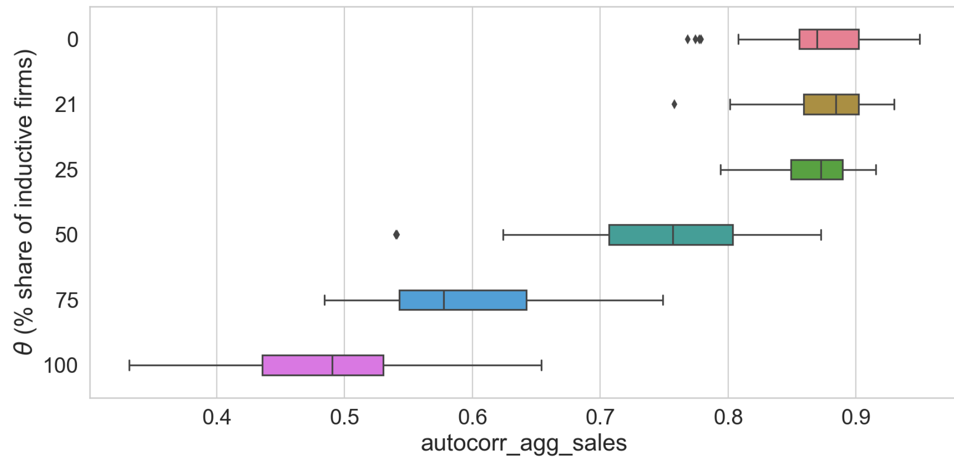

Notes: The statistic is computed over the 1000 periods in a given simulation. For each value of θ, the boxplot depicts the distribution of that statistic across the 40 simulation runs.

Figure 3.8. Standard Deviation of Unemployment Rate

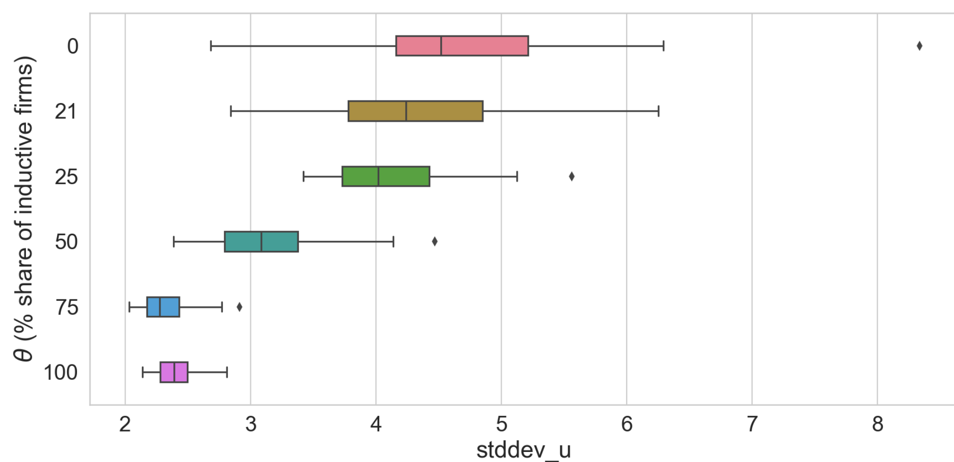

Notes: The statistic is computed over the 1000 periods in a given simulation. For each value of θ, the boxplot depicts the distribution of that statistic across the 40 simulation runs.

Prices and wages somewhat behave similarly, with dramatically lower standard deviation as the share of inductive firms increase (figure 3.11, 3.11. But persistence remains high and close to 1 in all parametrizations.

All in all, firms face a dramatically different economies when their competitors are also using inductive reasoning, and it leads to non-trivial aggregation effects.

Figure 3.9. Auto-correlation of Unemployment Rate

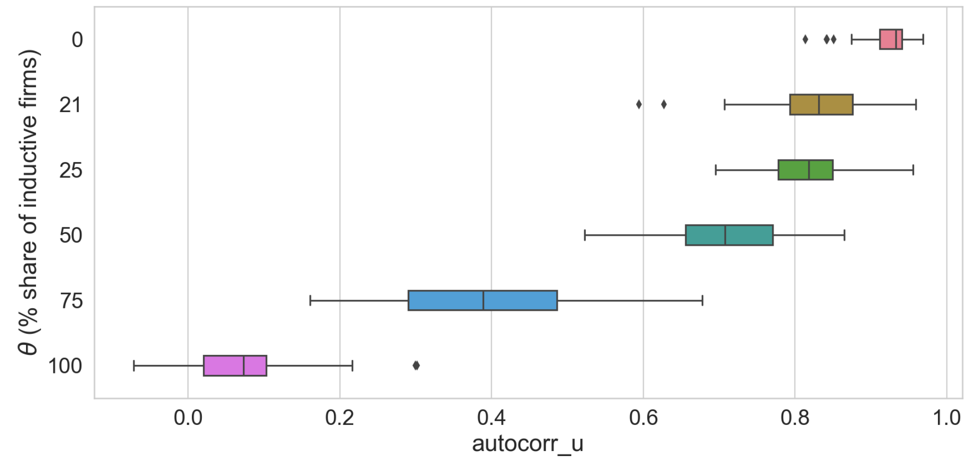

Notes: The statistic is computed over the 1000 periods in a given simulation. For each value of θ, the boxplot depicts the distribution of that statistic across the 40 simulation runs.

Figure 3.10. Standard Deviation of Wage Index

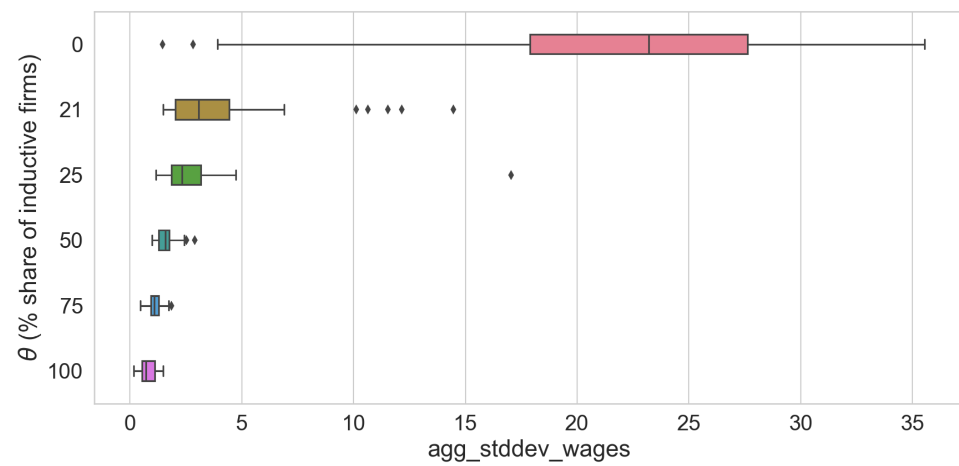

Notes: The statistic is computed over the 1000 periods in a given simulation. For each value of θ, the boxplot depicts the distribution of that statistic across the 40 simulation runs.

Figure 3.11. Standard Deviation of Price Index

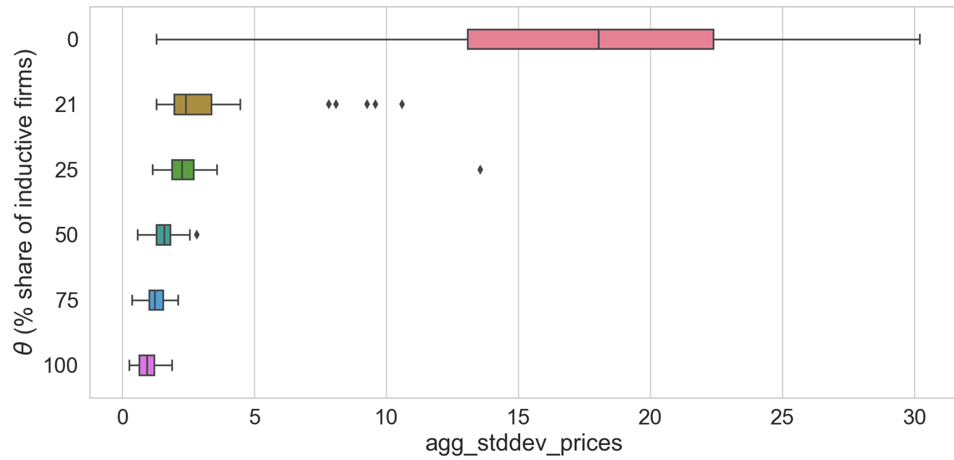

Notes: The statistic is computed over the 1000 periods in a given simulation. For each value of θ, the boxplot depicts the distribution of that statistic across the 40 simulation runs.

Impact on cross-sectional heterogeneity. The previous section has shown that the share of inductive firms has a significant impact on business cycle aggregates. What about the distribution of firm outcomes in the cross-section? This section shows that it also has a significant impact on the cross-sectional distribution of firm sales, prices, wages and markups.

Despite lower standard deviation in the aggregate, firm sales end up being more dispersed in presence of inductive reasoning (figure 3.12). This is not surprising given that naive firms all have the same forecasting strategy (putting a weight of 1 on past sales), while all inductive firms draw different sets of 17 strategies to choose from. In that sense increasing the share of inductive firms also increases the number of strategies used to forecast demand in the population of firms.

Figure 3.12. Average Standard Deviation of Sales in the Cross-Section of Firms

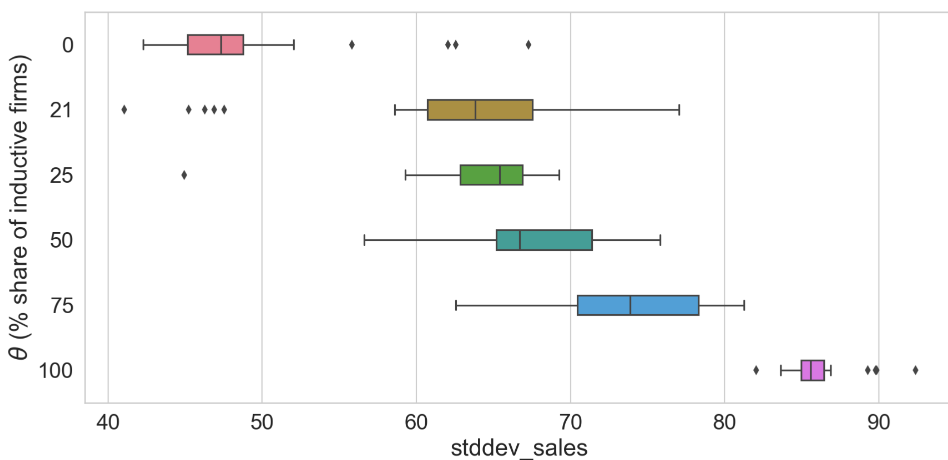

Notes: The statistic is computed each periods over the 40 firms in the sample and averaged over the 1000 periods in a given simulation. For each value of θ, the boxplot depicts the distribution of that statistic across the 40 simulation runs.

Yet, the variance of posted prices and markups shrinks, suggesting that labor and goods market converge to unique prices and wages. This is shown in figure 3.13 and 3.14.

Figure 3.13. Average Standard Deviation of Listed Prices in the Cross-Section of Firms

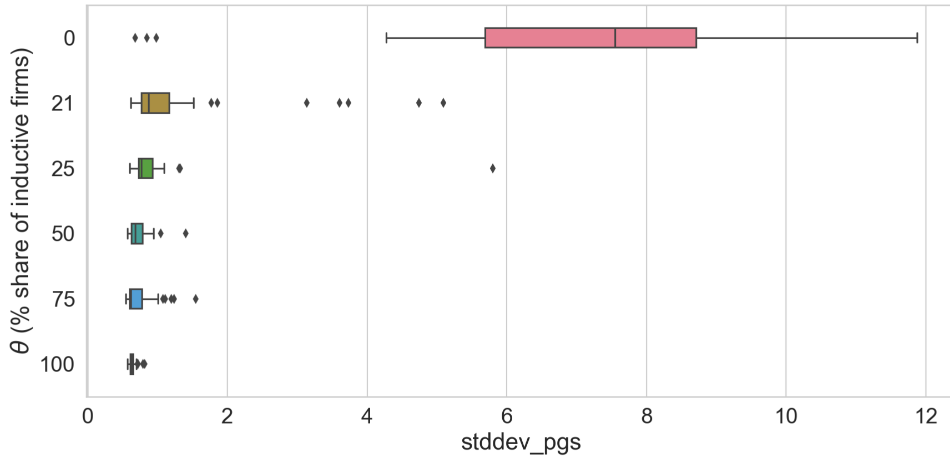

Notes: The statistic is computed each periods over the 40 firms in the sample and averaged over the 1000 periods in a given simulation. For each value of θ, the boxplot depicts the distribution of that statistic across the 40 simulation runs.

Figure 3.14. Average Standard Deviation of Markups in the Cross-Section of Firms

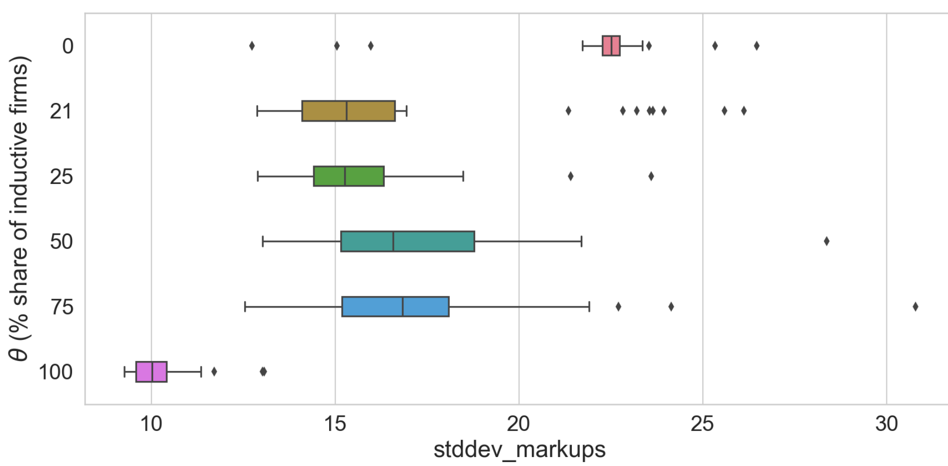

Notes: The statistic is computed each periods over the 40 firms in the sample and averaged over the 1000 periods in a given simulation. For each value of θ, the boxplot depicts the distribution of that statistic across the 40 simulation runs.

Why are firms paying different wages to identical workers? As observed in the previous section, wages offered are not unique but have a cross-sectional distribution. There are several reasons why wages can differ. First of all, different wages can compensate for different worker abilities. So far in the model all workers are identical with respect to labor productivity and we will therefore abstract from this dimension of heterogeneity.

The number of vacancies varies with the business cycles, giving workers a varying bargaining power. Hence the bargaining power of workers is a time-varying endogenous variable to the system. Wage differences therefore account for the macroeconomic conditions under which workers entered the labor market. This captures a well documented that workers entering the labor force during a recession tend to do worse than others in terms of lifetime income.

Besides bargaining power considerations, the return on a worker is not only the result of the technology determined productivity level of the firm. It is also the result of how successful the firm is at selling products. Some firms become much more successful than others with higher sales, accumulate higher equity, more financial capacity and therefore can afford to pay higher salaries to make sure no vacancies are unfilled. Let's take the example of two identical firms with the same number of workers, the same production function and same productivity levels. Suppose one firm manages to sell all its inventories while the other does not. One firm will have generated positive profits, and will expect higher demand in the future. The other one will have generated negative profits, will get closer to the financial constraint and expect lower demand. The first firm will value new workers more than the second because it expects to sell the goods it will produce with the new workers. The second will not value new workers as much, and might even have to reduce its labor force. Hence wage offers will differ even though the firms have similar production capacities and workers are identical.

Correlation across Business Cycle Variables. Another way to characterize the impact of θ on business cycle characteristics is to look at the correlation matrix for aggregate variables. the properties of the business cycles evolve dramatically. As shown in figure 3.15, the correlation structure of the economy weakens (red) when the share of inductive firms in the economy increases. This is consistent with the previous findings that increasing the share of inductive firms reduces the persistence of business cycles, making all variables more like white noise.

Figure 3.15. Correlation Matrix for Aggregate Variables

(a) $\theta = 0\%$

(b) $\theta = 50\%$

(c) $\theta = 75\%$

(d) $\theta = 100\%$

Notes: Each (i, j) cell of the matrix represents the correlation between variables i and j across simulations for a given value of θ. The values are expressed relative to the baseline calibration ($\theta = 21\%$). Negative (red) values imply that the correlation is lower than in the baseline calibration, while positive (blue) values imply that the correlation is higher than in the baseline calibration.

Explaining the reduction in aggregate level persistence. One possible explanation for the reduction in persistence caused by the increase in the share of inductive firms is that introducing memory at the firm level could make discrepancies between supply and demand less dependent on past observation, since it has been taken into account by firms. The same way that algorithmic reduces the ability to forecast stock prices using past observations, inductive firms could make the economy less dependent on past shocks.

To test this hypothesis, I run a second experiment where all parameters are held constant at the baseline calibration except for the memory of inductive firms. Figure 3.17 shows how the autocorrelation of sales decreases with the memory available to firms. Suggesting that this mechanism could be at play.

Figure 3.17. Auto-correlation of Aggregate Sales

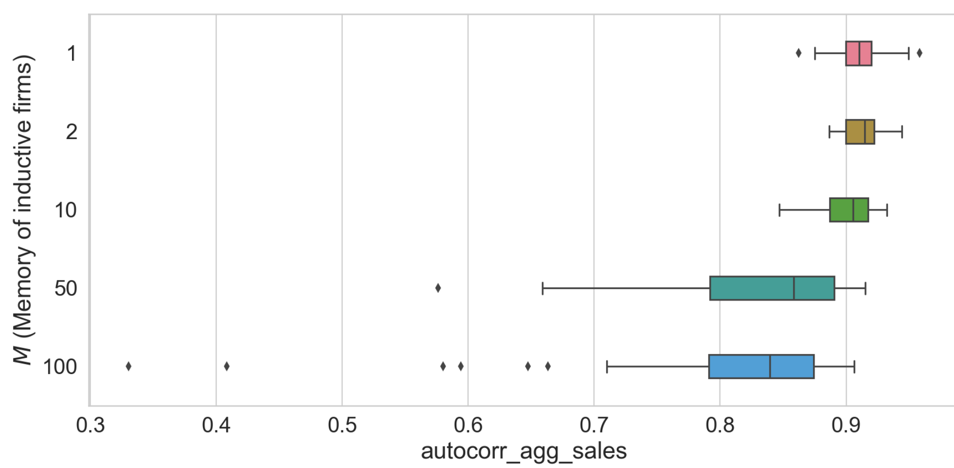

Notes: The statistic is computed over the 1000 periods in a given simulation. For each value of M, the boxplot depicts the distribution of that statistic across the 40 simulation runs.

Measures of passthrough. A different way to look at the simulations is to test for some relationships between variables and observe how they differ with variations in inductive reasoning parameters (θ, M, J).

A key result is that the share of inductive firms can qualitatively affect the relationship between aggregate variables in the model. For instance, the Phillips curve between inflation and output growth, although positively sloped with naive firms, turns negatively slopped when the share of inductive firms is high. This is shown in table 3.4.

Table 3.4. Regression results for outcome inflation varying θ

	inflation					
θ	0	21	25	50	75	100
	(1)	(2)	(3)	(4)	(5)	(6)
gdpgrowth	0.05***	0.02***	0.01	-0.06***	-0.11***	-0.16***
	(0.003)	(0.005)	(0.008)	(0.01)	(0.01)	(0.01)
Observations	39,880	40,040	40,040	40,040	40,040	40,040
R^2	0.07186	0.80135	0.43919	0.26850	0.63163	0.83493
Within R^2	0.02717	0.00551	0.00046	0.01408	0.09170	0.06191
simulation fixed effects	✓	✓	✓	✓	✓	✓
periods fixed effects	✓	✓	✓	✓	✓	✓

Notes: Panel regression estimates using fixest OLS with fixed effects for the simulation index (between 1 and 40) and fixed effects for each period (1 to 1001). This way estimations are based on a 40×10001 *panel regression sample.*

Similarly the Beveridge curve only displays the expected negative sign between vacancy rate and unemployment when the share of inductive firms is high enough (table 3.5)). Additionally the relationship between wage inflation and output growth is not robust to the share of inductive firms in the economy (table 3.6).

Table 3.5. Regression results for outcome vacancy_rate varying θ

	vacancy_rate					
θ	0	21	25	50	75	100
	(1)	(2)	(3)	(4)	(5)	(6)
unemployment_rate	0.18***	-0.31***	-0.30***	-0.57***	-0.71***	-0.34***
	(0.06)	(0.06)	(0.09)	(0.05)	(0.02)	(0.03)
Observations	40,040	40,040	40,040	40,040	40,040	40,040
R^2	0.12054	0.23243	0.27921	0.22837	0.16872	0.04136
Within R^2	0.00963	0.01396	0.01143	0.02237	0.02587	0.00819
simulation fixed effects	✓	✓	✓	✓	✓	✓
periods fixed effects	✓	✓	✓	✓	✓	✓

Table 3.6. Regression results for outcome wage_inflation varying θ

	wage_inflation					
θ	0	21	25	50	75	100
	(1)	(2)	(3)	(4)	(5)	(6)
gdpgrowth	0.04***	0.03***	0.03***	0.0008	0.010***	0.01
	(0.002)	(0.003)	(0.007)	(0.009)	(0.002)	(0.009)
Observations	39,880	40,040	40,040	40,040	40,040	40,040
R^2	0.11402	0.93097	0.31607	0.02102	0.30350	0.91715
Within R^2	0.04871	0.02430	0.00539	3.19×10^{-6}	0.00183	0.00049
simulation fixed effects	✓	✓	✓	✓	✓	✓
periods fixed effects	✓	✓	✓	✓	✓	✓

Notes: Panel regression estimates using fixest OLS with fixed effects for the simulation index (between 1 and 40) and fixed effects for each period (1 to 1001). This way estimations are based on a 40×10001 *panel regression sample.*

One relationship is robust to changes in the share of inductive firms: Okun's law. The relationship between output growth and unemployment change is always negative, regardless of the share of inductive firms (table 3.7), although the magnitude of the passthrough from output to unemployment is three times as strong in presence of inductive firms.

Table 3.7. Regression results for outcome unemployment_change varying θ

	unemployment_change					
θ	0	21	25	50	75	100
	(1)	(2)	(3)	(4)	(5)	(6)
gdpgrowth	-0.48***	-2.1***	-2.1***	-2.1***	-1.8***	-2.0***
	(0.14)	(0.10)	(0.09)	(0.07)	(0.10)	(0.12)
Observations	39,880	40,040	40,040	40,040	40,040	40,040
R^2	0.03274	0.10493	0.09363	0.07697	0.04739	0.03669
Within R^2	0.00500	0.07624	0.06903	0.05330	0.02194	0.01099
simulation fixed effects	✓	✓	✓	✓	✓	✓
periods fixed effects	✓	✓	✓	✓	✓	✓

Notes: Panel regression estimates using fixest OLS with fixed effects for the simulation index (between 1 and 40) and fixed effects for each period (1 to 1001). This way estimations are based on a 40×10001 *panel regression sample.*

3.6. Conclusion

A simple agent-based model with limited ex ante heterogeneity and no aggregate shocks is able to match key business cycle like behaviors in the aggregate, while providing a rich heterogeneity in the cross-section. Agents do not have a preference for fluctuations. Yet, local coordination failures in the goods and labor market accumulate enough to make fluctuations emerge in the aggregate.

A key finding is the realization that business cycle characteristics are fragile to firms' forecasting processes, because firms production decision conditional on their forecasts are critical on both goods and labor markets, and therefore aggregate outcomes.

In particular, the share of inductive firms in the economy can completely flip the sign of traditional economic relationships like the Phillips curve. This could be part of the explanation for apparently varying empirical estimates for the Phillips curve in the US economy, as firms have improved their forecasting process with the advent of new technologies.

The model also provides a rich cross-sectional distribution of firms' characteristics. In particular, it is able to generate right skewed firm level sales and firms' size. Cross-sectional heterogeneity is also significantly affected by firms' forecasting process. Prices and wages tend to converge to a unique value when inductive reasoning is more elaborate and common across firms.

Computational resources limited the scale and realism of the model. A possible extension of the paper would be to introduce a more realistic production function, with capital and intermediates inputs supplied in a production network. It is very well possible that introducing capital accumulation, the memory available to inductive firms might be of better use. I leave the study of these features to future researchers.

Bibliography

Acemoglu, Daron and Pablo D. Azar (2020), "Endogenous production networks." *Econometrica*, 88(1), 33–82.

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2012a), "The network origins of aggregate fluctuations." *Econometrica*, 80(5), 1977–2016.

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2012b), "The network origins of aggregate fluctuations." *Econometrica*, 80(5), 1977–2016.

Acemoglu, Daron and Alireza Tahbaz-Salehi (2020), "Firms, failures, and fluctuations: The macroeconomics of supply chain disruptions." NBER Working Paper No. 27565.

Alchian, A. (1950), "Uncertainty, Evolution, and Economic Theory." *Journal of Political Economy*, 58(3), 211–221.

Angeletos, George-Marios and Zhen Huo (2021), "Myopia and anchoring." *American Economic Review*, 111(4), 1166–1200.

Angeletos, George-Marios, Luigi Iovino, and Jennifer La'O (2016), "Real rigidity, nominal rigidity, and the social value of information." *American Economic Review*, 106(1), 200–227.

Angeletos, George-Marios and Jennifer La'O (2020), "Optimal monetary policy with informational frictions." *Journal of Political Economy*, 128(3), 1027–1064.

Angeletos, George-Marios and Chen Lian (2016), "Incomplete information in macroeconomics: Accommodating frictions in coordination." In *Handbook of Macroeconomics* (John B. Taylor and Harald Uhlig, eds.), volume 2, 1065–1240, Elsevier.

Angeletos, George-Marios and Chen Lian (2018), "Forward guidance without common knowledge." *American Economic Review*, 108(9), 2477–2512.

Angeletos, George-Marios and Alessandro Pavan (2007), "Efficient use of information and social value of information." *Econometrica*, 75(4), 1103–1142.

Arthur, W. Brian (1994), "Inductive reasoning and bounded rationality." *The American Economic R*, 84(2), 406–411.

Autio, Jaakko (1996), "Korot Suomessa 1862-1 952." *Bank of Finland Discussion Papers.*

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen (2020), "The Fall of the Labor Share and the Rise of Superstar Firms*." *The Quarterly Journal of Economics*, 135(2), 645–709, URL `https://doi.org/10.1093/qje/qjaa004`.

Axtell, Robert L. (2006), "Multi-agent systems macro: A prospectus." *Post Walrasian Macroeconomics: Beyond the Dynamic Stochastic General Equilibrium Model*, 203–220.

Axtell, Robert L and J Doyne Farmer (2023), "Agent-Based Modeling in Economics and Finance: Past, Present, and Future." *Journal of Economic Literature*, Forthcoming.

Baqaee, David and Emmanuel Farhi (2019), "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem." *Econometrica*, 87(4), 1155–1203.

Baqaee, David and Emmanuel Farhi (2021), "Entry vs. rents: Aggregation with economies of scale." NBER Working Paper No. 27140.

Baqaee, David and Emmanuel Farhi (2022), "Supply and demand in disaggregated Keynesian economies with an application to the COVID-19 crisis." *American Economic Review*, 112(5), 1397–1436.

Baqaee, David and Elisa Rubbo (2022), "Micro propagation and macro aggregation." NBER Working Paper No. 30538.

Barrot, Jean-Noël and Julien Sauvagnat (2016), "Input specificity and the propagation of idiosyncratic shocks in production networks." *Quarterly Journal of Economics*, 131(3), 1543–1592.

Beaudry, Paul, Dana Galizia, and Franck Portier (2015), "Reviving the limit cycle view of macroeconomic fluctuations." *University of British Columbia mimeo*.

Beaudry, Paul, Dana Galizia, and Franck Portier (2020a), "Emergent Equilibrium and the role of Expectations in a Dynamic Macroeconomic Model with Weak Complementarities." (December).

Beaudry, Paul, Dana Galizia, and Franck Portier (2020b), "Putting the cycle back into business cycle analysis." *American Economic Review*, 110(1), 1–47.

Bergemann, Dirk, Tibor Heumann, and Stephen Morris (2017), "Information and interaction." Cowles Foundation Discussion Paper No. 2088.

Bernard, Andrew B., Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes (2022), "The origins of firm heterogeneity: A production network approach." *Journal of Political Economy*, 130(7), 1765–1804.

Bigio, Saki and Jennifer La'O (2020), "Distortions in production networks." *Quarterly Journal of Economics*, 135(4), 2187–2253.

Blanchard, Olivier J. (1983), "The Production and Inventory Behavior of the American Automobile Industry." *Journal of Political Economy*, 91(3), 365–400. Publisher: University of Chicago Press.

Boehm, Christoph, Aaron Flaaen, and Nitya Pandalai-Nayar (2019), "Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tōhoku earthquake." *Review of Economics and Statistics*, 101(2), 60–75.

Boser, B., I. Guyon, and V. Vapnik (1992), "A training algorithm for optimal margin classifiers." *Proceedings of the Fifth Annual Workshop on Computational Learning Theory, Pittsburgh*.

Bottazzi, Giulio and Angelo Secchi (2006), "Explaining the distribution of firm growth rates." *The RAND Journal of Economics*, 37(2), 235–256, URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1756-2171.2006.tb00014.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1756-2171.2006.tb00014.x.

Breiman, Leo (2001), "Random Forests." *Machine Learning*, 45, 5–32.

Bryzgalova, Svetlana, Markus Pelger, and Jason Zhu (2019), "Forest Through the Trees: Building Cross-Sections of Stock Returns." *SSRN Electronic Journal*, URL https://

www.ssrn.com/abstract=3493458.

Bui, Ha, Zhen Huo, Andrei A. Levchenko, and Nitya Pandalai-Nayar (2022), "Information frictions and news media in global value chains." NBER Working Paper No. 30033.

Bui, Ha, Zhen Huo, Andrei A. Levchenko, and Nitya Pandalai-Nayar (2023), "Information frictions and news media in global value chains." NBER Working Paper No. 30033.

Burdett, Kenneth and Kenneth L. Judd (1983), "Equilibrium price dispersion." *Econometrica*, 51(4), 955–969, URL http://www.jstor.org/stable/1912045.

Calvó-Armengol, Antoni, Joan de Martí, and Andrea Prat (2015), "Communication and influence." *Theoretical Economics*, 10(2), 649–690.

Carvalho, Vasco M. (2014), "From micro to macro via production networks." *The Journal of Economic Perspectives*, 28(4), 23–47.

Carvalho, Vasco M., Makoto Nirei, Yukiko Saito, and Alireza Tahbaz-Salehi (2021), "Supply chain disruptions: Evidence from the Great East Japan Earthquake." *Quarterly Journal of Economics*, 136(2), 1255–1321.

Carvalho, Vasco M. and Alireza Tahbaz-Salehi (2019), "Production networks: A primer." *Annual Review of Economics*, 11(1), 635–663.

Chen, Tianqi and Carlos Guestrin (2016), "XGBoost: A Scalable Tree Boosting System." *In Krishnapuram, Balaji; Shah, Mohak; Smola, Alexander J.; Aggarwal, Charu C.; Shen, Dou; Rastogi, Rajeev (eds.). Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, 785–794.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005), "Nominal rigidities and the dynamic effects of a shock to monetary policy." *Journal of Political Economy*, 113(1), 1–45.

Correia, Sergio, Paulo Guimar aes, and Tom Zylkin (2020), "PPMLHDFE: Fast Poisson Estimation with High-Dimensional Fixed Effects." *The Stata Journal*.

Danmarks Statistik (1969), "Kreditmarkedstatistik." *Statistik Undersogelser*.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger (2020), "The Rise of Market Power and the Macroeconomic Implications*." *The Quarterly Journal of Economics*, 135(2), 561–644, URL https://doi.org/10.1093/qje/qjz041.

de Martí, Joan and Yves Zenou (2015), "Network games with incomplete information." *Journal of Mathematical Economics*, 61, 221–240.

Dedinger, Beatrice and Paul Girard (2017), "Exploring Trade Globalization in the Long Run: The RICardo Project." *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 50(1), 30–48.

Delli Gatti, D., G. Fagiolo, Gallegati M., and A. Russo M., Richiardi (2019), *Agent-Based Models in Economics: A Toolkit*, volume 57.

Delli Gatti, Domenico, Saul Desiderio, Edoardo Gaffeo, Pasquale Cirillo, and Mauro Gallegati (2011), *Macroeconomics from the bottom-up*. Springer.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Toshiaki Komatsu, Magne Mogstad, and Felix Tintelnot (2022), "Foreign demand shocks to production networks: Firm responses and worker impacts." NBER Working Paper No. 30447.

Eichenbaum, Martin (1989), "Some empirical evidence on the production level and production cost smoothing models of inventory investment." *The American Economic R*, 79(4), 853–864. Publisher: American Economic Association.

Einaudi, Luca L (2000), “From the Franc to the ’Europe’: The Attempted Transformation of the Latin Monetary Union into a European Monetary Union, 1865-1873.” *The Economic History Review*, 53(2), 284–308.

Einaudi, Luca L (2001), “Money and Politics. European Monetary Unification and the Internatiobnal Gold Standard (1865-1873).” *Oxford: Oxford University Press.*

Einaudi, Luca L (2007), *The Experience of Exchange Rate Regimes in Southeastern Europe in a Historical and Comparative Perspective, Workshop Proceedings of OeNB Workshop N13. Second Conference of the South-Eastern European Monetary History Network, April.*

Federal Reserve Bank of Richmond and Nicholas Trachter (2017), “Price Dispersion When Stores Sell Multiple Goods.” *Economic Quarterly*, 102(02), 127–146.

Ferrari, Alessandro (2023), “Inventories, demand shocks propagation and amplification in supply chains.” Working paper.

Fisher, Irving (1933), “The Debt-Deflation Theory of Great Depressions.” *Econometrica*, 1(4), 337–357.

Flandreau, Marc (2000), “The Economics and Politics of Monetary Unions: A Reassessment of the Latin Monetary Union, 1865-71.” *Financial History Review*, 25–43.

Flandreau, Marc and K Oosterlinck (2012), “Was the Emergence of the International Gold Standard Expected? Evidence from Indian Government Securities.” *Journal of Monetary Economics*, 59, 649–669.

Fouquin, Michel and Jules Hugot (2016), “Two Centuries of Bilateral Trade and Gravity data: 1827- 2012.” *CEPII Working paper*, 14.

Freund, Yoav and Robert E. Schapire (1999), “A Short Introduction to Boosting.” *Journal of Japanese Society for Artificial Intelligence*, 14(5), 771–780.

Fuster, Andreas, David Laibson, and Brock Mendel (2010), “Natural Expectations and Macroeconomic Fluctuations.” *Journal of Economic Perspectives*, 24(4), 67–84, URL https://pubs.aeaweb.org/doi/10.1257/jep.24.4.67.

Gabaix, Xavier (2011), “The Granular Origins of Aggregate Fluctuations.” *Econometrica*, 79(3), 733–772.

Gartner, Inc. (2022), “CMO Spend Survey 2022.” https://www.gartner.com/en/articles/what-marketing-budgets-look-like-in-2022.

Geanakoplos, John, Robert Axtell, J. Doyne Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan M. Palmer, and Chun-Yi Yang (2012), “Getting at systemic risk via an agent-based model of the housing market.” *American Economic Review*, 102(3), 53–58.

Geurts, Pierre, Damien Ernst, and Louis Wehenkel (2006), “Extremely randomized trees.” *Machine Learning*, 63, 3–42.

Gigerenzer, Gerd and Henry Brighton (2009), “Homo Heuristicus: Why Biased Minds Make Better Inferences.” *Topics in Cognitive Science*, 1(1), 107–143, URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1756-8765.2008.01006.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1756-8765.2008.01006.x.

Golub, Benjamin and Stephen Morris (2018), “Expectations, networks, and conventions.” Working paper.

Goodwin, R. M. (1951), “The nonlinear accelerator and the persistence of business cycles.” *Econometrica*, 19(1), 1–17.

Grossman, Gene M., Elhanan Helpman, and Hugo Lhuillier (2023), "Supply chain resilience: Should policy promote diversification or reshoring?" NBER Working Paper No. 29330.

Gualdi, Stanislao, Marco Tarzia, Francesco Zamponi, and Jean-Philippe Bouchaud (2015), "Monetary Policy and Dark Corners in a Stylized Agent-Based Model." *SSRN Electronic Journal*, 1–21.

Guesnerie, Roger and Michael Woodford (1993), "Endogenous fluctuations." In *Advances in Economic Theory* (Jean-Jacques Laffont, ed.), volume 2, 289–412, Cambridge University Press.

Helleiner (2003), *The Making of National Money: Territorial Currencies in Historical Perspective*. New York: Cornell University Press.

Helper, Susan and Evan Soltas (2021), "Why the pandemic has disrupted supply chains." White House Council of Economic Advisors, https://www.whitehouse.gov/cea/written-materials/2021/06/17/why-the-pandemic-has-disrupted-supply-chains/.

Hoerl, Arthur E. and Robert W. Kennard (2000), "Ridge Regression: Biased Estimation for Nonorthogonal Problems." *Technometrics*, 42, 80–86.

Horvath, Michael (2000), "Sectoral shocks and aggregate fluctuations." *Journal of Monetary Economics*, 45(1), 69 – 106.

Institute for Supply Management (2023a), "Manufacturing ISM Report On Business." https://www.ismworld.org/supply-management-news-and-reports/reports/ism-report-on-business/pmi/january/.

Institute for Supply Management (2023b), "Manufacturing ISM Report On Business." https://www.ismworld.org/supply-management-news-and-reports/reports/ism-report-on-business/pmi/march/.

Jasova, Martina, Luc Laeven, Caterina Mendicino, Jose-Luis Peydro, and Dominik Supera (2021), "Systemic risk and monetary policy: The haircut gap channel of the lender of last resort." *SSRN Electronic Journal*.

Kaldor, Nicholas (1940), "A model of the trade cycle." *The Economic Journal*, 50(197), 78–92.

Kalecki, M. (1935), "A macrodynamic theory of business cycles." *Econometrica*, 3(3), 327–344.

Kaniel, Ron, Zihan Lin, Markus Pelger, and Stijn Van Nieuwerburgh (2021), "Machine-Learning the Skill of Mutual Fund Managers." *SSRN Electronic Journal*, URL https://www.ssrn.com/abstract=3977883.

Khan, Aubhik and Julia K Thomas (2007), "Inventories and the business cycle: An equilibrium analysis of (s, s) policies." *The American Economic Review*, 97(4), 28.

Knight, F. (1921), *Risk, Uncertainty and Profit*.

Kopytov, Alexandr, Kristoffer Nimark, Bineet Mishra, and Mathieu Taschereau-Dumouchel (2022a), "Endogenous production networks under supply chain uncertainty." Working paper.

Kopytov, Alexandr, Kristoffer Nimark, Bineet Mishra, and Mathieu Taschereau-Dumouchel (2022b), "Endogenous production networks under supply chain uncertainty." Working paper.

Korovkin, V and A Makarin (2022), "Production networks and war: Evidence from Ukraine." CEPR Press Discussion Paper No. 16759.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh (2020), “Shrinking the cross-section.” *Journal of Financial Economics*, 135(2), 271–292, URL https://linkinghub.elsevier.com/retrieve/pii/S0304405X19301655.

Kryvtsov, O. and V. Midrigan (2013), “Inventories, Markups, and Real Rigidities in Menu Cost Models.” *The Review of Economic Studies*, 80(1), 249–276. Publisher: Oxford University Press.

Kydland, Finn E and Edward C Prescott (1982), “Time to Build and Aggregate Fluctuations.” *Source: Econometrica*, 50(6), 1345–1370.

La’O, Jennifer and Alireza Tahbaz-Salehi (2022), “Optimal monetary policy in production networks.” *Econometrica*, 90(3), 1295–1336.

Larsen-Hallock, Eugene, Adam Rej, and David Thesmar (2022), “Expectations Formation with Fat-tailed Processes: Evidence from Sales Forecasts.”

Lian, Chen and Yueran Ma (2021), “Anatomy of Corporate Borrowing Constraints.” *The Quarterly Journal of Economics*, 136(1), 229–291.

Liu, Ernest and Aleh Tsyvinski (2023), “Dynamical structure and spectral properties of input-output networks.” NBER Working Paper No. 28178.

Long, John B. Jr and Charles I. Plosser (1983), “Real Business Cycles.” *Journal of Political Economy*, 91(1), 39–69.

Lucas, Robert E. and Edward C. Prescott (1971), “Investment under uncertainty.” *Econometrica*, 39(5), 659–681.

Maćkowiak, Bartosz and Mirko Wiederholt (2009), “Optimal sticky prices under rational inattention.” *American Economic Review*, 99(3), 769–803.

Mankiw, N. Gregory and Ricardo Reis (2002), “Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips Curve.” *Quarterly Journal of Economics*, 117(4), 1295–1328.

Meissner, C (2015), “A New World Order: Explaining the International Diffusion of the Gold Standard, 1870-1913.” *Journal of International Economics*, 66, 385–406.

Pellet, Thomas and Alireza Tahbaz-Salehi (2023), “Rigid production networks.” *Journal of Monetary Economics, forthcoming*.

Reinhart, Carmen M., Vincent Reinhart, and Christoph Trebesch (2016), “Global Cycles: Capital Flows, Commodities, and Sovereign Defaults, 1815–2015.” *American Economic Review*, 106(5), 574–580, URL https://pubs.aeaweb.org/doi/10.1257/aer.p20161014.

Roughgarden, Tim (2010), “Computing Equilibria: A Computational Complexity Perspective.” *Economic Theory*, 42(1), 193–236, URL https://www.jstor.org/stable/25619983. Publisher: Springer.

Rubbo, Elisa (2022), “Networks, Phillips curves, and monetary policy.” Working paper.

Schelling, T. (1971), “Models of segregation.” *American Economic Review*, 59(2).

Smets, Frank and Rafael Wouters (2007), “Shocks and frictions in us business cycles: A bayesian dsge approach.” *American Economic Review*, 97(3), 586–606.

Tibshirani, R (1996), “Regression Shrinkage and Selection via the Lasso.” *Journal of the Royal Statistical Society*, 58, 267–288.

Timini, Jacopo (2018), “Currency Unions and Heterogeneous Trade Effects: The Case of the Latin Monetary Union.” *European Review of Economic History*, 22, 322–348.

Townsend, Robert M (1983), “Forecasting the forecasts of others.” *The Journal of Political Economy*, 546–588.

vom Lehn, Christian and Thomas Winberry (2022), "The investment network, sectoral comovement, and the changing U.S. business cycle." *Quarterly Journal of Economics*, 37, 387–433.

Willis, Henry Parker (1901), "A History of the Latin Monetary Union — A Study of International Monetary Action." *Chicago: University of Chicago Press.*

Woodford, Michael (2003), "Imperfect common knowledge and the effects of monetary policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps* (Philippe Aghion, Roman Frydman, Joseph E. Stiglitz, and Michael Woodford, eds.), chapter 1, 25–58, Princeton University Press.

Woodford, Michael (2013), "Macroeconomic analysis without the rational expectations hypothesis." *Annual Review of Economics*, 5, 303–346.

Yotov, Yoto V., Roberta Piermartini, José-Antonio Monteiro, and Mario Larch (2016), "An Advanced Guide to Trade Policy Analysis: The Structural Gravity Model." *WTO Publications.*

APPENDIX A

Appendix to Chapter 1

A.1. LMU membership

The following table and map provide a summary of the countries that participated to the Latin Monetary Union, together with the time period during which they were part of it.

Table A.1. LMU Membership

Country	Condition	Date	Period
Belgium	LMU founding member	23 December, 1865	1865-1927
France	LMU founding member	23 December, 1865	1865-1927
Italy	LMU founding member	23 December, 1865	1865-1927
Switzerland	LMU founding member	23 December, 1865	1865-1927
Greece	LMU member	18 November, 1868	1867-1927
Algeria (French colony)	Shadowing	23 December, 1865	n.a.
Austria-Hungary	Shadowing	n.a.	1870-1914
Bulgaria	Shadowing	9 August, 1877	1878-1914
Peru	Shadowing	31 July, 1863	n.a.
Poland	Shadowing	1926	1926
Pontifical State	Shadowing	1866	1866-1870
Romania	Shadowing	1 January, 1868	1867-1914
Russia	Shadowing	n.a.	1886-1865
Serbia	Shadowing	11 November, 1878	187*-1914
Spain	Shadowing	19 October, 1868	1868-1914
Sweden	Shadowing	n.a.	1868-1872
Tunisia (French colony)	Shadowing	23 December, 1865	n.a.
Venezuela	Shadowing	11 May, 1871	n.a.

Notes: This table is taken from Appendix II in Timini (2018), and is here reported for simplicity. The sources of the table are Willis (1901); Einaudi (2007); Helleiner (2003).

Figure A.1. LMU membership by year of accession (1880 administrative boundaries)

A.2. Tradehist data

Table A.2. Variables from Tradehist

Variable	Dimension	Description
iso	country	Origin (destination) country
year	year	Year
FLOW	country-pair-direction-year	Bilateral trade flow
GDP_o(d)	country-year	GDP of the country
SH_PRIM_o(d)	country-year	Share of primary sector in the country's GDP
SH_SECD_o(d)	country-year	Share of secondary sector in the country's GDP
IPTOT_o(d)	country-year	Total imports
XPTOT_o(d)	country-year	Total exports
BITARIFF	country-pair-direction-year	Tariff imposed by country d on imports from country o
TARIFF_o(d)	country-year	Average tariff imposed by country o(d)
Distw	country pair	Population-weighted-great-circle distance
Dist_coord	country pair	Great-circle distance between main cities
Dist_o(d)	country	Internal distance of the origin (destination) country
SeaDist_SHRT	country-pair-year	Shortest bilateral sea distance
SeaDist_2CST	country-pair-year	Shortest bilateral sea distance
Comlang	country-pair	=1 if at least one language is spoken by more than 9% of the population in both countries
Contig	country-pair	=1 if the countries are contiguous
Curcol	country-pair-year	=1 if the origin and the dest. are in a colonial relationship
Curcol_o(d)	country-year	=1 if the country is a colony
Evercol	country pair	=1 if countries ever were in a colonial relationship
XCH_RATE_o(d)	country-year	British pounds per local currency unit
POP_o(d)	country-year	Population of the country
CONTI_o(d)	country	Continent of the country
REGIO_o(d)	country	Sub-continental region of the country
OECD_o(d)	country-year	=1 if the country belongs to the OECD
EU_o(d)	country-year	=1 if the country belongs to the E.U.
GATT_o(d)	country-year	=1 if the country belongs to the GATT/WTO

Notes: The description of the variables follows Fouquin and Hugot (2016).

A.3. CPIS Statistics

Table A.3. CPIS Statistics

	Countries	Observations	FF (Mean)	FF (StD)
Total	93	258459	2650.25$	108148.27$
Advanced Economies	31	16765	64800.65$	188875.05$
Non-Advanced Economies	62	121091	575.75$	47257.14$
Advanced/Non-Advanced		120603	4179.84$	157339.48$
Timini	15	4368	6766.55$	16927.48$

Notes: FF stands for Financial Flows. The rows "Advanced Economies" and "Non-Advanced Economies" report value where bilateral financial flows involve only advanced or non-advanced economies, respectively. The row "Advanced/Non-Advanced" reports value for bilateral financial flows among advanced and non-advanced entities. The row "Timini" reports values for bilateral financial flows the subsection of countries considered in Timini (2018).

A.4. Long-term Interest Rates

Since the Tradehist dataset does not contain many financial variables, we supplement it with long-term interest rate data assembled using different sources. The tables below provide summary statics for our reconstructed variable, and a a description of the sources used.

Table A.4. Long-Run Interest Rate Series: Statistics

Country	Mean	StD
Austria-Hungary	5.65%	2.46%
Belgium	4.81%	2.43%
Denmark	5.62%	3.65%
Finland	5.50%	1.30%
France	4.97%	2.79%
Germany	4.81%	2.11%
Greece	9.45%	4.86%
Italy	6.40%	3.51%
Netherlands	4.37%	2.08%
Norway	5.05%	2.58%
Portugal	6.38%	4.12%
Spain	7.09%	4.24%
Sweden	5.00%	2.73%
Switzerland	3.88%	1.22%
United Kingdom	4.87%	3.18%

Table A.5. Long-Run Interest Rate Series: Sources

Country	Source	Series
Austria-Hungary	GFD	10y government bond yield (close), 1861-2017
Belgium	GFD	10y government bond yield (close), 1861-2017
Denmark	DS & GFD	DS: *Kursog rentetabeler for obligationsmarkedet, Tabel 6*
		GFD: 10y government bond yield (close), 1861-2017
Finland	Autio & JST	Autio: *Liite 1, Oblig. Tuotto* 1863-1869
		JST: Long-term rates 1870-2017
France	GFD	10y government bond yield (close), 1861-2017
Germany	GFD	10y government bond yield (close), 1861-2017
Greece	GFD & GCB	GFD: Mortgage lending rate (close) 1861-1941, 2003-2013;
		GCB: Long-term loans by commercial banks 1951-2002
Italy	GFD	10y government bond yield (close), 1861-2017
Netherlands	GFD	10y government bond yield (close), 1861-2017
Norway	GFD	10y government bond yield (close), 1861-2017
Portugal	GFD	10y government bond yield (close), 1861-2017
Spain	GFD	10y government bond yield (close), 1861-2017
Sweden	GFD	10y government bond yield (close), 1861-2017
Switzerland	SNB & JST	SNB: mortgage rates 1861-1880
		JST: Long-term rates 1881-2017
United Kingdom	GFD	10y government bond yield (close), 1861-2017

Notes: GFD stands for Global Financial Data, available at https://globalfinancialdata.com. JST stands for the Jordà-Schularick-Taylor Macrohistory Database, available at https://www.macrohistory.net/database/. For Finland, Autio refers to Autio (1996). For Greece, GCB stands for the Greek Central Bank, whose historical interest rate data is available at https://www.bankofgreece.gr/en/statistics/financial-markets-and-interest-rates/bank-deposit-and-loan-interest-rates. For Switzerland, SNB stands for the Swiss National Bank, whose historical interest rate data is available at: https://www.snb.ch/en/iabout/stat/statrep/statpubdis/id/statpub histz arch#t2. For Denmark, DS stands for Danmarks Statistik (1969), available at https://www.dst.dk/Site/Dst/Udgivelser/GetPubFile.aspx?id=19918&sid=kreditm.

A.5. Description of ML Methodologies Used

Our goal is to reconstruct bilateral financial flows during the second half of the 19^{th} century as accurately as possible. In order to achieve this goal, we rely on several machine learning techniques, which have been developed precisely to obtain high performance forecasts. In this section, we briefly summarize the characteristics of the methods we use in our analysis[1].

Lasso and Ridge. The first two methods we use are those of standard Lasso and Ridge regressions Tibshirani (1996); Hoerl and Kennard (2000). These are well known penalized regression methods whose prediction accuracy, when the set of regressors is large relative to the amount of available observations, is enhanced through variable selection (in the case of Lasso) or variable shrinkage (in the case of Ridge). In both cases, the goal is to increase out-of-sample prediction accuracy by limiting the in-sample fit of the model.

Support Vector Machine. Moving away from linear methods, the Support Vector Machine algorithm can implement non-linear regression analyses Boser et al. (1992) and achieve higher prediction accuracy. The idea behind this method is to classify the training data by creating hyperplanes in a high-dimensional space, which are then used to predict observations out-of-sample in a flexible way.

Random Forest and Extra Trees. Both the Random Forest algorithm (Breiman, 2001) and the Extra Trees algorithm (Geurts et al., 2006) consist in creating several independent regression trees, and then averaging across their predictions. Each regression tree implements a classification of the data through recursive binary partitions of it. The difference between the two methods relies on the fact that, in Extra Trees, each tree is

[1]This is in no way a detailed description of the algorithms we are using but, rather, an intuitive description of their main characteristics. We provide references to studies providing a more formal description of these methods.

trained using the whole sample while, in Random Forest, trees are trained on a random subset of the sample.

AdaBoost, LightGBM and XGBoost. Similar to Random Forest and Extra Trees, these methods also rely on averaging the results from independent regression trees (Freund and Schapire, 1999; Chen and Guestrin, 2016). Albeit with some minor differences in the way the algorithms are implemented, all three of them sequentially evaluate the performance of regression trees, and assign a weight to these based on the accuracy of their forecasts. Through this iterative procedure, the algorithms build a model as a weighted sum of the predictions of the independent trees, enhancing their individual forecasting ability. The main difference across the algorithms is indeed linked to the way in which the weighting is implemented.

Neural Networks. Multi-layer Perceptrons (MLP) regressors are function approximators characterized by hidden layers of basis functions stacked on top of each other between an input layer and the output layer. Each layer is composed of neurons, which are weighted linear summations of the output of previous layer's neurons plus a non-linear activation function. We use up to 4 hidden layers and 100 neurons per layer in the cross-validation step of the algorithm.

Table A.6 below provides a summary of the main pros and cons of the ML methods we use.

Table A.6. Characteristics of ML Models

Method	Category	Pros	Cons
Lasso	Regularization Algorithm	Model selection	Linear
Ridge	Regularization Algorithm	Model shrinkage	Linear
Support Vector Machine	Instance-based Algorithm	Memory-efficient classification	Unsuited for very large datasets
Random Forest	Ensemble Algorithm	Effective large data handling	Expensive cross-validation
Extra Trees	Ensemble Algorithm	Faster than Random Forest	Expensive cross-validation
AdaBoost	Ensemble Algorithm	Low overfit	Sensible to noise
XGBoost	Ensemble Algorithm	High-accuracy	Overfitting
LightGBM	Ensemble Algorithm	Faster than XGBoost	Overfitting
Neural Networks	Artificial Neural Network	High-accuracy	Difficult interpretability

A.6. Graphic Representation of Models' Performance

Figure A.2. Performance on CPIS (In-sample)

A. XGBoost

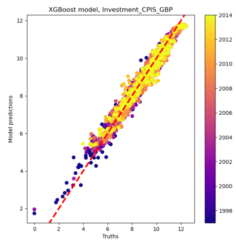

B. LGBM

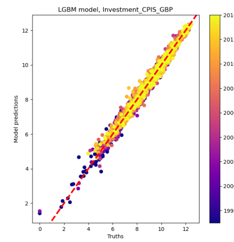

C. Extra Trees

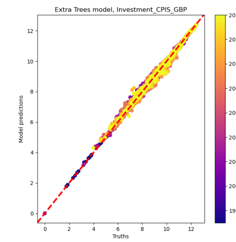

D. Lasso

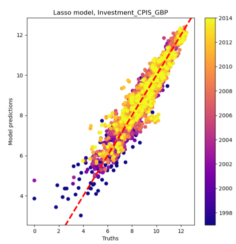

E. SVM

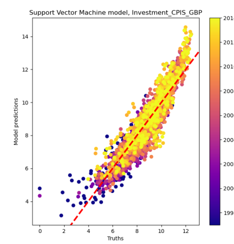

F. Ridge

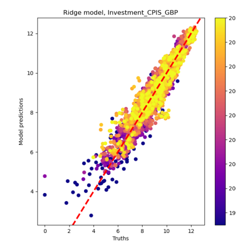

G. Random Forest

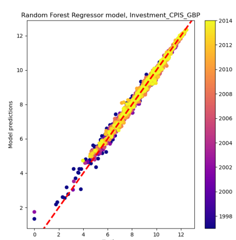

H. Neural Network

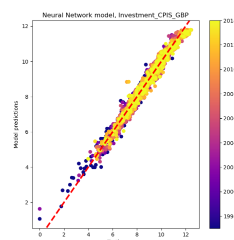

I. AdaBoost

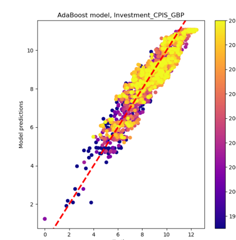

Notes: Axes are in log-scale.

Figure A.3. Performance on Trade Flows (In-sample)

A. XGBoost

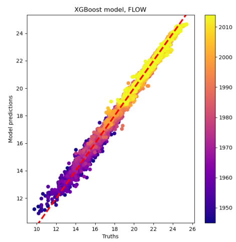

B. LGBM

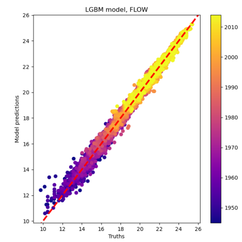

C. Extra Trees

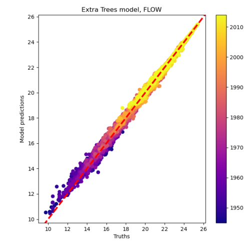

D. Lasso

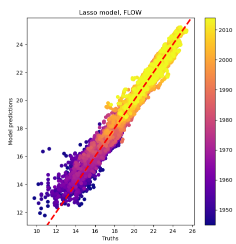

E. SVM

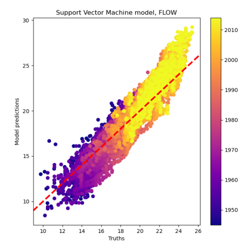

F. Ridge

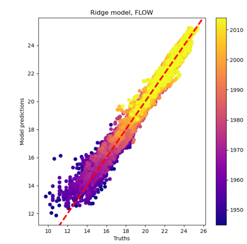

G. Random Forest

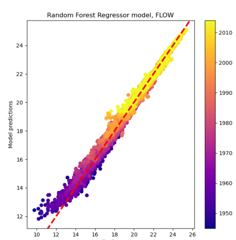

H. Neural Network

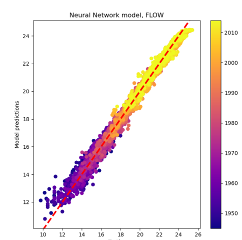

I. AdaBoost

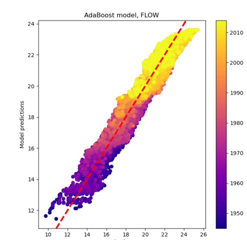

Notes: Axes are in log-scale.

Figure A.4. Performance on Trade Flows (Out-sample)

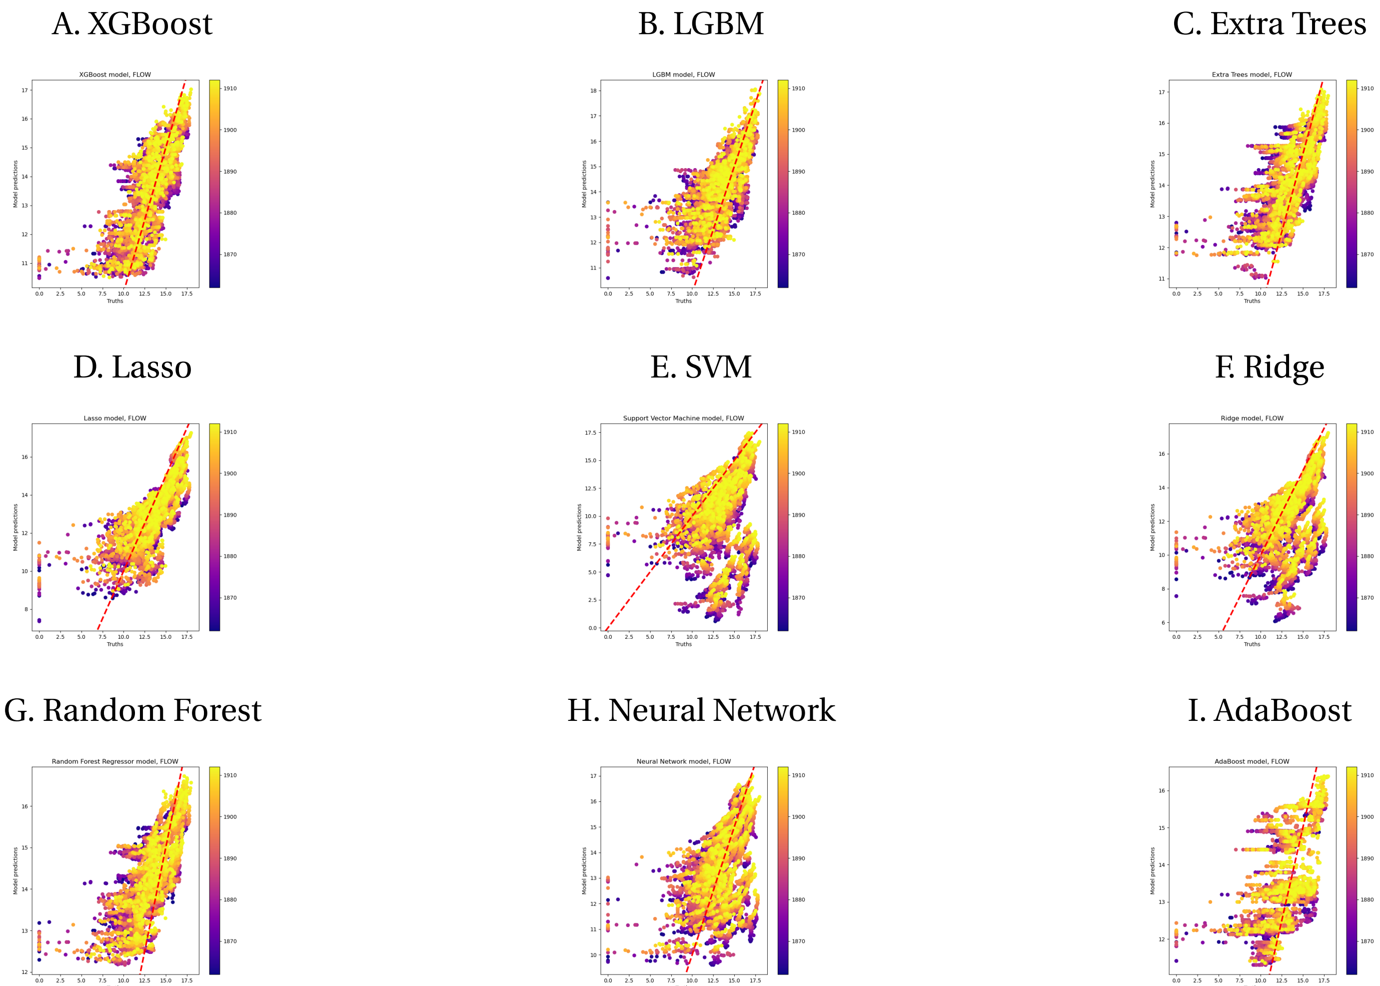

Notes: Axes are in log-scale.

A.7. Main Predictors of Bilateral Financial Flows

The charts below display the 50 most important variables in the forecasting exercise of bilateral financial flows according to Lasso (figure A.5) and XGBoost (figure A.6). The variables are displayed with increasing importance. In both charts, it is possible to note that variables referring to trade flows (FLOW and FLOW_1, the trade flow lag value) are very important predictors.

Figure A.5. Ranking of Features in Lasso

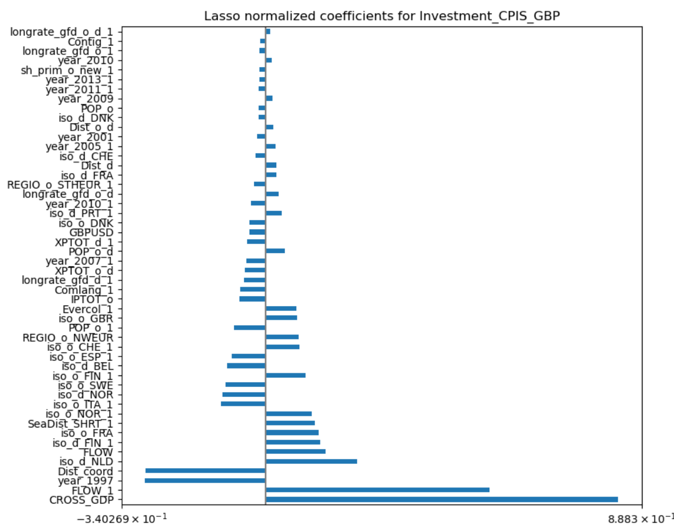

Figure A.6. Ranking of Features in XGBoost

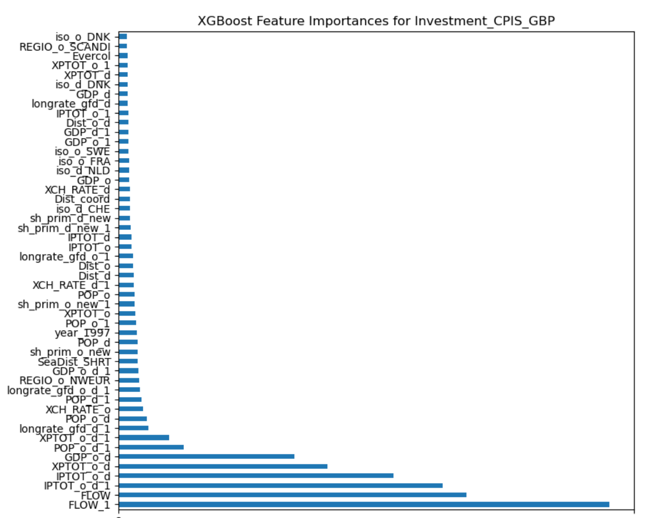

A.8. Results from Timini (2018)

Chart A.7 below is taken directly from Timini (2018), and is provided here to ease comparison with our results.

Figure A.7. Bilateral trade flows and Monetary agreements, 1861-1913

Table 2. *Bilateral trade flows and monetary agreements, 1861–1913*

	(1) LMU 1861–1913	(2) LMU 1861–1913	(3) LMU 1861–1885	(4) LMU 1861–1885	(5) LMU 1861–1874	(6) LMU 1861–1874
LMU	−0.127		−0.182*		−0.158*	
	(0.094)		(0.095)		(0.094)	
LMUFrance		−0.132		−0.209*		−0.147*
		(0.093)		(0.095)		(0.094)
LMURest		0.0817		0.129		0.0810
		(0.159)		(0.166)		(0.163)
LMU1885			0.155***			
			(0.0336)			
LMUFrance1885				0.167***		
				(0.0330)		
LMURest1885				−0.222***		
				(0.059)		
LMU1874					0.205***	
					(0.055)	
LMUFrance1874						0.205***
						(0.055)
LMURest1874						−0.105
						(0.093)
lnPOP	1.665***	1.665***	1.664***	1.654***	1.663***	1.658***
	(0.190)	(0.190)	(0.190)	(0.189)	(0.190)	(0.190)
SCU	−0.441***	−0.441***	−0.459***	−0.449***	−0.473***	−0.467***
	(0.092)	(0.092)	(0.093)	(0.093)	(0.094)	(0.094)
GS	0.295***	0.295***	0.259***	0.253***	0.262***	0.264***
	(0.040)	(0.040)	(0.039)	(0.039)	(0.039)	(0.039)
AllianceTreaty	−0.157***	−0.156***	−0.158***	−0.132***	−0.155***	−0.140***
	(0.025)	(0.025)	(0.024)	(0.025)	(0.025)	(0.025)
N	6,503	6,503	6,503	6,503	6,503	6,503

Notes: Poisson regressions. Dependent variable: Imports (value). All regressions include a constant, importer-year, exporter-year, and dyad fixed effects, not reported for the sake of simplicity. Robust standard errors in parentheses; $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.
Source: Author's elaboration.

A.9. Results Using Other Models

Table A.7. LMU dummy regression

	(1) Lasso	(2) XGBoost	(3) LGBM	(4) AdaBoost	(5) ET	(6) RF	(7) NN	(8) Ridge	(9) SVM
LMU	0.051*	0.046*	-0.017	0.027	-0.017*	-0.002	0.170*	0.028**	-0.061
	(0.021)	(0.019)	(0.039)	(0.037)	(0.008)	(0.019)	(0.077)	(0.010)	(0.099)
GS	0.248***	-0.027*	-0.019	0.037	0.009	-0.012	-0.184*	0.152*	0.409***
	(0.042)	(0.011)	(0.047)	(0.040)	(0.012)	(0.011)	(0.091)	(0.068)	(0.066)
SMU	-0.249***	-0.026	-0.134	-0.107*	-0.023	-0.074**	0.301*	-0.171***	-0.054
	(0.048)	(0.019)	(0.109)	(0.047)	(0.017)	(0.025)	(0.130)	(0.019)	(0.171)
N	7169	7169	7169	7169	7169	7169	7169	7169	7169

Notes: $^{***}p < 0.001, ^{**}p < 0.01, ^{*}p < 0.05$. Dependent variable: Estimated bilateral financial flows. All regressions include a constant, importer-year, exporter-year and importer-exporter fixed-effects. Clustered standard errors.

Table A.8. Comparing LMU effect on France and Rest of LMU

	(1) Lasso	(2) XGBoost	(3) LGBM	(4) AdaBoost	(5) ET	(6) RF	(7) NN	(8) Ridge	(9) SVM
LMU_France	0.047	0.005	-0.032	0.041	-0.016	-0.009	0.257*	0.014	0.225**
	(0.031)	(0.007)	(0.043)	(0.034)	(0.009)	(0.015)	(0.107)	(0.019)	(0.078)
LMU_Rest	0.087***	0.082**	0.007	0.007	-0.018	0.004	0.052	0.058**	-0.063
	(0.016)	(0.031)	(0.081)	(0.058)	(0.010)	(0.024)	(0.072)	(0.022)	(0.103)
GS	0.247***	-0.030**	-0.020	0.037	0.009	-0.013	-0.175	0.152*	0.409***
	(0.042)	(0.011)	(0.045)	(0.038)	(0.012)	(0.011)	(0.091)	(0.068)	(0.066)
SMU	-0.250***	-0.026	-0.134	-0.108*	-0.023	-0.074**	0.301*	-0.172***	-0.054
	(0.045)	(0.016)	(0.109)	(0.047)	(0.017)	(0.025)	(0.129)	(0.020)	(0.179)
N	7169	7169	7169	7169	7169	7169	7169	7169	7169

Notes: $^{***}p < 0.001, ^{**}p < 0.01, ^{*}p < 0.05$. Dependent variable: Estimated bilateral financial flows. All regressions include a constant, importer-year, exporter-year and importer-exporter fixed-effects. Clustered standard errors.

Table A.9. Comparing Effect of LMU Before and After 1885

	(1) Lasso	(2) XGBoost	(3) LGBM	(4) AdaBoost	(5) ET	(6) RF	(7) NN	(8) Ridge	(9) SVM
LMU	-0.049***	0.011	-0.028	0.035	-0.038***	-0.026	0.369***	-0.068***	-0.976***
	(0.014)	(0.019)	(0.052)	(0.048)	(0.011)	(0.022)	(0.090)	(0.017)	(0.188)
LMU_1885	0.204***	0.058**	0.019	-0.014	0.035*	0.041*	-0.305**	0.221***	1.083***
	(0.036)	(0.019)	(0.024)	(0.026)	(0.015)	(0.020)	(0.116)	(0.067)	(0.138)
GS	0.131**	-0.039**	-0.023	0.039	0.002	-0.021	-0.096	0.055	0.171*
	(0.047)	(0.012)	(0.045)	(0.036)	(0.012)	(0.012)	(0.064)	(0.068)	(0.071)
SMU	-0.256***	-0.020	-0.133	-0.109*	-0.019	-0.070**	0.259*	-0.190***	-0.030
	(0.015)	(0.017)	(0.110)	(0.048)	(0.016)	(0.023)	(0.131)	(0.026)	(0.161)
N	7169	7169	7169	7169	7169	7169	7169	7169	7169

Notes: $^{***}p < 0.001, ^{**}p < 0.01, ^{*}p < 0.05$. Dependent variable: Estimated bilateral financial flows. All regressions include a constant, importer-year, exporter-year and importer-exporter fixed-effects. Clustered standard errors.

Table A.10. Interacting France and 1885 dummies

	(1) Lasso	(2) XGBoost	(3) LGBM	(4) AdaBoost	(5) ET	(6) RF	(7) NN	(8) Ridge	(9) SVM
LMU_France	-0.059*	-0.032***	-0.038	0.049	-0.036***	-0.037*	0.528***	-0.097***	-0.186
	(0.026)	(0.010)	(0.054)	(0.044)	(0.009)	(0.019)	(0.141)	(0.019)	(0.106)
LMU_Rest	0.084***	0.052	-0.011	0.013	-0.040**	-0.015	0.219***	0.019	-0.980***
	(0.008)	(0.036)	(0.103)	(0.073)	(0.014)	(0.028)	(0.043)	(0.038)	(0.190)
LMU_France_1885	0.222***	0.065***	0.010	-0.015	0.035**	0.049**	-0.406***	0.271***	0.445***
	(0.045)	(0.012)	(0.024)	(0.023)	(0.012)	(0.017)	(0.105)	(0.055)	(0.083)
LMU_Rest_1885	-0.033	0.051	0.031	-0.011	0.036	0.033	-0.257*	0.071	1.086***
	(0.045)	(0.028)	(0.037)	(0.055)	(0.019)	(0.025)	(0.131)	(0.094)	(0.142)
GS	0.124**	-0.041***	-0.024	0.040	0.002	-0.021	-0.081	0.049	0.172*
	(0.047)	(0.012)	(0.043)	(0.034)	(0.012)	(0.012)	(0.067)	(0.069)	(0.073)
SMU	-0.252***	-0.019	-0.133	-0.109*	-0.019	-0.069**	0.250	-0.188***	-0.031
	(0.036)	(0.015)	(0.111)	(0.048)	(0.016)	(0.024)	(0.130)	(0.028)	(0.161)
N	7169	7169	7169	7169	7169	7169	7169	7169	7169

Notes: $^{***}p < 0.001,^{**}p < 0.01,^{*}p < 0.05$. Dependent variable: Estimated bilateral financial flows. All regressions include a constant, importer-year, exporter-year and importer-exporter fixed-effects. Clustered standard errors.

Table A.11. Comparing Effects of LMU Before and After 1874

	(1) Lasso	(2) XGBoost	(3) LGBM	(4) AdaBoost	(5) ET	(6) RF	(7) NN	(8) Ridge	(9) SVM
LMU	-0.003	0.025	-0.025	0.031	-0.029***	-0.020	0.314***	-0.035***	-0.315
	(0.008)	(0.019)	(0.044)	(0.038)	(0.008)	(0.020)	(0.092)	(0.007)	(0.162)
LMU_1874	0.147**	0.055*	0.021	-0.009	0.030*	0.047*	-0.369***	0.201***	0.684***
	(0.048)	(0.022)	(0.018)	(0.026)	(0.015)	(0.021)	(0.088)	(0.059)	(0.159)
GS	0.207***	-0.031**	-0.021	0.037	0.007	-0.015	-0.144	0.112	0.381***
	(0.049)	(0.012)	(0.046)	(0.038)	(0.012)	(0.010)	(0.089)	(0.067)	(0.053)
SMU	-0.256***	-0.021	-0.132	-0.108*	-0.020	-0.069**	0.245	-0.190***	-0.120
	(0.031)	(0.017)	(0.110)	(0.046)	(0.015)	(0.022)	(0.136)	(0.019)	(0.146)
N	7169	7169	7169	7169	7169	7169	7169	7169	7169

Notes: $^{***}p < 0.001,^{**}p < 0.01,^{*}p < 0.05$. Dependent variable: Estimated bilateral financial flows. All regressions include a constant, importer-year, exporter-year and importer-exporter fixed-effects. Clustered standard errors.

Table A.12. Interacting 1874 and France Dummies

	(1) Lasso	(2) XGBoost	(3) LGBM	(4) AdaBoost	(5) ET	(6) RF	(7) NN	(8) Ridge	(9) SVM
LMU_France	-0.008	-0.019	-0.035	0.054	-0.028***	-0.031*	0.472***	-0.054*	0.119
	(0.030)	(0.011)	(0.047)	(0.037)	(0.007)	(0.015)	(0.134)	(0.022)	(0.097)
LMU_Rest	0.097***	0.066*	-0.008	-0.014	-0.029*	-0.007	0.152***	0.037	-0.317
	(0.015)	(0.031)	(0.096)	(0.061)	(0.011)	(0.026)	(0.045)	(0.033)	(0.164)
LMU_France_1874	0.161**	0.064***	0.007	-0.040	0.034**	0.062***	-0.514***	0.251***	0.276***
	(0.055)	(0.015)	(0.023)	(0.027)	(0.012)	(0.017)	(0.118)	(0.059)	(0.071)
LMU_Rest_1874	-0.105	0.042	0.038	0.064	0.026	0.028	-0.264**	0.017	0.688***
	(0.066)	(0.036)	(0.053)	(0.054)	(0.023)	(0.026)	(0.087)	(0.087)	(0.169)
GS	0.203***	-0.033**	-0.020	0.039	0.007	-0.016	-0.126	0.109	0.382***
	(0.047)	(0.012)	(0.044)	(0.036)	(0.012)	(0.010)	(0.091)	(0.067)	(0.054)
SMU	-0.250***	-0.019	-0.134	-0.118*	-0.019	-0.067**	0.225	-0.186***	-0.121
	(0.035)	(0.014)	(0.110)	(0.052)	(0.015)	(0.022)	(0.133)	(0.028)	(0.149)
N	7169	7169	7169	7169	7169	7169	7169	7169	7169

Notes: $^{***}p < 0.001, ^{**}p < 0.01, ^{*}p < 0.05$. Dependent variable: Estimated bilateral financial flows. All regressions include a constant, importer-year, exporter-year and importer-exporter fixed-effects. Clustered standard errors.

APPENDIX B

Appendix to Chapter 2

This appendix contains the proofs and derivations omitted from the main body of the paper.

B.1. Proof of Proposition 1

We characterize the equilibrium via backward induction. Starting with the firms' decisions at $t = 1$, recall that firms optimally choose their labor input and flexible intermediate input quantities to meet the realized demand. Taking the prices, its realized demand, and their rigid input demands as given, firm k in industry i faces the following cost-minimization problem:

$$\min_{l_{ik},\{x_{ij,k}\}_{j\in\mathcal{F}_i}} \quad wl_{ik} + \sum_{j\in\mathcal{F}_i} p_j x_{ij,k}$$

$$\text{subject to} \quad y_{ik} = z_i \zeta_i l_{ik}^{\alpha_i} \prod_{j=1}^{n} x_{ij,k}^{a_{ij}}.$$

Solving this problem implies that the firm's expenditure on labor and flexible input demands are given by

$$wl_{ik} = \alpha_i (y_{ik}/Q_{ik})^{1/(1-\sum_{j\in\mathcal{R}_i} a_{ij})} \tag{B.1}$$

$$p_j x_{ij,k} = a_{ij} (y_{ik}/Q_{ik})^{1/(1-\sum_{j\in\mathcal{R}_i} a_{ij})} \qquad \text{for all } j \in \mathcal{F}_i \tag{B.2}$$

respectively, where Q_{ik} only depends on the firm's productivity, its input prices, the nominal wage, and the intermediate input decisions that are sunk by $t = 1$:

$$Q_{ik} = z_i w^{-\alpha_i} \prod_{j\in\mathcal{F}_i} p_j^{-a_{ij}} \prod_{j\in\mathcal{R}_i} (x_{ij,k}/a_{ij})^{a_{ij}}. \tag{B.3}$$

Therefore, the firm faces the following problem when deciding on its nominal price at $t = 1$:

$$\max_{p_{ik}} \quad (1-\tau_i)p_{ik}y_{ik} - wl_{ik} - \sum_{j\in\mathcal{F}_i} p_j x_{ij,k} \tag{B.4}$$

$$\text{subject to} \quad y_{ik} = (p_{ik}/p_i)^{-\theta_i} y_i \tag{B.5}$$

as well as the labor and intermediate input demand constraints (B.1) and (B.2). The first-order conditions of this optimization implies that

$$(1-\tau_i)(1-\theta_i)(p_{ik}/p_i)^{-\theta_i} y_i - (y_{ik}/Q_{ik})^{1/(1-\sum_{j\in\mathcal{R}_i} a_{ij})} \frac{1}{y_{ik}} \frac{\mathrm{d}y_{ik}}{\mathrm{d}p_{ik}} = 0. \tag{B.6}$$

Solving this optimization problem implies that the nominal price set by firm k in industry i is given by

$$p_{ik} = \left((p_i^{\theta_i} y_i)^{\sum_{j\in\mathcal{R}_i} a_{ij}} Q_{ik}^{-1} \right)^{1/(1+(\theta_i-1)\sum_{j\in\mathcal{R}_i} a_{ij})}, \tag{B.7}$$

where Q_{ik} is given by (B.3) and we are using the assumption that $\tau_i = 1/(1-\theta_i)$. With the firm's price and quantity decisions at $t = 1$ in hand, we can now turn to the rigid intermediate input decisions of the firm at $t = 0$. Recall that firms choose their rigid intermediate inputs in order to maximize the expected real value of their profits given their information set. Therefore, firm k in industry i faces the following optimization problem at $t = 0$:

$$\max_{\{x_{ij,k}\}_{j \in \mathcal{R}_i}} \quad \mathbb{E}_i \left[\frac{U'(C)}{P} \left((1-\tau_i) p_{ik} y_{ik} - w l_{ik} - \sum_{j=1}^{n} p_j x_{ij,k} \right) \right]$$

subject to constraints (B.1)–(B.2), (B.5), and (B.7), where $\mathbb{E}_i[\cdot]$ denotes the expectation operator with respect to the information set of firms in industry i, $U'(C) = 1/C$ is the household's marginal utility, and P is the price of the consumption good bundle. Note that, $PC = m$. Therefore, the first-order condition of the firm's problem at $t = 0$ is given by

$$\mathbb{E}_i \left[\frac{1}{m} \left((1-\tau_i)(1-\theta_i)(p_{ik}/p_i)^{-\theta_i} y_i - (y_{ik}/Q_{ik})^{1/(1-\sum_{j \in \mathcal{R}_i} a_{ij})} \frac{1}{y_{ik}} \frac{\mathrm{d}y_{ik}}{\mathrm{d}p_{ik}} \right) \frac{\mathrm{d}p_{ik}}{\mathrm{d}x_{ij,k}} \right] + \mathbb{E}_i \left[\frac{1}{m} \left((y_{ik}/Q_{ik})^{1/(1-\sum_{j \in \mathcal{R}_i} a_{ij})} \frac{1}{Q_{ik}} \frac{\mathrm{d}Q_{ik}}{\mathrm{d}x_{ij,k}} - p_j \right) \right] = 0.$$

Equation (B.6) implies that the first term on the right-hand side of the above equation is equal to zero. Furthermore, note that (B.3) implies that $\mathrm{d}Q_{ik}/\mathrm{d}x_{ij,k} = a_{ij} Q_{ik}/x_{ij,k}$. Therefore,

$$\text{(B.8)} \qquad x_{ij,k} = \frac{a_{ij}}{\mathbb{E}_i[p_j/m]} \mathbb{E}_i \left[\frac{1}{m} (y_{ik}/Q_{ik})^{1/(1-\sum_{j \in \mathcal{R}_i} a_{ij})} \right] \qquad \text{for all } j \in \mathcal{R}_i.$$

To simplify the above, note that given that all firms within the same industry are symmetric, they all set the same prices and produce the same quantities, that is, $p_{ik} = p_i$ and $y_{ik} = y_i$. Therefore, we can drop the firm index k from (B.7) and solve for Q_{ik} in terms of the price of firms in industry i:

$$\text{(B.9)} \qquad Q_{ik} = (p_i y_i)^{\sum_{j \in \mathcal{R}_i} a_{ij}} / p_i.$$

Plugging this expression back into (B.2) and (B.8), we obtain

$$x_{ij,k} = \begin{cases} a_{ij}\lambda_i m/p_j & \text{if } j \in \mathcal{F}_i \\ a_{ij}\mathbb{E}_i[\lambda_i]/\mathbb{E}_i[p_j/m] & \text{if } j \in \mathcal{R}_i, \end{cases} \tag{B.10}$$

where we are using the fact that the Domar weight of industry i is given by $\lambda_i = p_i y_i/m$. This expression together with the market-clearing condition (2.6) for sectoral good i implies that

$$y_i = c_i + \sum_{j\in\mathcal{F}_i} a_{ji}\frac{\lambda_j}{p_i/m} + \sum_{j\in\mathcal{R}_i} a_{ji}\frac{\mathbb{E}_j[\lambda_j]}{\mathbb{E}_j[p_i/m]}.$$

Multiplying both sides of the above equation by p_i/m and using the fact that $c_i = \beta_i m/p_i$—which is a consequence of the household's optimization problem—then establishes (2.11).

We next establish (2.10). To this end, note that equations (B.3) and (B.10) imply that

$$Q_{ik} = z_i w^{-\alpha_i} \prod_{j\in\mathcal{F}_i} p_j^{-a_{ij}} \prod_{j\in\mathcal{R}_i} (\mathbb{E}_i[\lambda_i]/\mathbb{E}_i[p_j/m])^{a_{ij}}.$$

Combining the above equation with the expression for Q_{ik} in (B.9) then establishes (2.10). $\square$

B.2. Proof of Lemma 1

As a first observation, note that combining (B.1) with the expression for Q_{ik} in (B.9) implies that the labor demand of firm k in industry i is given by $l_{ik} = \alpha_i\lambda_i m/w$. Therefore, aggregate demand for labor in the economy is equal to

$$\sum_{i=1}^{n}\int_0^1 l_{ik}\mathrm{d}k = (m/w)\sum_{i=1}^{n}\alpha_i\lambda_i.$$

Furthermore, note that the first-order conditions of the household's problem imply that total labor supply is given by $L = (m\chi/w)^{-\eta}$. Combining the above two equations therefore implies that the labor market equilibrium condition (2.5) is given by

$$(w-\bar{w})\left((m\chi/w)^{-\eta}-(m/w)\sum_{i=1}^{n}\alpha_i\lambda_i\right)=0, \qquad w\geq\bar{w}, \qquad \chi m/w\leq\left(\frac{1}{\chi}\sum_{i=1}^{n}\alpha_i\lambda_i\right)^{-1/(1+\eta)}.$$

We consider two separate cases. First, suppose that $w>\bar{w}$. The first condition above implies that $w=m\chi^{\eta/(1+\eta)}\left(\sum_{i=1}^{n}\alpha_i\lambda_i\right)^{1/(1+\eta)}$. This is consistent with the original conjecture as long as $\bar{w}<m\chi^{\eta/(1+\eta)}\left(\sum_{i=1}^{n}\alpha_i\lambda_i\right)^{1/(1+\eta)}$. As the second case, suppose $w=\bar{w}$. In that case, the last inequality above implies that $\bar{w}\geq m\chi^{\eta/(1+\eta)}\left(\sum_{i=1}^{n}\alpha_i\lambda_i\right)^{1/(1+\eta)}$. Putting the two cases together establishes (2.13). Finally, note that taking $\eta\to\infty$ in (2.13) implies that $w=\max\{\chi m,\bar{w}\}$. □

B.3. Proof of Proposition 2

We prove this result by establishing that the optimality conditions corresponding to the planner's problem coincide with the equilibrium conditions in equations (2.10)–(2.13). As a first observation, note that since all firms in the same industry have identical production technologies and information sets, we can drop the firm index k in the planner's problem.

To express the planner's problem, let

$$s=(z,m,(\omega_1,\ldots,\omega_n))\in S=\mathbb{R}_{+}^{n+1}\times\Omega_1\times\cdots\times\Omega_n$$

denote the aggregate state of the economy, consisting of all realized productivity and demand shocks, as well as the cross-sectional profile of signals, where $\omega_i\in\Omega_i$ denotes the component of the state observable to firms in industry i. To ensure that the planner is subject to the same information and quantity adjustment frictions as the firms, we impose the following measurability constraint on the quantities: if j is a rigid input of industry i (so that $j\in\mathcal{R}_i$), then x_{ij} can be contingent on ω_i, but not on the aggregate state s. We capture this measurability constraint by denoting corresponding input quantity by $x_{ij}(\omega_i)$. In contrast, if j is a flexible input for firms in industry i (so that $j \in \mathcal{F}_i$), then x_{ij} can be contingent on the economy's aggregate state, in which case we denote this quantity by $x_{ij}(s)$. Finally, note that since labor supply, labor demand, and consumption are not subject to informational frictions, they can depend on the economy's aggregate state. We therefore denote the corresponding quantities by $l_i(s)$, $L(s)$, and $c_i(s)$, respectively.

Using the above notation, we can now express the planner's problem as follows. The planner maximizes the household's expected utility

$$\int_{s\in S}\left(\sum_{i=1}^{n}\beta_i \log c_i(s) - \chi\frac{L^{1+1/\eta}(s)}{1+1/\eta}\right)\mathrm{d}G(s) \tag{B.11}$$

subject to the following resource and technology constraints:

$$y_i(s) = c_i(s) + \sum_{j:i\in\mathcal{R}_j} x_{ji}(\omega_j) + \sum_{j:i\in\mathcal{F}_j} x_{ji}(s) \tag{B.12}$$

$$L(s) = \sum_{i=1}^{n} l_i(s) \tag{B.13}$$

$$y_i(s) = z_i F_i(l_i(s), \{x_{ij}(s)\}_{j\in\mathcal{F}_i}, \{x_{ij}(\omega_i)\}_{j\in\mathcal{R}_i}), , \tag{B.14}$$

where $G(s)$ denotes the probability distribution of the economy's aggregate state and F_i denotes the production function of firms in industry i and is given by (2.1). The Lagrangian corresponding to the above problem is thus given by

$$\begin{aligned}\mathcal{L} = &\int_{s\in S}\left(\sum_{i=1}^{n}\beta_i \log c_i(s) - \chi\frac{L^{1+1/\eta}(s)}{1+1/\eta}\right)\mathrm{d}G(s) + \int_{s\in S}\nu_0(s)\left(L(s) - \sum_{i=1}^{n} l_i(s)\right)\mathrm{d}G(s)\\ &+ \sum_{i=1}^{n}\int_{s\in S}\psi_i(s)\Big(y_i(s) - c_i(s) - \sum_{j:i\in\mathcal{R}_j} x_{ji}(\omega_j) - \sum_{j:i\in\mathcal{F}_j} x_{ji}(s)\Big)\mathrm{d}G(s)\\ &+ \sum_{i=1}^{n}\int_{s\in S}\nu_i(s)\Big(z_i F_i(l_i(s), \{x_{ij}(s)\}_{j\in\mathcal{F}_i}, \{x_{ij}(\omega_i)\}_{j\in\mathcal{R}_i}) - y_i(s)\Big)\mathrm{d}G(s).\end{aligned}$$

where $\nu_i(s)\mathrm{d}G(s)$ is the Lagrange multiplier corresponding to good i's resource constraint (B.12), $\nu_0(s)\mathrm{d}G(s)$ is the multiplier corresponding to labor resource constraint (B.13), and $\psi_i(s)\mathrm{d}G(s)$ is the multiplier for industry i's technology constraint, (B.14). Therefore, the first-order conditions with respect to $c_i(s)$, $L(s)$, and $y_i(s)$ are given by

$$\beta_i/c_i(s) = \psi_i(s), \qquad \chi L^{1/\eta}(s) = \nu_0(s), \qquad \psi_i(s) = \nu_i(s), \tag{B.15}$$

respectively, whereas the first-order conditions with respect to $l_i(s)$ and $x_{ij}(s)$ for $j \in \mathcal{F}_i$ are given by

$$\nu_0(s) = \nu_i(s) z_i \frac{\partial F_i}{\partial l_i}(s) = \alpha_i \nu_i(s) y_i(s)/l_i(s) \tag{B.16}$$

$$\psi_j(s) = \nu_i(s) z_i \frac{\partial F_i}{\partial x_{ij}}(s) = a_{ij} \nu_i(s) y_i(s)/x_{ij}(s), \tag{B.17}$$

respectively. Finally, the first-order condition with respect to the rigid input $x_{ij}(\omega_i)$ is given by

$$\int_{s\in\Omega_i} \psi_j(s)\mathrm{d}G(s) = a_{ij} \int_{s\in\Omega_i} \nu_i(s) y_i(s)/x_{ij}(\omega_i)\mathrm{d}G(s),$$

where $\Omega_i \subseteq S$ denotes the subset of states with corresponding element ω_i. Note that dividing both sides of the above equation by $G(\Omega_i)$ leads to

$$\mathbb{E}_i[\psi_j(s)] = a_{ij}\mathbb{E}_i[\nu_i(s) y_i(s)]/x_{ij}(\omega_i). \tag{B.18}$$

Plugging in the expressions for $c_i(s)$, $x_{ij}(s)$, and $x_{ij}(\omega_i)$ in (B.15), (B.17), and (B.18) into the resource constraint (B.12) implies that

$$\psi_i(s) y_i(s) = \beta_i + \sum_{j:i\in\mathcal{R}_j} a_{ji}\psi_i(s)\mathbb{E}_j[\psi_j(s) y_j(s)]/\mathbb{E}_j[\psi_i(s)] + \sum_{j:i\in\mathcal{F}_j} a_{ji}\psi_j(s) y_j(s), \tag{B.19}$$

where we are using the fact that $\nu_i(s) = \psi_i(s)$, established in (B.15). Next, note that plugging the same expressions and the expression for $l_i(s)$ in (B.16) into the technology constraint in (B.14) leads to

(B.20)
$$y_i(s) = z_i \left(\psi_i(s)y_i(s)/\nu_0(s)\right)^{\alpha_i} \prod_{j\in\mathcal{F}_i} \left(\psi_i(s)y_i(s)/\psi_j(s)\right)^{a_{ij}} \prod_{j\in\mathcal{R}_i} \left(\mathbb{E}_i[\psi_i(s)y_i(s)]/\mathbb{E}_i[\psi_j(s)]\right)^{a_{ij}}.$$

Finally, plugging the expressions for $L(s)$ and $l_i(s)$ in (B.15) and (B.16) into the resource constraint for labor (B.13) implies that

$$\sum_{i=1}^{n} \alpha_i \psi_i(s) y_i(s) = v_0^{1+\eta}(s)/\chi^{\eta}. \tag{B.21}$$

The proof is complete once we verify that equations (B.19)–(B.21) coincide with equilibrium conditions (2.10)–(2.13). We do so by a simple change of variables. Let

$$\lambda_i(s) = \psi_i(s)y_i(s), \qquad p_i(s) = \psi_i(s)m(s), \qquad w(s) = \nu_0(s)m(s),$$

where $m(s)$ is an arbitrary function. Using this change of variables, it is then immediate to verify that, as long as the downward nominal wage rigidity constraint does not bind (that is $w > \bar{w}$), then equations (B.19)–(B.21) reduce to (2.10)–(2.13). □

B.4. An Auxiliary Result

We now state and prove a result that provides an exact expression for aggregate output in terms of model primitives and the nominal wage when there is only a single rigid industry. We will use this result in proving Propositions 3 and 4.

Proposition B.4.1. *If r is the only industry that is subject to frictions and Assumption 1 is satisfied, then,*

$$\log C = \sum_{j=1}^{n} \lambda_j^{\text{ss}} \log z_j - \log(w/m) - \lambda_r^{\text{ss}} \sum_{j \in \mathcal{R}_r} a_{rj} \mathbb{K}_r \left(\log(w/m) - \sum_{s=1}^{n} \ell_{js} \log z_s \right), \tag{B.22}$$

where $\mathbb{K}_r(x) = \log \mathbb{E}_r[\exp(x)] - x$.

PROOF. We first show that $\lambda_r = \lambda_r^{\text{ss}}$. Since industry r is the only industry subject to informational frictions, equation (2.11) implies that

$$\lambda_i = \beta_i + \sum_{j=1}^{n} a_{ji}\lambda_j + a_{ri}\left(\mathbb{E}_r[\lambda_r] \frac{p_i/m}{\mathbb{E}_r[p_i/m]} - \lambda_r \right) \mathbb{I}_{\{i \in \mathcal{R}_r\}}. \tag{B.23}$$

Taking expectations from both sides of the above equation with respect to the information set of industry r implies that $\mathbb{E}_r[\lambda_i] = \beta_i + \sum_{j=1}^{n} a_{ji}\mathbb{E}_r[\lambda_j]$ for all i. Solving this system of equations for $\mathbb{E}_r[\lambda_i]$ implies that $\mathbb{E}_r[\lambda_i] = \lambda_i^{\text{ss}}$, where is the steady-state Domar weight of industry i. Consequently, we can rewrite equation (B.23) as follows:

$$\lambda_i = \beta_i + \sum_{j=1}^{n} a_{ji}\lambda_j + a_{ri}\left(\lambda_r^{\text{ss}} \frac{p_i/m}{\mathbb{E}_r[p_i/m]} - \lambda_r \right) \mathbb{I}_{\{i \in \mathcal{R}_r\}},$$

Furthermore, note that the steady-state Domar weights satisfy the following system of equations: $\lambda_i^{\text{ss}} = \beta_i + \sum_{j=1}^{n} a_{ji}\lambda_j^{\text{ss}}$ for all i. Subtracting this equation from the previous one therefore implies that

$$\Delta_i = \sum_{j=1}^{n} a_{ji}\Delta_j + a_{ri}\left(\lambda_r^{\text{ss}} \frac{p_i/m}{\mathbb{E}_r[p_i/m]} - \lambda_r \right) \mathbb{I}_{\{i \in \mathcal{R}_r\}}.$$

where $\Delta_i = \lambda_i - \lambda_i^{\text{ss}}$. Solving the above system of equations for Δ_i implies that

$$\Delta_i = \sum_{j=1}^{n} \ell_{ji} a_{rj} \left(\lambda_r^{\text{ss}} \frac{p_j/m}{\mathbb{E}_r[p_j/m]} - \lambda_r \right) \mathbb{I}_{\{j \in \mathcal{R}_r\}}, \tag{B.24}$$

where ℓ_{ji} denotes the (j,i) element of the economy's Leontief inverse $\mathbf{L} = (\mathbf{I} - \mathbf{A})^{-1}$. Setting $i = r$ in the above equation and using Assumption 1 then implies that the right-hand side of the above equation is equal to zero, thus establishing that $\lambda_r = \lambda_r^{\text{ss}}$.

Next, note that since industry r is the only industry that is subject to frictions, equation (2.10) implies that the (log) nominal price of industry $i \neq r$ is given by

$$\log p_i = -\log z_i + \alpha_i \log w + \sum_{j=1}^{n} a_{ij} \log p_j.$$

Let $\tilde{p} \in \mathbb{R}^{n-1}$ denote the vector of nominal prices for all industries $i \neq r$ and let $\tilde{\mathbf{A}} \in \mathbb{R}^{(n-1)\times(n-1)}$ denote the sub-block of the input-output matrix $\mathbf{A}$ corresponding to all industries except for r. Writing the above equation in vector form therefore implies that $\log \tilde{p} = -\log \tilde{z} + \tilde{\alpha} \log w + \tilde{\mathbf{A}} \log \tilde{p} + \tilde{a}_r \log p_r$, where $\tilde{\alpha}$ and $\tilde{z}$ denote the vectors of labor shares and productivity shocks for all $i \neq r$ and $\tilde{a}_r \in \mathbb{R}^{n-1}$ is a vector with elements a_{is} for all $i \neq r$. Consequently,

$$\log \tilde{p} = \tilde{\mathbf{L}}\tilde{\alpha} \log w - \tilde{\mathbf{L}} \log \tilde{z} + \tilde{\mathbf{L}}\tilde{a}_r \log p_r,$$

where $\tilde{\mathbf{L}} = (\mathbf{I} - \tilde{\mathbf{A}})^{-1}$. Under Assumption 1, the elements of $\tilde{\mathbf{L}}$ can be expressed in terms of the elements of the economy's Leontief inverse $\mathbf{L}$. In particular, $\tilde{\ell}_{ij} = \ell_{ij} - \ell_{ir}\ell_{rj}$ for all $i, j \neq r$. Hence,

$$\log p_i = \log w \sum_{j \neq r} (\ell_{ij} - \ell_{ir}\ell_{rj})\alpha_j - \sum_{j \neq r} (\ell_{ij} - \ell_{ir}\ell_{rj}) \log z_j + \log p_r \sum_{j \neq r} (\ell_{ij} - \ell_{ir}\ell_{rj}) a_{jr}$$

for all $i \neq r$. Consequently,

$$\log p_i = (1 - \ell_{ir}) \log w + \ell_{ir} \log p_r - \sum_{j \neq r} (\ell_{ij} - \ell_{ir}\ell_{rj}) \log z_j \tag{B.25}$$

for all $i \neq r$, where we are using the fact that $\sum_{j=1}^{n} \ell_{ij}\alpha_j = 1$ for all i and $\ell_{rr} = 1$, the latter of which is a consequence of Assumption 1. The above equation expresses all prices in terms of the price of industry r and the nominal wage. With the above in hand, we can therefore obtain an expression for aggregate output in terms of the nominal price of industry r. In particular, the fact that $\log C = \log m - \sum_{i=1}^{n} \beta_i \log p_i$ together with (B.25) implies that

$$\log C = \log(m/w) - \lambda_r^{\text{ss}} \log(p_r/w) + \sum_{j \neq r} (\lambda_j^{\text{ss}} - \lambda_r^{\text{ss}} \ell_{rj}) \log z_j, \tag{B.26}$$

where λ_j^{ss} denotes the steady-state Domar weight of industry j. Therefore, to obtain the expression for aggregate output is sufficient to characterize $\log p_r$. To this end, note that setting $i = r$ in equation (2.10) implies that

$$\log p_r = -\log z_r + \alpha_r \log w + \sum_{j \in \mathcal{F}_r} a_{rj} \log p_j + \log m \sum_{j \in \mathcal{R}_r} a_{rj} + \sum_{j \in \mathcal{R}_r} a_{rj} \log \mathbb{E}_r[p_j/m],$$

where we are also using the fact that $\lambda_r = \lambda_r^{\text{ss}}$. Replacing for $\log p_j$ from (B.25) for all $j \neq r$ into the above equation and using the implication of Assumption 1 that $a_{rj}\ell_{jr} = 0$ for all $j \in \mathcal{F}_r$ implies that

$$\begin{aligned} \log(p_r/m) = &-\log z_r + \left(\alpha_r + \sum_{j \in \mathcal{F}_r} a_{rj}\right) \log(w/m) - \sum_{j \in \mathcal{F}_r} \sum_{s=1}^{n} a_{rj} \ell_{js} \log z_s \\ &+ \sum_{j \in \mathcal{R}_r} a_{rj} \log \mathbb{E}_r \left[\exp\left(\log(w/m) - \sum_{s=1}^{n} \ell_{js} \log z_s\right)\right]. \end{aligned} \tag{B.27}$$

Plugging the above into the expression for $\log C$ in (B.26) and using Assumption 1 then establishes (B.22). $\square$

B.5. Proof of Proposition 3

Proof of part (a). Recall from the proof of Proposition B.4.1 that $\Delta_i = \lambda_i - \lambda_i^{\text{ss}}$ satisfies (B.24). As a result,

$$\sum_{i=1}^{n} \alpha_i \lambda_i = \sum_{i=1}^{n} \alpha_i (\lambda_i^{\text{ss}} + \Delta_i) = 1 + \lambda_r^{\text{ss}} \sum_{j \in \mathcal{R}_r} a_{rj} \left(\frac{p_j/m}{\mathbb{E}_r[p_j/m]} - 1 \right),$$

where we are using $\sum_{i=1}^{n} \alpha_i \lambda_i^{\text{ss}} = \sum_{i=1}^{n} \alpha_i \ell_{ji} = 1$ and the fact that $\lambda_r = \lambda_r^{\text{ss}}$, established in the proof of Proposition B.4.1. Therefore, to a first-order approximation

$$\log\left(\sum_{i=1}^{n} \alpha_i \lambda_i\right) = \lambda_r^{\text{ss}} \sum_{j \in \mathcal{R}_r} a_{rj} \left(\frac{p_j/m}{\mathbb{E}_r[p_j/m]} - 1 \right) = \lambda_r^{\text{ss}} \sum_{j \in \mathcal{R}_r} a_{rj} \left(\log(p_j/m) - \mathbb{E}_r[\log(p_j/m)] \right).$$

Combining the above with equation (2.13), together with the assumption that the downward constraint on nominal wage does not bind, implies that

$$\log(w/m) = \frac{\eta}{1+\eta} \log \chi + \frac{1}{1+\eta} \lambda_r^{\text{ss}} \sum_{j \in \mathcal{R}_r} a_{rj} \Big(\log(p_j/m) - \mathbb{E}_r[\log(p_j/m)] \Big). \tag{B.28}$$

Next, recall from the proof of Proposition B.4.1 that $\log(p_r/m)$ is given by (B.27). Thus, to a first-order approximation,

$$\log(p_r/m) = -\log z_r + \left(\alpha_r + \sum_{j \in \mathcal{F}_r} a_{rj} \right) \log(w/m) - \sum_{j \in \mathcal{F}_r} \sum_{s=1}^{n} a_{rj} \ell_{js} \log z_s$$
$$+ \sum_{j \in \mathcal{R}_r} a_{rj} \mathbb{E}_r \left[\log(w/m) - \sum_{s=1}^{n} \ell_{js} \log z_s \right].$$

Plugging this back into the expression for $\log p_i$ in (B.25) we get

$$\log(p_i/m) = \log(w/m) - \sum_{j=1}^{n} \ell_{ij} \log z_j - \ell_{ir} \left(\sum_{j \in \mathcal{R}_r} a_{rj} \right) \Big(\log(w/m) - \mathbb{E}_r[\log(w/m)] \Big)$$
$$- \ell_{ir} \sum_{j \in \mathcal{R}_r} a_{rj} \sum_{p=1}^{n} \ell_{jp} \mathbb{E}_r \left([\log z_p] - \log z_p \right)$$

for all $i \neq r$. Taking expectations from both sides of the above equation and subtracting it from both sides therefore implies that

$$\sum_{i\in\mathcal{R}_r} a_{ri}(\log(p_i/m) - \mathbb{E}_r[\log(p_i/m)]) = \sum_{i\in\mathcal{R}_r} a_{ri}\Big(\log(w/m) - \mathbb{E}_r[\log(w/m)]\Big) - \sum_{i\in\mathcal{R}_r} a_{ri}\sum_{j=1}^{n}\ell_{ij}(\log z_j - \mathbb{E}_r[\log z_j]),$$

Note that (B.28) implies that $\mathbb{E}_r[\log(w/m)] = \frac{\eta}{1+\eta}\log\chi$. Therefore, we can rewrite the above equation as follows:

$$\sum_{i\in\mathcal{R}_r} a_{ri}(\log(p_i/m) - \mathbb{E}_r[\log(p_i/m)]) = \sum_{i\in\mathcal{R}_r} a_{ri}\Big(\log(w/m) - \frac{1}{1+1/\eta}\log\chi\Big) - \sum_{i\in\mathcal{R}_r}\sum_{j=1}^{n} a_{ri}\ell_{ij}(\log z_j - \mathbb{E}_r[\log z_j]),$$

Combining the above equation with (B.28) and solving for $\log(w/m)$ we obtain,

$$\log(w/m) = \frac{1}{1+1/\eta}\log\chi - \frac{\lambda_r^{\mathrm{ss}}}{1+\eta-\lambda_r^{\mathrm{ss}}\sum_{i\in\mathcal{R}_r} a_{ri}}\sum_{i\in\mathcal{R}_r}\sum_{j=1}^{n} a_{ri}\ell_{ij}(\log z_j - \mathbb{E}_r[\log z_j]).$$

Now, plugging the above expression into the expression for $\log C$ in (B.22) and performing a first-order approximation establishes (2.14). □

Proof of part (b). Recall from Proposition B.4.1 that log aggregate output is given by (B.22). Furthermore, note that by Lemma 1, when labor supply is fully elastic and the downward constraint on the nominal wage does not bind, $\log(w/m) = \log\chi$. Therefore, the expression in (B.22) simplifies as follows:

$$\log C = \sum_{j=1}^{n}\lambda_j^{\mathrm{ss}}\log z_j - \log\chi - \lambda_r^{\mathrm{ss}}\sum_{j\in\mathcal{R}_r} a_{rj}\mathbb{K}_r\left(-\sum_{s=1}^{n}\ell_{js}\log z_s\right),$$

where we are using the fact that $\mathbb{K}_r(x+a) = \mathbb{K}_r(x)$ for any constant a. Noting that $\log C^* = \sum_{j=1}^{n}\lambda_j^{\mathrm{ss}}\log z_j - \log\chi$ then establishes (2.15). □

B.6. Proof of Proposition 4

Recall from Proposition B.4.1 that log aggregate output is given by (B.22). Therefore, when the downward constraint on the nominal wage binds (that is, $w = \bar{w}$) and the absence of productivity shocks, the expression for log aggregate output reduces to

$$\log C = \log m - \log \bar{w} - \mathbb{K}_r\left(-\log m\right)\lambda_r^{\mathrm{ss}} \sum_{j \in \mathcal{R}_r} a_{rj},$$

which coincides with the expression in (2.19). Also note that (2.20) follows immediately from the observation that $\log P = \log m - \log C$.

B.7. Proof of Propositions 5

Let $\mathbb{E}_\omega[\cdot]$ denote the expectation operator conditional on the public signal, ω. Taking conditional expectations from both sides of (2.11) implies that $\mathbb{E}_\omega[\lambda_i] = \beta_i + \sum_{j=1}^n a_{ji}\mathbb{E}_\omega[\lambda_j]$ for all i. On the other hand, note that the steady-state Domar weights of all industries also satisfy the following system of equations: $\lambda_i^{\mathrm{ss}} = \beta_i + \sum_{j=1}^n a_{ji}\lambda_j^{\mathrm{ss}}$. Comparing the two equations then implies that

$$\mathbb{E}_\omega[\lambda_i] = \lambda_i^{\mathrm{ss}} \qquad \text{for all } i.$$

Plugging this into equation (2.10) and taking logarithms from both sides then implies that

$$\log(p_i/m) = -\log z_i + \alpha_i \log(w/m) + \sum_{j \in \mathcal{F}_i} a_{ij} \log(p_j/m) + \sum_{j \in \mathcal{R}_i} a_{ij}\left(\log \mathbb{E}_\omega[p_j/m] + \log(\lambda_i/\lambda_i^{\mathrm{ss}})\right)$$

To simplify notation, define $\hat{p}_i = p_i/m$ and $\hat{w} = w/m$. Writing the above equation in matrix form, we get

$$\log \hat{p} = -\log z + \alpha \log \hat{w} + \mathbf{A}_f \log \hat{p} + \mathbf{A}_r \log \mathbb{E}_\omega[\hat{p}] + \mathrm{diag}(\mathbf{A}_r \mathbf{1}) \log(\lambda/\lambda^{\mathrm{ss}}),$$

where $\mathbf{A}_f$ is the matrix whose (i,j) element is equal to a_{ij} if $j \in \mathcal{F}_i$ and is equal to zero otherwise and $\mathbf{A}_r = \mathbf{A} - \mathbf{A}_f$. Consequently,

$$\log \hat{p} = \xi + \mathbf{L}_f \alpha \log \hat{w} + \mathbf{L}_f \mathbf{A}_r \log \mathbb{E}_\omega[\hat{p}], \tag{B.29}$$

where $\mathbf{L}_f = (\mathbf{I} - \mathbf{A}_f)^{-1}$ and

$$\xi = -\mathbf{L}_f \log z + \mathbf{L}_f \operatorname{diag}(\mathbf{A}_r \mathbf{1}) \log(\lambda/\lambda^{\text{ss}}). \tag{B.30}$$

Exponentiating both sides of (B.29), taking conditional expectations, and then taking logarithms implies that

$$\log \mathbb{E}_\omega[\hat{p}] = \log \mathbb{E}_\omega[e^\xi] + \mathbf{L}_f \alpha \log \hat{w} + \mathbf{L}_f \mathbf{A}_r \log \mathbb{E}_\omega[\hat{p}],$$

where note that since $\eta \to \infty$, Lemma 1 implies that $\hat{w} = w/m = \chi$, which is deterministic and hence is measurable with the respect to the firms' common information structure. Solving for $\log \mathbb{E}_\omega[\hat{p}]$ and using the observation that $(\mathbf{I} - \mathbf{L}_f \mathbf{A}_r)^{-1} = (\mathbf{I} - \mathbf{L}_f(\mathbf{A} - \mathbf{A}_f))^{-1} = \mathbf{L}(\mathbf{I} - \mathbf{A}_f)$, we can rewrite the above equation as follows:

$$\log \mathbb{E}_\omega[\hat{p}] = \mathbf{L}(\mathbf{I} - \mathbf{A}_f) \log \mathbb{E}_\omega[e^\xi] + \mathbf{1} \log \chi,$$

Plugging the above expression back into (B.29) leads to the following expression for log prices in terms of vector ξ defined in (B.30):

$$\log \hat{p} = \mathbf{1} \log \chi + \xi + \mathbf{L}\mathbf{A}_r \log \mathbb{E}_\omega[e^\xi].$$

Combining this equation with the observation that $\log C = \log m - \sum_{i=1}^n \beta_i \log p_i$ we get the following expression for log aggregate output in terms of vector ξ:

$$\log C = -\log \chi - \lambda^{\text{ss}\prime}(\mathbf{I} - \mathbf{A})\xi - \lambda^{\text{ss}\prime} \mathbf{A}_r \log \mathbb{E}_\omega[e^\xi],$$

which to a second-order approximation is equal to

$$\log C = -\log \chi - \lambda^{\text{ss}\prime}(\mathbf{I} - \mathbf{A})\xi - \lambda^{\text{ss}\prime}\mathbf{A}_r \left(\mathbb{E}_\omega[\xi] + \frac{1}{2}\text{var}_\omega(\xi)\right). \tag{B.31}$$

To express log output in (B.31) in terms of model primitives, we next need to solve for ξ and its first two conditional moments. We thus turn to (2.11), which can be rewritten as follows:

$$\lambda = \beta + \mathbf{A}_f'\lambda + \text{diag}(\mathbf{A}_r'\lambda^{\text{ss}})\frac{\hat{p}}{\mathbb{E}_\omega[\hat{p}]}.$$

Solving for the vector of Domar weights and using (B.29), we get

$$\lambda = \mathbf{L}_f'\beta + \mathbf{L}_f' \,\text{diag}(\mathbf{A}_r'\lambda^{\text{ss}})\frac{e^\xi}{\mathbb{E}_\omega[e^\xi]} = \lambda^{\text{ss}} + \mathbf{L}_f' \,\text{diag}(\mathbf{A}_r'\lambda^{\text{ss}})\left(\frac{e^\xi}{\mathbb{E}_\omega[e^\xi]} - \mathbf{1}\right),$$

and as a result,

$$\lambda/\lambda^{\text{ss}} = \mathbf{1} + \text{diag}^{-1}(\lambda^{\text{ss}})\mathbf{L}_f' \,\text{diag}(\mathbf{A}_r'\lambda^{\text{ss}})\left(e^{\xi - \log \mathbb{E}_\omega[e^\xi]} - \mathbf{1}\right). \tag{B.32}$$

Therefore, to a second-order approximation,

$$\begin{aligned}\lambda/\lambda^{\text{ss}} &= \mathbf{1} + \mathbf{H}'\left(\xi - \log \mathbb{E}_\omega[e^\xi] + \frac{1}{2}\,\text{diag}\left((\xi - \log \mathbb{E}_\omega[e^\xi])(\xi - \log \mathbb{E}_\omega[e^\xi])'\right)\right)\\ &= \mathbf{1} + \mathbf{H}'\left(\xi - \mathbb{E}_\omega[\xi] - \frac{1}{2}\text{var}_\omega(\xi) + \frac{1}{2}\,\text{diag}\left((\xi - \mathbb{E}_\omega[\xi])(\xi - \mathbb{E}_\omega[\xi])'\right)\right),\end{aligned}$$

where $\mathbf{H}' = \text{diag}^{-1}(\lambda^{\text{ss}})\mathbf{L}_f' \,\text{diag}(\mathbf{A}_r'\lambda^{\text{ss}})$. Plugging the above expression into equation (B.30) and performing a second-order approximation, we get

$$\begin{aligned}\xi = &-\mathbf{L}_f \log z + \mathbf{L}_f \,\text{diag}(\mathbf{A}_r\mathbf{1})\mathbf{H}'\left(\xi - \mathbb{E}_\omega[\xi] - \frac{1}{2}\text{var}_\omega(\xi) + \frac{1}{2}\,\text{diag}((\xi - \mathbb{E}_\omega[\xi])(\xi - \mathbb{E}_\omega[\xi])')\right)\\ &- \frac{1}{2}\mathbf{L}_f \,\text{diag}(\mathbf{A}_r\mathbf{1})\,\text{diag}(\mathbf{H}'(\xi - \mathbb{E}_\omega[\xi])(\xi - \mathbb{E}_\omega[\xi])'\mathbf{H}).\end{aligned} \tag{B.33}$$

Taking conditional expectations from both sides of the above equation implies that

$$\mathbb{E}_\omega[\xi] = -\mathbf{L}_f \mathbb{E}_\omega[\log z] - \frac{1}{2}\mathbf{L}_f \operatorname{diag}(\mathbf{A}_r \mathbf{1}) \operatorname{diag}(\mathbf{H}' \mathrm{var}_\omega(\xi) \mathbf{H}), \tag{B.34}$$

where $\mathrm{var}_\omega(\xi)$ denotes the variance-covariance matrix of ξ conditional on the common signal ω. Subtracting the above equation from (B.33) leads to

$$\begin{aligned}
\xi - \mathbb{E}_\omega[\xi] = -&\mathbf{L}_f(\log z - \mathbb{E}_\omega[\log z]) \\
&+ \mathbf{L}_f \operatorname{diag}(\mathbf{A}_r \mathbf{1})\mathbf{H}' \left(\xi - \mathbb{E}_\omega[\xi] - \frac{1}{2}\mathrm{var}_\omega(\xi) + \frac{1}{2}\operatorname{diag}((\xi - \mathbb{E}_\omega[\xi])(\xi - \mathbb{E}_\omega[\xi])')\right) \\
&- \frac{1}{2}\mathbf{L}_f \operatorname{diag}(\mathbf{A}_r \mathbf{1}) \operatorname{diag}(\mathbf{H}'(\xi - \mathbb{E}_\omega[\xi])(\xi - \mathbb{E}_\omega[\xi])'\mathbf{H}) + \frac{1}{2}\mathbf{L}_f \operatorname{diag}(\mathbf{A}_r \mathbf{1}) \operatorname{diag}(\mathbf{H}' \mathrm{var}_\omega(\xi) \mathbf{H}).
\end{aligned}$$

As a result,

$$\mathrm{var}_\omega(\xi) = \mathrm{var}_\omega(\mathbf{L}_f \log z) + \mathbf{L}_f \operatorname{diag}(\mathbf{A}_r \mathbf{1})\mathbf{H}' \mathrm{var}_\omega(\xi) \mathbf{H} \operatorname{diag}(\mathbf{A}_r \mathbf{1})\mathbf{L}_f'.$$

Solving for $\mathrm{var}_\omega(\xi)$ from the above equation, we get

$$\mathrm{var}_\omega(\xi) = \sum_{k=0}^{\infty} (\mathbf{L}_f \operatorname{diag}(\mathbf{A}_r \mathbf{1})\mathbf{H}')^k \mathrm{var}_\omega(\mathbf{L}_f \log z)(\mathbf{H} \operatorname{diag}(\mathbf{A}_r \mathbf{1})\mathbf{L}_f')^k. \tag{B.35}$$

Plugging the expressions in (B.34) and (B.35) into (B.31) and taking conditional expectations then implies that

$$\mathbb{E}_\omega[\log C] = \mathbb{E}_\omega[\log C^*] + \frac{1}{2}\lambda^{\mathrm{ss}\prime} \operatorname{diag}(\mathbf{A}_r \mathbf{1}) \operatorname{diag}(\mathbf{H}' \mathrm{var}_\omega(\xi) \mathbf{H}) - \frac{1}{2}\lambda^{\mathrm{ss}\prime} \mathbf{A}_r \mathrm{var}_\omega(\xi).$$

Taking unconditional expectations from both side of the above equation and letting $\mathbf{Q} = \mathbb{E}[\mathrm{var}_\omega(\xi)]$ then establishes the result. □

B.8. Proof of Propositions 6

Recall from the proof of Proposition 5 that $\log \hat{p}_i = \log p_i - \log m$ satisfies equation (B.29), where vector ξ is given by (B.30). In the absence of productivity shocks, and when the downward constraint on the nominal wage binds, this implies that

$$\log \hat{p} = \theta + \mathbf{L}_f \alpha \log \bar{w} + \mathbf{L}_f \mathbf{A}_r \log \mathbb{E}_\omega[\hat{p}],$$

where

$$\theta = -\mathbf{L}_f \alpha \log m + \mathbf{L}_f \operatorname{diag}(\mathbf{A}_r \mathbf{1}) \log(\lambda/\lambda^{\text{ss}}). \tag{B.36}$$

Exponentiating both sides of the above equation, taking conditional expectations, and then taking logarithms implies that

$$\log \mathbb{E}_\omega[\hat{p}] = \log \mathbb{E}_\omega[e^\theta] + \mathbf{L}_f \alpha \log \bar{w} + \mathbf{L}_f \mathbf{A}_r \log \mathbb{E}_\omega[\hat{p}].$$

Solving for $\log \mathbb{E}_\omega[\hat{p}]$ and using the observation that $(\mathbf{I} - \mathbf{L}_f \mathbf{A}_r)^{-1} = (\mathbf{I} - \mathbf{L}_f(\mathbf{A} - \mathbf{A}_f))^{-1} = \mathbf{L}(\mathbf{I} - \mathbf{A}_f)$, we can rewrite the above equation as follows:

$$\log \mathbb{E}_\omega[\hat{p}] = \mathbf{L}(\mathbf{I} - \mathbf{A}_f) \log \mathbb{E}_\omega[e^\theta] + \mathbf{1} \log \bar{w}.$$

Plugging the above back into the expression for $\log \hat{p}$ leads to the following expression for log prices:

$$\log \hat{p} = \mathbf{1} \log \bar{w} + \theta + \mathbf{L}\mathbf{A}_r \log \mathbb{E}_\omega[e^\theta].$$

Combining this equation with the observation that $\log C = -\sum_{i=1}^{n} \beta_i \log \hat{p}_i$ we get the following expression for log aggregate output:

$$\log C = -\log \bar{w} - \lambda^{\text{ss}\prime}(\mathbf{I} - \mathbf{A})\theta - \lambda^{\text{ss}\prime}\mathbf{A}_r \log \mathbb{E}_\omega[e^\theta].$$

Next, note that steps identical to those in the proof of Proposition 5 imply that we can write the ratio of Domar weights to their steady-state values as follows:

$$\lambda/\lambda^{\text{ss}} = \mathbf{1} + \text{diag}^{-1}(\lambda^{\text{ss}})\mathbf{L}_f' \,\text{diag}(\mathbf{A}_r'\lambda^{\text{ss}}) \left(e^{\theta - \log \mathbb{E}_\omega[e^\theta]} - \mathbf{1}\right),$$

and hence, to a second-order approximation,

$$\lambda/\lambda^{\text{ss}} = \mathbf{1} + \mathbf{H}' \left(\theta - \mathbb{E}_\omega[\theta] - \frac{1}{2}\text{var}_\omega(\theta) + \frac{1}{2}\,\text{diag}\left((\theta - \mathbb{E}_\omega[\theta])(\theta - \mathbb{E}_\omega[\theta])'\right)\right),$$

Plugging the above expression into equation (B.36) and performing a second-order approximation, we get

(B.37)

$$\theta = -\mathbf{L}_f\alpha \log m + \mathbf{L}_f \,\text{diag}(\mathbf{A}_r\mathbf{1})\mathbf{H}' \left(\theta - \mathbb{E}_\omega[\theta] - \frac{1}{2}\text{var}_\omega(\theta) + \frac{1}{2}\,\text{diag}((\theta - \mathbb{E}_\omega[\theta])(\theta - \mathbb{E}_\omega[\theta])')\right)$$
$$- \frac{1}{2}\mathbf{L}_f \,\text{diag}(\mathbf{A}_r\mathbf{1})\,\text{diag}(\mathbf{H}'(\theta - \mathbb{E}_\omega[\theta])(\theta - \mathbb{E}_\omega[\theta])'\mathbf{H}).$$

Taking conditional expectations from both sides of the above equation implies that

$$\text{(B.38)} \qquad \mathbb{E}_\omega[\theta] = -\mathbf{L}_f\alpha\mathbb{E}_\omega[\log m] - \frac{1}{2}\mathbf{L}_f \,\text{diag}(\mathbf{A}_r\mathbf{1})\,\text{diag}(\mathbf{H}'\text{var}_\omega(\theta)\mathbf{H}).$$

Subtracting the above equation from (B.37) leads to

$$\theta - \mathbb{E}_\omega[\theta] = -\mathbf{L}_f\alpha(\log m - \mathbb{E}_\omega[\log m])$$
$$+ \mathbf{L}_f \,\text{diag}(\mathbf{A}_r\mathbf{1})\mathbf{H}' \left(\theta - \mathbb{E}_\omega[\theta] - \frac{1}{2}\text{var}_\omega(\theta) + \frac{1}{2}\,\text{diag}((\theta - \mathbb{E}_\omega[\theta])(\theta - \mathbb{E}_\omega[\theta])')\right)$$
$$- \frac{1}{2}\mathbf{L}_f \,\text{diag}(\mathbf{A}_r\mathbf{1})\,\text{diag}(\mathbf{H}'(\theta - \mathbb{E}_\omega[\theta])(\theta - \mathbb{E}_\omega[\theta])'\mathbf{H}) + \frac{1}{2}\mathbf{L}_f \,\text{diag}(\mathbf{A}_r\mathbf{1})\,\text{diag}(\mathbf{H}'\text{var}_\omega(\theta)\mathbf{H}).$$

As a result,

$$\mathrm{var}_\omega(\theta) = \mathrm{var}_\omega(\mathbf{L}_f \alpha \log m) + \mathbf{L}_f \,\mathrm{diag}(\mathbf{A}_r \mathbf{1}) \mathbf{H}' \mathrm{var}_\omega(\theta) \mathbf{H} \,\mathrm{diag}(\mathbf{A}_r \mathbf{1}) \mathbf{L}_f'.$$

Solving for $\mathrm{var}_\omega(\theta)$ from the above equation, we get

$$\mathrm{var}_\omega(\theta) = \sum_{k=0}^{\infty} (\mathbf{L}_f \,\mathrm{diag}(\mathbf{A}_r \mathbf{1}) \mathbf{H}')^k \mathrm{var}_\omega(\mathbf{L}_f \alpha \log m) (\mathbf{H} \,\mathrm{diag}(\mathbf{A}_r \mathbf{1}) \mathbf{L}_f')^k. \tag{B.39}$$

Plugging the expressions in (B.38) and (B.39) into the expression for aggregate output and taking conditional expectations then implies that

$$\mathbb{E}_\omega[\log C] = \mathbb{E}_\omega[\log C^*] + \frac{1}{2} \lambda^{\mathrm{ss}\prime} \,\mathrm{diag}(\mathbf{A}_r \mathbf{1}) \,\mathrm{diag}(\mathbf{H}' \mathrm{var}_\omega(\theta) \mathbf{H}) - \frac{1}{2} \lambda^{\mathrm{ss}\prime} \mathbf{A}_r \mathrm{var}_\omega(\theta).$$

Taking unconditional expectations from both side then establishes the result. □

B.9. Proof of Proposition 7

We consider a more general vertical production network with industries 1 through n arranged on a chain, with industry 1 as the final good producer. We then specialize this economy to the case of $n = 3$ in Proposition 7.

Recall from Proposition 1 that nominal prices and Domar weights satisfy the system of equations in (2.10) and (2.11). Applying these equations to the vertical production network economy, we obtain

$$p_i = \frac{1}{z_i} w^{1-a_i} \left(m \frac{\mathbb{E}_i[p_{i+1}/m]}{\mathbb{E}_i[\lambda_i]/\lambda_i} \right)^{a_i} \qquad \text{for } 1 \leq i \leq n, \tag{B.40}$$

$$\lambda_{i+1} = a_i \mathbb{E}_i[\lambda_i] \frac{p_{i+1}/m}{\mathbb{E}_i[p_{i+1}/m]} \qquad \text{for } 1 \leq i \leq n-1 \tag{B.41}$$

with the initial conditions that $\lambda_1 = 1$ and the convention that $a_n = 0$. Solving for $\mathbb{E}_i[\lambda_i]/\mathbb{E}_i[p_{i+1}/m]$ from (B.41) and plugging it back into (B.40) implies that $p_i = \frac{1}{z_i} a_i^{a_i} \left(\lambda_i/\lambda_{i+1}\right)^{a_i} w^{1-a_i} p_{i+1}^{a_i}$

Hence,

$$\log p_i = a_i(\varphi_i - \varphi_{i+1}) - \log z_i + a_i \log p_{i+1} + (1 - a_i) \log w,$$

for $1 \leq i \leq n$, where $\varphi_i = \log \lambda_i - \log \lambda_i^{\text{ss}}$ and we are using the fact that $\lambda_{i+1}^{\text{ss}} = a_i \lambda_i^{\text{ss}}$. Solving the above recursion, we can express nominal prices in terms of Domar weights:

$$\log p_i = \log w - \sum_{j=i}^{n} (a_i a_{i+1} \ldots a_{j-1}) \log z_j + \sum_{j=i}^{n-1} (a_i a_{i+1} \ldots a_j)(\varphi_j - \varphi_{j+1}). \tag{B.42}$$

Next, note that, to a first-order approximation, (B.41) can be expressed as

$$\log \lambda_{i+1} = \log a_i + \log(p_{i+1}/m) + \mathbb{E}_i[\log \lambda_i] - \mathbb{E}_i[\log p_{i+1}/m],$$

and as a result,

$$\varphi_{i+1} = \mathbb{E}_i[\varphi_i] + \log(p_{i+1}/m) - \mathbb{E}_i[\log p_{i+1}/m]. \tag{B.43}$$

We now have a system of linear expectations (B.42) and (B.43) that fully pins down equilibrium nominal prices and Domar weights in terms of the productivity shocks, nominal aggregate demand, and the nominal wage.

Specializing these equations to the case of $n = 3$ and shutting off all productivity shocks, it is immediate that $\log p_3 = \log w$, and hence, we get the following equations:

$$\log p_2 = \log w + a_2(\varphi_2 - \varphi_3)$$

$$\varphi_2 = \log p_2 - \mathbb{E}_1[\log p_2] - \log m + \mathbb{E}_1[\log m]$$

$$\varphi_3 = \mathbb{E}_2[\varphi_2] + \log(w/m) - \mathbb{E}_2[\log(w/m)].$$

Replacing for φ_3 into the expression for $\log p_2$, we get

$$\log p_2 = a_2(\varphi_2 - \mathbb{E}_2[\varphi_2]) + Q + \log m, \tag{B.44}$$

where

$$Q = (1 - a_2)\log(w/m) + a_2\mathbb{E}_2[\log(w/m)]. \tag{B.45}$$

Consequently, we get the following equation for φ_2:

$$\varphi_2 = a_2(\varphi_2 - \mathbb{E}_2[\varphi_2]) - a_2(\mathbb{E}_1[\varphi_2] - \mathbb{E}_1\mathbb{E}_2[\varphi_2]) + Q - \mathbb{E}_1[Q]$$

Noting that $\mathbb{E}_1[\varphi_2] = 0$, we get

$$\varphi_2 = -\frac{a_2}{1-a_2}\mathbb{E}_2[\varphi_2] + \frac{a_2}{1-a_2}\mathbb{E}_1\mathbb{E}_2[\varphi_2] + \frac{1}{1-a_2}(Q - \mathbb{E}_1[Q]). \tag{B.46}$$

Taking expectations with respect to the information set of firms in industry 2 from both sides of (B.46) and solving for $\mathbb{E}_2[\varphi_2]$ implies that

$$\mathbb{E}_2[\varphi_2] = a_2\mathbb{E}_2\mathbb{E}_1\mathbb{E}_2[\varphi_2] + \mathbb{E}_2[Q] - \mathbb{E}_2\mathbb{E}_1[Q].$$

We can thus solve for $\mathbb{E}_2[\varphi_2]$ in terms of the infinite regress of expectations as follows:

$$\mathbb{E}_2[\varphi_2] = \sum_{s=0}^{\infty} a_2^s(\mathbb{E}_2\mathbb{E}_1)^s(\mathbb{E}_2[Q] - \mathbb{E}_2\mathbb{E}_1[Q]) = \sum_{s=0}^{\infty} a_2^s(\mathbb{E}_2\mathbb{E}_1)^s\mathbb{E}_2[Q] - \sum_{s=0}^{\infty} a_2^s(\mathbb{E}_2\mathbb{E}_1)^{s+1}[Q].$$

Plugging the above expression into (B.46), we get

$$\varphi_2 = \frac{1}{1-a_2}\left(\sum_{s=0}^{\infty} a_2^{s+1}(\mathbb{E}_1\mathbb{E}_2)^{s+1}[Q] - \sum_{s=0}^{\infty} a_2^{s+1}(\mathbb{E}_2\mathbb{E}_1)^s\mathbb{E}_2[Q] + \sum_{s=0}^{\infty} a_2^s(\mathbb{E}_2\mathbb{E}_1)^s[Q] - \sum_{s=0}^{\infty} a_2^s\mathbb{E}_1(\mathbb{E}_2\mathbb{E}_1)^s[Q]\right).$$

Hence, combining this equation with (B.44) and using the observations that $\log p_1 = a_1(\varphi_1 - \varphi_2) + a_1\log p_2 + (1 - a_1)\log w$ and $\varphi_1 = 0$, we get the following expression for

$\log p_1$:

$$\log(p_1/m) = a_1 \left(\sum_{s=0}^{\infty} a_2^s \mathbb{E}_1 (\mathbb{E}_2 \mathbb{E}_1)^s [Q] - \sum_{s=0}^{\infty} a_2^{s+1} (\mathbb{E}_1 \mathbb{E}_2)^{s+1} [Q] \right) + (1 - a_1) \log(w/m).$$

Rearranging terms, we get

$$\log(p_1/m) = a_1 \sum_{s=0}^{\infty} a_2^s (\mathbb{E}_1 \mathbb{E}_2)^s \mathbb{E}_1 \Big[Q - a_2 \mathbb{E}_2[Q] \Big] + (1 - a_1) \log(w/m).$$

On the other hand, note that (B.45) implies that $Q - a_2 \mathbb{E}_2[Q] = (1 - a_2) \log(w/m)$. Therefore,

$$\log(p_1/m) = a_1 (1 - a_2) \sum_{s=0}^{\infty} a_2^s (\mathbb{E}_1 \mathbb{E}_2)^s \mathbb{E}_1 [\log(w/m)] + (1 - a_1) \log(w/m).$$

By Lemma 1, the assumption that $m < \bar{w}/\chi$ implies that $w = \bar{w}$, in which case the above equation immediately reduces to (2.28). Furthermore, noting that $\log m = \log(PC)$ then establishes (2.27). $\square$

APPENDIX C

Appendix to Chapter 3

C.1. Summary Statistics

Table C.1. Summary Statistics for the Baseline Calibration

Variable	**Statistic**	Value
agg-borrowing	**mean**	203.885
	std	552.350
agg-equity	**mean**	859295.888
	std	424812.277
agg-expected-demand	**mean**	1448.483
	std	127.743
agg-forecast-error	**mean**	36.758
	std	31.498
agg-inventories	**mean**	2394.449
	std	1453.078
agg-sales	**mean**	1473.495
	std	112.720
average-wage	**mean**	39.438
	std	14.356
gdpgrowth	**mean**	0.001
	std	0.058
goods-clearing-rate	**mean**	0.008
	std	0.000
inflation	**mean**	0.005
	std	0.021
kurtosis-average-cost	**mean**	3.543
	std	4.886
kurtosis-markups	**mean**	0.995
	std	2.290
kurtosis-pgs	**mean**	1.719
	std	3.091
kurtosis-sales	**mean**	4.787
	std	8.058
log-output	**mean**	7.317
	std	0.063
marginal-wage	**mean**	36.688
	std	12.968
mean-forecast-error-inductive	**mean**	21.885
	std	36.901
mean-forecast-error-simple	**mean**	35.983
	std	32.799
output-gap	**mean**	-0.013

Table C.2. Summary Statistics for the Baseline Calibration, continuing

Variable	Statistic	Value
output-gap	**std**	0.079
price-index	**mean**	27.014
	std	9.500
prodgrowth	**mean**	0.000
	std	0.031
real-production	**mean**	1508.280
	std	92.794
relative-performance-smart-firms	**mean**	0.053
	std	0.092
rolling-output	**mean**	7.330
	std	0.039
share-financial-constraint	**mean**	0.000
	std	0.000
skewness-average-cost	**mean**	-0.155
	std	1.404
skewness-markups	**mean**	0.215
	std	0.530
skewness-pgs	**mean**	-0.511
	std	0.943
skewness-sales	**mean**	1.442
	std	1.531
stddev-average-cost	**mean**	8.899
	std	4.291
stddev-markups	**mean**	25.660
	std	3.138
stddev-pgs	**mean**	5.174
	std	2.570
stddev-sales	**mean**	41.239
	std	17.455
unemployment-change	**mean**	-0.004
	std	0.390
unemployment-rate	**mean**	9.140
	std	5.590
vacancy-rate	**mean**	15.288
	std	7.360
wage-inflation	**mean**	0.005
	std	0.012

Vita

Thomas Pellet is part of the Data Science and Technologies team at Bloomberg L.P. and was a PhD student at Northwestern University from 2019 to 2023 under the supervision of Martin Eichenbaum, Matt Rognlie and Alireza Tahbaz-Salehi.

Pellet earned his bachelor's degree in political science from Sciences Po, Paris, and his bachelor's degree in mathematics from Sorbonne University, Paris, with an exchange year scholarship from the University of Chicago. He earned his master's degree in the joint "Grandes Écoles" program from HEC Paris and Sciences Po Paris.

Prior to his tenure as PhD candidate, Pellet was a research consultant at the World Bank from 208 to 2019 on issues related to internet access in Africa. Pellet was also a research analyst at the Peterson Institute for International Economics from 2017 to 2018. He worked with C. Fred Bergsten Senior Fellow Olivier Blanchard on issues related to the euro area, macroeconomic policy, and capital flows. Previously, Pellet worked as a research assistant to the chief economist at Rothschild & Co., in Paris, France and produced debt sustainability assessments for the euro area, as well as global financial stability reports and economic forecasts for executive committees.