THE TRANSMISSION OF COMMODITY PRICE SUPER-CYCLES

FELIPE BENGURIA [†]   FELIPE SAFFIE [‡]   SERGIO URZUA[§]

JULY 22, 2023

ABSTRACT

We examine two key channels through which commodity price super-cycles affect the economy: a *wealth* channel, through which higher commodity prices increase domestic demand, and a *cost* channel, through which they induce wage increases. By exploiting regional variation in exposure to commodity price shocks and administrative firm-level data from Brazil, we empirically disentangle these transmission channels. We introduce a dynamic model with heterogeneous firms and workers to further quantify the mechanisms and evaluate welfare. A counterfactual economy in which commodity booms are purely endowment shocks experiences only 30% of the intersectoral labor reallocation between tradables and nontradables, and 40% of the within-tradables labor reallocation between domestic and exported production. Finally, the consumption–equivalent welfare gain of a commodity super–cycle is twice as large in the counterfactual economy.

Keywords: Commodity shocks, local labor markets, skill premium, heterogeneous firms.

JEL classification: E32, F16, F41

[†] Department of Economics, Gatton College of Business and Economics, University of Kentucky. Email: felipe.benguria@uky.edu.

[‡] Darden School of Business, University of Virginia. Email: saffieF@darden.virginia.edu

[§] Department of Economics, University of Maryland and NBER. Email: urzua@econ.umd.edu

This paper was previously circulated under the title "Commodity Shocks, Firm-level Responses and Labor Market Dynamics". We are thankful to the editor and several anonymous referees. We are also thankful to our discussant, Philippe Martin, and to seminar and conference participants at the 2017 NBER Summer Institute, the 2017 Meetings of the Society of Economic Dynamics, the 2017 WEA meetings, Boston College, the Bank of Spain, the Central Bank of Chile, Drexel University, Emory University, the Federal Reserve Bank of Dallas, the Federal Reserve Bank of New York, the Federal Reserve Bank of Philadelphia, the Federal Reserve Bank of San Francisco, the Federal Reserve Board of Governors, Indiana University, the International Monetary Fund, North Carolina State University, Santa Clara University, the Small Business Administration, the University of Florida, the University of Houston, the University of Illinois at Urbana–Champaign, the University of Kentucky, the University of Rochester, and the University of Virginia. We benefited from comments and suggestions from George Alessandria, Boragan Aruoba, Javier Bianchi, Andres Fernandez, Sebnem Kalemli-Ozcan, Eunhee Lee, Julian Messina, Nuno Limao, Liliana Varela, Marcelo Veracierto and Kei-Mu Yi. We thank the Brazilian Labor Ministry (MTE) for data access. We recognize excellent research assistance by Camila Galindo, Rodrigo Heresi, Karam Jo and Eugene Oue. All mistakes are our own.

1 Introduction

Global commodity prices are a fundamental source of volatility for emerging markets. Commodity price super-cycles evolve at lower frequencies than regular business cycles with booms often lasting more than a decade. These super-cycles explain a third of the variance in output [Fernández et al., 2017b]. The literature has typically abstained from micro-founding their transmission channels to the economy. In this paper, we fill this gap by empirically identifying two key transmission channels of commodity super-cycles and studying their effects on labor reallocation and aggregate outcomes.

Firm and sector heterogeneity play a central role in our strategy to identify these transmission channels. First, during a commodity price boom, the increase in wealth leads to higher spending, and the response to this *wealth* channel is heterogeneous as it benefits more firms that sell their output locally (nonexporters and firms in low-tradability industries). Second, a commodity price boom induces a *cost* channel, operating through an increase in wages in response to an expansion in the labor demand of the commodity sector. Because the commodity sector in our setting is relatively unskilled-intensive, a commodity price boom lowers the skill premium, increasing the relative cost of less skill-intensive industries. While these two channels can be traced back to the characterization of the Dutch disease by Corden and Neary [1982], ours is the first study to introduce this heterogeneity and to exploit it to identify the existence of these channels using granular data for all sectors in the economy.[1]

Our empirical results are drawn from the experience of Brazil during 1999-2013, which is an ideal setting for understanding the mechanisms behind the transmission of commodity price cycles. During this period, the country faced large commodity price fluctuations. This economy has been historically characterized by substantial regional differences in the composition of commodity production. We exploit this regional variation and detailed administrative data for millions of firms to empirically identify the two channels. We motivate our identification strategy with a simple conceptual framework based on Jones [1975], Kovak [2013] and Dix-Carneiro and Kovak [2015], which delivers three testable results. First, an employment–weighted average of the prices of individual commodities is a sufficient statistic for the regional skill premium and employment

---

[1] Our cost channel is equivalent to Corden and Neary [1982]'s resource movement effect. In line with Alberola and Benigno [2017], we extend the original spending channel to allow for consumption smoothing with an endogenous current account and denote it wealth channel. The term 'Dutch disease' refers to the adverse effect of a boom in a resource sector on other sectors of the economy.

outcomes. Second, as a result of the cost channel, regions with a larger increase in this regional commodity price index will experience higher employment growth in high–skill relative to low–skill nontradable (i.e., service) industries. Third, as a result of the wealth channel, regions with a larger increase in the regional commodity price index face higher employment growth among firms selling their goods locally (nonexporters and firms in low–tradability industries) relative to exporters within the tradable (i.e., manufacturing) sector. We find strong evidence in favor of both channels in the data. We consider several alternative mechanisms as well as confounding shocks that may have affected manufacturing or service firms during this period. These include regional trends of trade liberalization, regional variation on import penetration from China, exchange rates dynamics, and intersectoral productions linkages, among several others. The magnitude and statistical significance of our main channels are robust to these other forces.

While these empirical results prove the existence of the cost and wealth transmission channels, they are silent about their aggregate implications. To assess these implications, we apply the intuition from our simple conceptual framework to a quantitative and dynamic general–equilibrium model. We introduce a three-sector (commodity, manufacturing, and services) small open economy model, which extends the quantitative international finance literature that studies commodity cycles (e.g., Alberola and Benigno [2017]) along two main dimensions. First, we deviate from the representative firm paradigm by including heterogeneous firms. Within the manufacturing sector, an export productivity cutoff divides firms into exporters and nonexporters, which allows us to capture the differential exposure to the wealth channel. Second, because the cost channel operates through differences in skill intensity across sectors, our model allows for heterogeneous workers (skilled and unskilled labor). The model is solved dynamically, taking into account the consumption smoothing behavior of the representative household, featuring endogenous trade balance and current account dynamics that are critical to account properly for the wealth channel.

We calibrate our model to the Brazilian economy. To discipline it, we explicitly target the heterogeneous skill intensities across sectors, the fraction of exporters in manufacturing, and several other firm-level and macroeconomic moments. We test the validity of the model by contrasting it to nontargeted outcomes. The model is able to replicate the dynamics of the regional skill premium, the reallocation of labor between sectors, and the differential evolution of employment in exporting versus nonexporting firms in response to a commodity price super-cycle. With the calibrated model, we perform a general equilibrium counterfactual to assess the importance of the labor market as a transmission mechanism. A counterfactual economy without transmission through the labor market (in which commodity booms are purely endowment shocks) sees 30% of the intersectoral labor reallocation between tradables and nontradables found in the baseline economy, and 40% of the within-tradables labor reallocation between domestic and exported production. From an aggregate perspective, the consumption–equivalent welfare gain of a commodity super–cycle is twice as large in the counterfactual economy relative to the baseline economy. Therefore, while the literature has frequently understood commodity price cycles as pure endowment shocks, our results show that incorporating contagion through the labor market is paramount. These quantitative results are robust to extending the model to a multiregion version with interregional trade and migration, and to using different modeling assumptions and parametrizations.

Related Literature. Our paper builds on the long-standing literature in international finance seeking to understand the impact of terms of trade on macroeconomic aggregates in emerging economies [Mendoza, 1995, Kose, 2002]. In particular, we contribute to the recent efforts to study and quantify the economic consequences of commodity price super-cycles [Reinhart et al., 2016, Shousha, 2016, Fernández et al., 2017a, Drechsel and Tenreyro, 2017, Alberola and Benigno, 2017, Fernández et al., 2020, Kohn et al., 2021].

On the empirical side, we identify the classic cost and wealth transmission channels proposed by Corden and Neary [1982]. To do so, we combine a novel empirical strategy, a quasi-natural experiment, and rich administrative microdata. Our empirical strategy extends previous studies that have focused on cross-country evidence [Sachs and Warner, 1995, Harding and Venables, 2016, Fernández et al., 2017b] and relies instead on variation in exposure to commodity price fluctuations across regions within a country. In addition, our identification strategy is based on the new insight that commodity price fluctuations have *heterogeneous* impacts across sectors and across firms. It is this heterogeneity, which we incorporate to the Dutch disease framework, that is key to the identification of the transmission mechanisms. A second useful element of our empirical work is the context. We study the large boom of the 2000s, during which Brazilian regions experienced different commodity prices due to differences in their economic structure. Finally, the microdata we use have the novelty of covering not only the manufacturing sector, but also the service and commodity sectors, left out in previous work on this topic but where most individuals are employed. Our identification strategy and microeconomic approach connects our work to the literature examining the consequences of local economic shocks [Autor et al., 2013, Dix-Carneiro and Kovak, 2017, Adão, 2015, Costa et al., 2016]. Closer to our study, Allcott and Keniston [2018] find a positive impact of local oil and gas booms in the U.S. on the manufacturing sector. In their paper, the positive effect on the manufacturing sector is driven by locally-traded subsectors and those with upstream linkages to oil and gas. We find a similar pattern in the manufacturing sector in the Brazilian data and show that our main results are robust to controlling for these mechanisms. Also related to our work, Faber and Gaubert [2019] study the impact of the expansion of tourism in Mexico on local manufacturing.

Our quantitative model features the same heterogeneity that enables our empirical identification. The international finance elements of our model are mostly related to Alberola and Benigno [2017]. They extend the framework of Corden and Neary [1982] to a dynamic multisector model with representative firms featuring endogenous growth to study structural long-run changes driven by temporary commodity booms. We abstract from endogenous growth and extend their framework by allowing heterogeneous firms and selection into exporting in the spirit of Ghironi and Melitz [2005] and Alessandria and Choi [2007].[2]

The rest of the paper is structured as follows. Section 2.1 describes the data used in our empirical analysis. Section 2.3 documents the empirical findings that guide our theory. Section 3 outlines the model, and Section 4 describes the results of our quantitative analysis. Section 5 concludes.

2 Empirical Analysis

2.1 Data Description

Our empirical analysis is carried out using linked employer-employee data from Brazil for the period 1999-2013. This dataset –Relação Anual de Informações Sociais (RAIS)– is an administrative census collected for social security purposes by Brazil's Labor Ministry.[3] It encompasses the

---

[2] More broadly, our paper belongs to the recent literature that uses dynamic heterogeneous-firms models and/or firm-level information to study classical problems in international economics such as the consequences of financial integration [Gopinath et al., 2017, Varela, 2015], the response of trade flows to devaluations [Alessandria et al., 2013], sudden stops [Ates and Saffie, 2021], and the effect of news shocks arising from commodity discoveries [Arezki et al., 2017].

[3] Differently than recent papers that have used worker-level records from these data [Alvarez et al., 2018, Dix-Carneiro and Kovak, 2017, Helpman et al., 2017], we use it to construct a comprehensive dataset of millions of firms in all sectors of the economy. No other such data exists for Brazil.

universe of formal-sector employees. The longitudinal information allows us to track both workers and firms over time. We observe detailed worker characteristics including educational attainment, age, gender, and occupation. The data also report each firm's industry and geographic location at the municipality level.[4] This allows us to exploit regional variation in exposure to commodity price changes in our empirical analysis.

For each job spell in each year, we observe the exact starting and ending day of a worker's employment in a firm. Based on this information, we compute quarterly indicators of employment. The data report individuals' mean monthly earnings in each job spell in each year, as well as earnings each December. In our analysis below, we use both annual earnings and monthly December earnings.

Regional Units. We define regional units that are a close approximation to the concept of local labor markets. While in the raw data we observe the municipality of each worker's employer, these municipalities are geographically too small to be considered a local labor market. Instead, we use 558 microregions as the regional unit that best reflects a local labor market. Throughout the paper we refer to microregions simply as regions. These units are defined by Brazil's Statistical Institute (IBGE) and are similar to U.S. commuting zones.[5] They are large enough such that commuting across them is fairly small. Aggregating the individual-level data, we compute sectoral and total regional employment, and regional measures of the skill premium.

Economic Sectors. The three sectors of interest to our paper are commodity, tradable, and nontradable. The commodity sector includes agriculture, mining, and fuels. The tradable sector includes all manufacturing industries. Nontradables include retail and wholesale trade, hospitality, construction, transportation, finance and real estate, among others.[6]

Heterogeneity in Skill Intensity. We divide workers into skilled and unskilled categories based on their educational attainment. The raw data records nine groups of attainment ranging from no formal education to tertiary education. We classify as unskilled workers those with at most

---

[4] Industry categories follow Brazil's National Classification of Economic Activities (CNAE).

[5] Commuting zones have been widely used as approximations to local labor markets in the US for the analysis of regional shocks, such as in Autor et al. [2013]. Brazil is divided into about 5500 municipalities, 558 microregions, 137 mesoregions, 27 states, and 5 macroregions.

[6] We exclude from our analysis government and quasi-public sectors. This includes federal and local governments, education and healthcare. We also exclude domestic service.

complete secondary education, and as skilled workers those with at least some tertiary education. Nationally, unskilled workers represent about 80 percent of total employment.

Differences in skill intensity across industries are key to our identification strategy, so we start by describing the distribution of skill intensity across industries within the commodity, tradable, and nontradable sectors. This allows us to characterize both the relative ordering of sectors as well as the variation in skill intensity within sectors. We measure an industry's skill intensity as the share of skilled workers in total employment using cross-sectional data in year 2000.[7] The distribution of skill intensity across industries for each sector is displayed in Figure 1 and moments of this distribution are reported in Appendix Table A3. Two observations emerge. First, the commodity sector is the most unskilled-intensive, followed by tradables and then nontradables. Second, while the distribution of skill intensity in commodity industries (and to a lesser extent within tradables) is fairly homogeneous, it varies substantially within the nontradable sector. This wide variation within the nontradable sector is exploited in our subsequent analysis to identify the transmission of commodity price cycles.

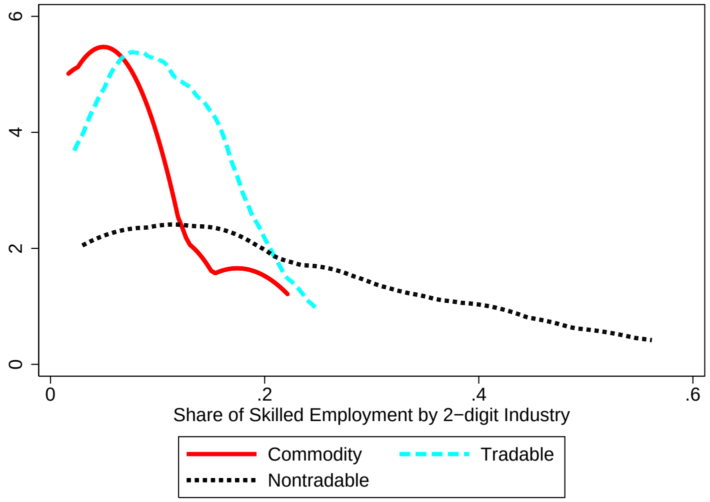

Figure 1: Skill Intensity Across Industries

Note: This figure displays kernel density estimates of the distribution of skill intensity (measured as the share of skilled workers in year 2000) of two-digit industries in the commodity, tradable, and nontradable sectors.

Firm-Level Panel. We construct a firm-level panel by aggregating our worker-level records. The panel consists of 2.3 million firms operating in the three sectors of interest. We compute firms'

---

[7] We have verified that these statistics are stable over time. While we report these distributions using data from RAIS, data from Brazil's demographic census show similar patterns.

quarterly employment in the last month of each quarter.[8] We complement our firm-level panel with additional information on exporting and importing firms obtained from Brazil's Secretariat of Foreign Trade (SECEX). We assign each firm in our panel a dummy variable for exporting and importing status at an annual frequency.

Commodity Prices. We use data on 14 commodity goods that capture the majority of commodity employment in Brazil. These commodities span agriculture, mining, and fuel industries and are chosen based on the following criteria: i) we must be able to match these categories to employment data to construct regional employment weights, and ii) we must be able to match these categories to data on commodity prices in world markets. The list of 14 commodities consists of cereals, cotton, sugarcane, soybeans, citrus, coffee, cacao, bovine meat, ovine meat, poultry meat, coal, oil and gas, a basket of metallic minerals, and a basket of precious metals. The largest commodities in terms of employment are cereals, bovine meat, coffee, sugarcane, and soybeans. There is considerable variation in the geographic distribution of employment in these commodities across Brazil. There is also geographic variation in the share of employment in commodities in each region as is illustrated in Figure A1 in Appendix A.1.2.

We obtain commodity prices for the period 1999-2013 from the World Bank's Global Economic Monitor - Commodities. This dataset has the advantage of tracking prices for a wide number of commodities over a long period of time and with systematic criteria. These prices are in nominal U.S. dollars, which we deflate using the U.S. Consumer Price Index. The data are reported at a monthly frequency, and we construct quarterly price indices based on the last month of each quarter.

Additional Data Sources. We complement our data with information from multiple additional sources. These include Brazil's Demographic Census, data on international trade flows, import tariffs, Brazil's input-output table, data on interregional trade flows, and industry-level measures of tradability and income elasticities. All of these are described in Appendix A.1.1.

Descriptive Statistics. Appendix A.1.2 provides descriptive statistics on the variation of employment, number of firms, and share of exporters across regions and sectors. It also discusses the

---

[8] In the case of multiplant firms, we define a firm's sector as that in which the firm has a majority of employment. Similarly, we define a firm's region as that in which most of the firm's employment is concentrated.

geographic variation in the skill premium and regional specialization in commodities.

Having described our data sources, we now proceed to the empirical analysis.

2.2 Conceptual Framework

We first provide a simple analytical framework to guide our empirical analysis. The quantitative model in Section 3 relaxes many of the assumptions used here. We use a specific factors model along the lines of Jones [1975], Kovak [2013] and Dix-Carneiro and Kovak [2015].[9] Consider a country with regions indexed by r and industries indexed by i. Production requires skilled and unskilled labor, and a specific factor. Each type of labor is mobile across industries and immobile across regions. The specific factor is immobile across industries and regions. There is perfect competition in all goods and factor markets.

The economy is comprised of I industries such that $\mathcal{I} = \{i = 1, \ldots, I\}$. The commodity sector is the subset of industries $C = \{i = 1, \ldots, I - 3\}$. The nontradable sector has only 2 industries ($I - 2$ and $I - 1$), and the tradable sector has a single industry (I). This is the simplest setup under which we can derive our testable implications.

For simplicity and without loss of generality, we assume the cost share of the specific factor, θ_R, is common across all industries. Skill intensities, on the other hand, vary across industries following the patterns described earlier in Figure 1. The commodity sector is the most unskilled intensive, with cost shares $(\theta_u^C,\ \theta_s^C)$. We assume these are constant across industries within the commodity sector, which is consistent with the small amount of variation in skill intensities across commodity industries documented in Figure 1. The tradable sector, with cost shares $(\theta_u^T,\ \theta_s^T)$, is skill intensive relative to the commodity sector. Finally, consistent with the larger amount of variation in skill intensity within the nontradable sector in Figure 1, we assume there is a low–skill nontradable industry with cost shares $(\theta_u^{N\ell},\ \theta_s^{N\ell})$ and a high–skill nontradable industry with cost shares $(\theta_u^{Nh},\ \theta_s^{Nh})$. To be able to compare industries based on their tradability (keeping everything else constant) we assume that the low–skill nontradable industry has the same skill intensity as the tradable industry:

$$\frac{\theta_s^C}{\theta_u^C} < \frac{\theta_s^T}{\theta_u^T} = \frac{\theta_s^{N\ell}}{\theta_u^{N\ell}} < \frac{\theta_s^{Nh}}{\theta_u^{Nh}} \tag{1}$$

Commodity output is sold in Brazil and the rest of the world, with prices that are exogenous to

---

[9] This type of model also motivates the empirical work in Dix-Carneiro and Kovak [2017].

the domestic economy. Output of the tradable industry is exported and faces a foreign demand with finite price elasticity. Prices in each nontradable industry are endogenous and clear the domestic market. In line with the literature, hats represent proportional changes.[10] We omit region subscripts for the sake of clarity. Our key results are the following.

Proposition 1: The change in the regional skill premium $\hat{w}_s - \hat{w}_u$ can be written as $\hat{w}_s - \hat{w}_u = \left(\sum_i \kappa_i \lambda_i \hat{P}_i\right)$, where $\lambda_i = \frac{L_i}{L}$ is the employment share of each industry in each region and $\hat{P}_i$ is the change in the price of industry i. Given that price changes in the tradable and nontradable industries can themselves be written as a function of commodity prices $\Big(\hat{P}_T = \left(\sum_{i \in C} \zeta_i^T \hat{P}_i\right)$, $\hat{P}_{Nh} = \left(\sum_{i \in C} \zeta_i^{Nh} \hat{P}_i\right)$, and $\hat{P}_{N\ell} = \left(\sum_{i \in C} \zeta_i^{N\ell} \hat{P}_i\right)\Big)$, we can approximate the change in the skill premium as $\hat{w}_s - \hat{w}_u \approx \kappa_C \left(\sum_{i \in C} \lambda_i \hat{P}_i\right)$, with $\kappa_C < 0$ due to the fact that commodities are unskilled intensive relative to the rest of the economy.

Proof: See Appendix C.

The mechanism behind this result is that an increase in commodity prices leads to an expansion in the commodity sector's labor demand. Because the commodity sector is relatively unskilled-intensive, this implies an increase in the relative demand for unskilled labor, and consequently, a decline in the skill premium. Thus, this proposition leads to the following testable implication:

Testable Implication 1: Regions with a larger increase in a regional commodity price index $\left(\sum_{i \in C} \lambda_i \hat{P}_i\right)$ will experience a larger decline in the skill premium $(\hat{w}_s - \hat{w}_u)$.

Proposition 2 (cost channel): We can write the relative change in employment between the high– and low–skill industries in the nontradable sector as: $\hat{L}_{Nh} - \hat{L}_{N\ell} = \left(\eta_U^{Nh} - \eta_U^{N\ell}\right)(\hat{w}_S - \hat{w}_U)$, where η_U^{Nh} and $\eta_U^{N\ell}$ are measures of the skill intensity in the high– and low–skill industries in the nontradable sector, respectively. Using Proposition 1, we can write $\hat{L}_{Nh} - \hat{L}_{N\ell} \approx \left(\eta_U^{Nh} - \eta_U^{N\ell}\right) \kappa_C \left(\sum_{i \in C} \lambda_i \hat{P}_i\right)$.

Proof: See Appendix C.

The increase in wages during a commodity boom represents a negative cost shock for all industries. Given that the unskilled wage increases more than the skilled wage, this cost shock is larger for unskilled-intensive industries, which in consequence face lower employment growth. This proposition leads to the following testable implication:

---

[10] Specifically, the proportional change in a variable x is denoted as $\hat{x}$ and is equal to $\frac{dx}{x}$, considering a small change dx.

Testable Implication 2: Regions facing higher relative growth in a regional commodity price index $\left(\sum_{i\in C}\lambda_i\hat{P}_i\right)$ will experience higher employment growth in high–skill relative to low–skill nontradable industries.

Proposition 3 (wealth channel): Let $\hat{L}_T$ be the employment change in the tradable industry. We can write the relative change in employment between the tradable industry and the low–skill nontradable industry (which by assumption have the same skill intensity) as: $\hat{L}_T - \hat{L}_{N\ell} = -\hat{m}$, where $\hat{m}$ is the change in the region's income. We can write the change in income as $\hat{m} = \sum_i \varphi_i\hat{P}_i$, where φ_i is the share of regional production in each industry. In consequence, $\hat{L}_T - \hat{L}_{N\ell} \approx \kappa_C\left(\sum_{i\in C}\lambda_i\hat{P}_i\right)$.

Proof: See Appendix C.

A positive commodity price shock triggers a wealth effect increasing the local demand for goods.[11] This wealth channel benefits producers of nontradables, which sell their goods locally. Since by assumption the tradable industry and the low–skill nontradable industry in the model have the same technology, we can interpret them as representing the export–oriented and domestically–oriented segments of a same manufacturing industry in the data. This proposition then leads to a final testable implication:

Testable Implication 3: Regions facing higher relative growth in a regional commodity price index $\left(\sum_{i\in C}\lambda_i\hat{P}_i\right)$ will experience higher employment growth among nonexporters relative to exporters within the manufacturing sector.

2.3 Empirical Characterization of Transmission Channels

Based on the testable implications of our simple framework, we now empirically identify the cost and wealth channels that transmit changes in commodity prices through the economy. But first, we discuss the construction of a regional commodity price index.

Regional Commodity Price Index. Following Proposition 1, we can define a regional commodity price index which drives labor market outcomes. This is defined as the weighted average of individual commodity prices. In each region, the weights are the base-period shares of employment

---

[11] The quantitative model in the next section features consumption smoothing with an endogenous current account in a dynamic setting, and thus we refer to a wealth effect.

in each individual commodity over total employment in the commodity sector. Formally,

$$p_{rt} = \frac{\sum\limits_{i \,\epsilon\, C} p_{it} \times e_{ir}}{\sum\limits_{i \,\epsilon\, C} e_{ir}}, \tag{2}$$

where p_{it} stands for the price of commodity i in period t, and e_{ir} represents the base-year employment of commodity i in region r.[12] We obtain employment data for these weights from Brazil's Demographic Census of year 2000, the earliest available year that is close to the start of our sample period. By using the Demographic Census for these weights instead of RAIS, we are using a wider measure of employment, including informal employment. This regional commodity price index is similar to the one used by Adão [2015], who exploits it for the different goal of quantifying the effects of global shocks on between– and within–group inequality.

Commodity super-cycles are low-frequency phenomena where most of the time series variation is explained by medium to long-run movements. Therefore, to better capture super-cycles, we use the trend of this price index rather than the short run deviations from it. We compute the logarithm of the price index in equation (2) and we extract the trend for each region's log price index using the Hodrick-Prescott filter, based on quarterly data for 1990-2015.

This regional price index varies across regions and time due to the interaction of (time-invariant) differences in the base-year composition of the commodity sector and the variation over time in the prices of individual commodities. Figure 2a depicts the evolution of the individual commodity price series over time. To a large extent, this period has seen a commodity price boom, but there is widespread variation, with for instance large growth in metal prices and fairly stable prices in cotton or sugarcane. At the national level, our 1999-2013 window captures falling prices up to mid-2002 followed by a commodity price boom with the price index at its peak in 2011 and the start of a bust in the last two years of our sample. This is shown in Figure 2b, which displays an aggregate commodity price index based on nation-wide employment weights. Figure 2c illustrates the variation in the residual regional price index after extracting region and time fixed effects, which is the source of our identification in the next section. The outer bars show the 10th and 90th percentiles of the residual price across regions in each quarter, and the inner bars show the 25th and 75th percentiles. The dashed (solid) line shows the path of the price for the

---

[12] For each region, there is a fraction of employment in commodity-related activities that we cannot directly link to one of the 14 commodities for which we have price data. In these cases, we distribute those workers equally across the commodities for which we do have price data found within the same 2-digit industry.

region with the price at the 10th (90th) percentile in the initial period.

We now proceed to examine our testable implications.

a) Individual Commodity Prices b) Aggregate Commodity Price Index c) Regional Commodity Prices

Figure 2: Commodity Prices

Note: The graph on the left shows the path of the prices in real U.S. dollars of the 14 commodities included in our commodity price index. These prices are normalized to one in 1990. The graph on the middle shows the path of the commodity price index for Brazil (dashed line) and its trend (solid line). The graph on the right shows the percentiles of the distribution of the residual regional commodity price index after extracting region and time fixed effects. The outer bars mark the 10th and 90th percentiles and the inner bars mark the 25th and 75th percentiles. The dashed (solid) line marks the path of the residual price for the region at the 10th (90th) percentile in the initial period.

Testable Implication 1: The Skill Premium and Commodity Prices. To establish the relationship between commodity prices and the regional skill premium, we estimate the following equation through OLS:

$$\text{SP}_{rt} = \pi_1 \cdot \text{Price}_{rt} + \gamma_r + \delta_{tm} + v_{rt} \ , \tag{3}$$

where SP_{rt} denotes the skill premium in region r in macroregion m at time t and v_{rt} is an error term. We use annual observations for each region during the period 1999-2013, with commodity prices measured in the last month of the year. The regional commodity price index (Price_{rt}) is the HP filtered trend of the log of the employment weighted index defined in equation (2). We measure the skill premium as the (log) difference in mean skilled and unskilled earnings in each region, for workers employed in the last month of the year. We report our results using both annual and monthly (December) earnings.[13] The regression includes region fixed effects as well as macroregion × year fixed effects. Observations are unweighted and standard errors are clustered by region.

---

[13] We estimate the regression with annual observations given the nature of how wage data are reported. As described in more detail in Section 2.1, we observe the annual earnings per job spell in each year as well as December earnings. Employment regressions below are estimated using quarterly data.

The results are reported in Table 1. Consistent with Testable Implication 1, the sign on the regional commodity price index is negative, such that a region experiencing a larger increase in the regional commodity price index will experience a larger decline in the skill premium.[14] Specifically, a region facing a one log point larger increase in the regional commodity price index experiences a 0.14 log point relative decline in the skill premium based on annual wages (column 1) and a 0.13 log point relative decline based on monthly wages (column 2). These are large elasticities; if a region faces an increase in commodity prices such as that seen in Brazil between the second quarter of 2002 and the last quarter of 2011 (a 55 log point trough-to-peak increase), it would experience a 7 log point relative decline in the skill premium (based on column 2). In Appendix A.1.4, we report the impact of commodity prices on the wages of unskilled and skilled workers separately, finding that the response of the skill premium is driven primarily by an increase in unskilled workers' wages.

Table 1: Commodity Prices and the Skill Premium

	(1)	(2)
Price_{rt}	-0.140**	-0.128***
	(0.057)	(0.058)
Observations	8363	7912
R^2	0.694	0.687

NOTE: This table reports the results of the estimation of equation (3). Column 1 corresponds to annual earnings. Column 2 corresponds to monthly (December) earnings. Each regression includes region as well as macroregion × time (year) fixed effects. Standard errors are clustered by region. ***, **, and * denote statistical significance at a 1, 5 and 10 percent confidence level, respectively.

Testable Implication 2: The Cost Channel. We now document that higher commodity prices lead to employment losses in unskilled relative to skilled industries. We estimate the following regression, linking firm–level employment to the regional commodity price index and its interaction with industry–level skill intensity:

$$\log(\text{Emp})_{ft} = \alpha_0 \cdot \text{Price}_{rt} + \alpha_c \cdot [\text{Price}_{rt} \times \text{Skill Intensity}_i] + \gamma_f + \delta_{st} + \epsilon_{ft} \,, \qquad (4)$$

---

[14] These elasticities do not allow us to infer aggregate effects, which are absorbed by the intercept or time fixed effects. This limitation is common to the empirical literature exploiting within–country variation to identify the effects of various shocks on relative regional outcomes [Autor et al., 2013, Faber and Gaubert, 2019].

where Emp_{ft} denotes employment of firm f in quarter t located in region r, Price_{rt} is the commodity price index of region r in period t, and Skill Intensity_i is the skill intensity of firm f's industry i measured in the baseline period. We include firm fixed effects (γ_f) as well as state × time (year–quarter) fixed effects (δ_{st}) to control for confounding time-varying shocks. Observations are unweighted and standard errors are clustered by region. We use the quarterly firm panel described in Section 2.1.[15] The parameter of interest is α_c. Testable Implication 2 entails $\alpha_c > 0$.

Note that following our conceptual framework, the wealth channel has a homogeneous impact across nontradable industries, so any variation in the impact of commodity prices on employment as a function of industries' skill intensity can be attributed to the cost channel. For this reason, and in line with Testable Implication 2, we estimate equation (4) for firms in the nontradable sector. This mechanism should still apply in the commodity and tradable sectors. Although in these two sectors we cannot necessarily isolate other mechanisms, we also report results for them for the sake of completeness.

Table 2 displays the results.[16] We find that in regions facing larger increases in the regional commodity price index, high–skill industries experience higher employment growth relative to low–skill industries, which is consistent with Testable Implication 2. For the nontradable sector (in column 1), the elasticity of firm employment to commodity prices for firms at the 25th and 75th percentiles of skill intensity is -0.039 and 0.005, respectively.[17,18] For the tradable and the commodity sectors (in columns 2 and 3, respectively), we also find a statistically significant and positive interaction term. Overall, our findings confirm that higher commodity prices lead to larger employment losses in relatively unskilled industries, a result consistent with the cost channel.[19]

Testable Implication 3: The Wealth Channel. The third testable implication predicts a gain

---

[15] Employment and commodity prices are measured in the last month of each quarter.

[16] Note that we demean the commodity price index and the measure of skill intensity.

[17] To benchmark this, note that if a region faces an increase in commodity prices such as that seen in Brazil between the second quarter of 2002 and the last quarter of 2011 (a 55 log point trough-to-peak increase), it would experience a 2.1 log point relative decline in employment for firms at the 25th percentile of skill intensity and a 0.4 log point relative increase for firms at the 75th percentile of skill intensity.

[18] Table A8 in Appendix A.1.5 reports the impact of the regional commodity price index on the employment in each sector without including interaction terms. We find that a one log point increase in the regional commodity price index is associated with a 0.13 log point increase and a 0.14 log point decline in employment for firms in the commodity and tradable sectors respectively. For firms in the nontradable sector, the impact is small and not statistically different from zero. The same appendix also provides an analysis of sectoral employment shares in each region.

[19] In Appendix Table A14 (column 2) we show these results hold when we identify them based on variation across three-digit industries within broader two-digit industries, which reduces concerns about unobserved industry characteristics shaping the response to commodity price shocks.

Table 2: Commodity Prices and Firm-Level Employment: The Cost Channel

	Nontradable Sector (1)	Tradable Sector (2)	Commodity Sector (3)
Price_{rt}	0.014	-0.136***	0.100*
	(0.021)	(0.030)	(0.056)
Skill Intensity_i x Price_{rt}	0.845***	0.795***	3.467***
	(0.114)	(0.205)	(0.476)
Observations	30,453,457	8,007,422	1,007,570
R^2	0.849	0.873	0.908

Note: This table reports the results of the estimation of equation (4) by sector. Each regression includes firm and state × time (year–quarter) fixed effects. Note we estimate equation (2) demeaning the commodity price index and skill intensity variables. Standard errors are clustered by region. ***, **, and * denote statistical significance at a 1, 5 and 10 percent confidence level, respectively.

in employment among nonexporters relative to exporters within the tradable sector. Exporter status, however, does not account for trade across regions within a country. For this reason, we provide further evidence in favor of this channel by also comparing industries with different degrees of tradability within the tradable sector, under the assumption that firms in low-tradability manufacturing industries are more likely to sell their output locally.

To identify the wealth channel, we estimate the following linear regression model:

$$\begin{aligned}\log(\text{Emp})_{ft} &= \beta_0 \cdot \text{Price}_{rt} + \beta_1 \cdot \text{Exporter}_{ft} + \beta_{w1} \cdot \left[\text{Exporter}_{ft} \times \text{Price}_{rt}\right] \\ &+ \beta_{w2} \cdot [\text{Tradability}_i \times \text{Price}_{rt}] + \gamma_f + \delta_{st} + u_{ft}\ , \end{aligned} \tag{5}$$

where Exporter_{ft} is a dummy variable taking a value of one if firm f operating in industry i exports in period t, Tradability_i is an industry-level measure of the extent to which a firm's output is tradable, and u_{ft} is the error term. We estimate (5) only for the tradable sector. We also control for firm and state × time fixed effects. As before, observations are unweighted and standard errors are clustered by region. The parameters of interest are β_{w1} and β_{w2}. The existence of a wealth channel (Testable Implication 3) implies $\beta_{w1} < 0$ and $\beta_{w2} < 0$.

We use a measure of industry tradability constructed by Holmes and Stevens [2014] and used in a similar context by Allcott and Keniston [2018].[20] Specifically, Holmes and Stevens [2014]'s measure, η, captures the transportation cost of different industries and is constructed based on data on shipment distances. For ease of interpretation, in our estimation of (5) we define $Tradability = -\eta$, such that highly tradable industries have low η values.

Table 3 displays the results for the estimation of equation (5). Column 1 includes only the interaction between the regional price index and the exporter dummy, and excludes the tradability measure. Given the well-documented link between exporting status and firm size, we additionally include interaction terms between firm–size bin dummies and the price index.[21] We find that in regions facing larger increases in the regional commodity price index, nonexporters experience higher employment growth relative to exporters, which is consistent with Testable Implication 3. Specifically, we find a much larger negative elasticity of firm employment to commodity prices for exporters (-0.10) than nonexporters (-0.042) in the tradable sector (see column 1). In column 2, we include the interaction term between the regional commodity price index and the industry-level tradability measure, and exclude the interaction of the price index with the exporter dummy. We find that the total elasticity of firm employment to commodity prices is -0.026 for firms at the 75th percentile of tradability and -0.007 at the 25th percentile. Finally, in column 3 we include both interaction terms, finding very similar coefficients to those in the previous two columns. These large and statistically significant differences illustrate the relevance of the wealth channel in the transmission of commodity price shocks.

Next, we verify the robustness of our findings to the existence of alternative transmission mechanisms and to other potentially confounding factors. We also provide further evidence in favor of the cost and wealth channels.

2.4 Alternative Transmission Mechanisms

Upstream and Downstream Linkages. A complementary transmission mechanism consists of upstream and downstream linkages between the commodity sector and other industries. This mechanism has been explored by Allcott and Keniston [2018] in the context of oil and gas booms in the U.S. We explore this mechanism in Appendix A.2.3. Based on Brazil's 2005 input-output table, we construct measures of these vertical linkages, categorizing manufacturing industries as

---

[20] Details on the tradability measure and summary statistics are provided in Appendix A.1.6.

[21] We include the interaction between the regional commodity price index and the following firm–size bin dummies: 5 to 50 workers, 50 to 100 workers, 100 to 500 workers, 500 to 1000 workers, and 1000 or more workers.

Table 3: Firm-Level Employment and Export Status in the Tradable Sector: The Wealth Channel

	(1)	(2)	(3)
Price_{rt}	-0.042*	-0.046**	-0.044**
	(0.022)	(0.022)	(0.022)
Exporter_{ft}	0.135***	0.135***	0.134***
	(0.005)	(0.005)	(0.005)
Exporter_{ft} x Price_{rt}	-0.058***		-0.055***
	(0.015)		(0.015)
Tradability_{i} x Price_{rt}		-0.042**	-0.039**
		(0.017)	(0.017)
Observations	8,086,964	8,005,174	8,005,174
R^2	0.909	0.909	0.909

Note: This table reports the results of the estimation of equation (5). Each regression includes firm and state × time (year–quarter) fixed effects. We include (but do not display) the interaction between the regional commodity price index and the following firm–size bin dummies (as well as the uninteracted dummies): 5 to 50, 50 to 100, 100 to 500, 500 to 1000, and 1000 or more workers. Standard errors are clustered by region. ***, **, and * denote statistical significance at a 1, 5 and 10 percent confidence level, respectively.

upstream, downstream, or unlinked to the commodity sector. We show that upstream and downstream linkages are relevant: an increase in commodity prices leads to an increase in employment in upstream industries, and a decline in employment in downstream industries. We also find that the coefficients estimated in equations (4) and (5) are robust to including interaction terms between the commodity price index and upstream and downstream dummy variables.

Interest Rates and Financial Conditions. Shousha [2016] and Fernández et al. [2017a] suggest that commodity cycles can affect the economy through an interest rate channel as spreads seem to be lower during booms and higher during busts. While we control for economy-wide shocks using state × time fixed effects, the pass-through of interest rates to firm-level employment could have a different impact on different firms. Although our data do not include financial information, the literature on financial frictions has documented that borrowing constraints are strongly tied to firm size [Gopinath et al., 2017]; small firms are more likely to be constrained and therefore less likely to benefit from the boom. Appendix Table A14 and Table 3, for the cost and wealth channels respectively, confirm our findings hold even after we control for the interaction between firm–size bin dummies and the regional commodity price index.

Political Economy. Commodity prices could potentially impact labor market outcomes through various political economy mechanisms. A first such channel could consist in increased revenue by local governments leading to increases in social transfers or a larger demand for employment associated with public works. Caselli and Michaels [2013] find a small (if at all) increase in social transfers and public goods provision to increases in oil-related revenue to regions. A second channel is that local governments could adjust the minimum wage in response to commodity prices. The following facts suggest our results are robust to these mechanisms. First, while oil resources might lead to local government revenue, other commodities do not. We find that our results are robust to excluding oil and gas from the commodity price index (Appendix A.2.4). Second, in Brazil the minimum wage is primarily defined at the national level. Although a handful of states have state-level minimum wages (that must be higher than the national-level minimum), Engbom and Moser [2022] argue that minimum wages vary little across states. Our firm-level regressions include state × time fixed effects, which would absorb changes in state-level minimum wages, alongside any other policies determined at this level of regional aggregation. Policies determined at more local levels are also controlled for with region × time fixed effects used in robustness checks.

2.5 Robustness Checks

Market Power in Commodity Markets. Our identification strategy for the cost and wealth channels in equations (4) and (5) does not rely on a lack of market power of the Brazilian economy (or specific Brazilian regions) on world markets for individual commodities. The reason is that these regressions exploit firm–level variation in the tradable (i.e., manufacturing) or nontradable (i.e., service) sectors. Thus, the identifying assumption is that unobserved shocks to the commodity sector that could impact world prices do not simultaneously have i) a differential impact on skilled–versus unskilled–intensive nontradable industries (in the case of the cost channel) or ii) a differential impact on exporting vs nonexporting firms in the manufacturing sector (in the case of the wealth channel).

Nevertheless, in Appendix A.1.7 we document that Brazil holds a relatively small world market share in each of the commodities used in the construction of this index. In addition, Appendix A.1.8 shows that employment in each of these commodities is regionally dispersed: in each commodity, even regions at the right tail of the distribution represent a small share of aggregate employment. Going further, Appendix A.2.4 shows that the regressions for the cost and wealth channels are robust to excluding the commodities in which Brazil has a larger world market share (soybeans, sugarcane, and coffee) from the regional commodity price index.[22] Furthermore, Appendix A.2.5 confirms that our results for the cost and wealth channel regressions are robust to excluding the states with larger world market shares of soybeans, sugarcane, and coffee. Thus, we use a sample in which no state has more than a 2% world market share in any commodity.

Another potential concern is that depending on the market structure of the commodity sector in Brazil, the prices received by commodity producers could differ from the world price. In Appendix A.1.9, we compare producer prices and world prices for important commodities in our sample, and find an extremely high correlation between them. Further, we document there are only small regional differences in the evolution of producer prices. Finally, we show that there is no concentration of employment among commodity producers, suggesting that these firms are price takers.

Trade Shocks. Brazil faced a large increase in imports from China in the period under study [Costa et al., 2016]. In Appendix A.2.6, we construct a measure of regional exposure to the increase in Chinese competition and show that our estimates are not confounded by this shock. Additionally, in the late 1980's and early 1990's, Brazil reduced its import tariffs substantially [Dix-Carneiro and Kovak, 2017]. While this occurred earlier than our sample period, the impacts could potentially be long-lasting. In Appendix A.2.7, we construct a measure of regional exposure to trade liberalization and show our estimates are robust to controlling for the effect of tariff liberalization.

Informality. While Brazil has a large share of informal labor, our firm-level data are limited to the formal sector of the economy. Because informality is more common among unskilled workers, a transition of informal workers toward the formal sector in response to an increase in commodity prices would bias our results against finding a decline in the skill premium. Thus, we interpret our results for the skill premium, and the consequent reallocation of labor, as lower bounds. Despite this point, in Appendix A.2.8 we construct regional measures of informal employment and show that our identification of the cost and wealth channels is not driven by regional variation in informality.

---

[22] We show equivalent results for the skill premium regression in Appendix Table A20.

Geographic Spillovers. We extend our empirical strategy to capture geographic spillovers through general equilibrium effects. These spillovers could occur due to trade linkages, migration, or other geographic linkages. Following Allcott and Keniston [2018], we assume that these spillovers are mostly local, meaning that a given region is only impacted by commodity price shocks to other regions within a certain radius. In Appendix A.2.4, we construct a commodity price index that takes into account the prices faced by nearby regions. The estimated impact of commodity prices on employment is very similar to that reported earlier, which suggests the magnitude of these spillovers is small. Moreover, in Appendix B.1 we develop a version of our quantitative model with interregional trade and migration to fully take into account these general equilibrium effects. The results from this model are similar to those from its more parsimonius version in Section 3.

Exchange Rates. Exchange rate fluctuations could impact employment of exporting and non-exporting firms in the tradable sector differentially. Several conditions would have to be met for exchange rate fluctuations to confound our estimates of the impact of regional commodity prices on firm-level employment in equation (5). While our regressions include time fixed effects that capture aggregate shocks, a spatial correlation between the share of exporters and the local composition of commodity baskets, however, could confound our results. Note that this spatial correlation would have to occur within states, as we include state × time fixed effects. We show in Appendix A.2.9 that this is not the case.

Regional Commodity Price Index: Trend vs. Cyclical Component. Our baseline empirical analysis is carried out using the trend of the regional commodity price index obtained with the Hodrick-Prescott (HP) filter. Tables A19 and A20 in the Appendix show our results are similar when using the raw price index instead of its trend. Moreover, by including both the trend and cyclical components simultaneously, we conclude that the trend of the price index is the source of variation driving our main findings.

Unobservable Regional or Region by Sector Shocks. Finally, we can modify equations (4) and (5) to include region × time (year–quarter) fixed effects. These control for any unobserved regional shock. In this case, the first term in each equation (the regional commodity price index by itself) is absorbed by these fixed effects. Further, we also include region × time × 2–digit industry fixed effects, absorbing regional shocks that vary by region and broad industries. Our regression for the cost channel is still feasible with these fixed effects given that we measure skill intensity at the 3–digit level. The results, in Appendix Sections A.2.1 (for the cost channel) and A.2.2 (for the wealth channel) show our results remain robust.

Standard Errors. Our approach consists of a shift-share design as we construct regional shocks weighting individual commodity price time series with initial regional employment shares. Adão et al. [2019] propose a correction to standard errors in this type of regression. In Appendix A.2.10, we report standard errors following this method. Because these corrected standard errors are smaller, we conservatively leave them to the Appendix and report the standard errors clustered by region in the main text.

2.6 Further Evidence in Favor of the Cost and Wealth Channels

Identification of the Wealth Channel based on Variation in Income Elasticities. An alternative approach to identifying the wealth channel could exploit variation across nontradable industries in income elasticities. If a commodity price increase leads to an increase in wealth and, consequently, an increase in the local demand for nontradables, this should lead to a relative expansion in industries with a higher income elasticity. Appendix A.2.11 shows that this is indeed the case. It also confirms that the identification of both the cost and wealth channels in equations (4) and (5), respectively, is robust to the inclusion of these income elasticities.

Importing Firms. Commodity prices could have a differential impact not only on exporters but also on importing firms. One would expect a positive impact of commodity prices on importers relative to nonimporters because importers purchase inputs instead of producing them, and are thus less exposed to the increase in labor costs induced by the cost channel. Appendix A.2.9 shows that this is the case. We also report that the negative coefficient on the interaction between the regional commodity price index and the exporter dummy variable is even larger in magnitude after controlling for an importer dummy variable. Finally, we also show that the response of importers to exchange rate movements does not confound our results.

3 Model

The stylized framework in the previous section allowed us to clearly lay out the conditions behind the identification of the cost and wealth channels. In this section, we develop a quantitative multisector small open economy model. The main differences with the previous framework are adding heterogeneous firms and a dynamic setting which allows for consumption smoothing with an endogenous current account. These features bring us closer to the data, and allow us to quantify the mechanisms involved in the transmission of commodity price super-cycles and to assess welfare.

The economy consists of commodity, nontradable (services), and tradable (manufacturing) sectors. In the commodity sector, a representative firm produces a homogeneous good using a scarce natural resource in fixed supply. The commodity price is exogenously determined in the world market. The commodity good can be exported. In the tradable and nontradable sectors, heterogeneous firms produce differentiated product varieties in a monopolistic competition setting.[23] Nontradable goods are consumed only domestically. Firms in the tradable sector can sell their varieties domestically or they can export subject to a fixed cost.[24] Firms in all sectors use both skilled and unskilled labor.

Skilled and unskilled labor are supplied inelastically by a representative household. This household demands the commodity good, the bundle of tradable varieties, the nontradable good, and an imported good; and it can borrow freely from the world.

We use the imported good as the numeraire. We index sectors by $k \in \{C, N, T\}$ (for commodity, nontradable, and tradable), skill types by $d \in \{U, S\}$ (for unskilled and skilled), and firms by f.

3.1 Representative Household

The representative household consumes imported (M), commodities (C^d), nontradable (N), and tradable (T) goods. It supplies skilled ($\bar{L}^s = (1-\kappa)\bar{L}$) and unskilled ($\bar{L}^u = \kappa\bar{L}$) labor inelastically.

---

[23] Because our results concern employment reallocation between the nontradable and tradable sectors, we impose the same market structure across them. In this way, the only differences between these two sectors are tradability and skill intensity in production. Note also that the two key differences between the commodity sector and the tradable (manufacturing) sector are the constant returns to scale (with a representative firm) in the commodity sector and that the commodity good is traded costlessly. These assumptions are immaterial to our results.

[24] In Appendix A.1.10 we document that fluctuations in regional commodity prices have an impact on firm entry and exit in the commodity, tradable and nontradable sectors, as well as on entry and exit in and out of exporting within the tradable sector. This motivates the inclusion of the entry and exit margins in the model.

Lifetime utility is:

$$U = \sum_{t=0}^{\infty} \beta^t \frac{(C_t)^{1-\nu}}{1-\nu}\,, \tag{6}$$

where $\nu > 1$ is the intertemporal elasticity of substitution and C_t is the final consumption bundle that is a nested-CES aggregate of M_t, C_t^d, N_t, and T_t:

$$C_t = \left\{ \left[\left[(M_t)^{\theta_1} + \alpha_T (T_t)^{\theta_1} \right]^{\frac{\theta_2}{\theta_1}} + \alpha_N (N_t)^{\theta_2} \right]^{\frac{\theta_4}{\theta_2}} + \alpha_C \left[C_t^d \right]^{\theta_4} \right\}^{\frac{1}{\theta_4}} . \tag{7}$$

The tradable bundle T_t and nontradable bundle N_t are both CES aggregators of individual varieties according to:

$$T_t = \left[\int_{\zeta \in \mathfrak{Z}_t} t_t(\zeta)^\rho d\zeta \right]^{\frac{1}{\rho}} \qquad \text{and} \qquad N_t = \left[\int_{\varphi \in \mathfrak{P}_t} n_t(\varphi)^\rho d\varphi \right]^{\frac{1}{\rho}} ,$$

where ζ and φ are indices for individual varieties and $\mathfrak{Z}_t$ and $\mathfrak{P}_t$ denote the respective sets of individual varieties sold at time t.

Let P_t^C be the exogenous price of the commodity good, P_t^R the price of the fixed natural resource, P_t^N the price of the nontradable bundle, and P_t^T the price of the tradable bundle. With the exogenous price of the imported good normalized to one and used as the numéraire, the representative household budget constraint becomes:

$$M_t + P_t^N N_t + P_t^T T_t + P_t^C C_t^d + B_{t+1} \leq w_t^s \bar{L}^s + w_t^u \bar{L}^u + P_t^R \bar{R} + (1+r) B_t + \Pi_t \,, \tag{8}$$

where bonds are in units of the imported good and the interest rate r is exogenous. Π_t collects total profits from firms in the commodity, nontradable, and tradable sectors.

Therefore, the representative household takes as given the sequence $\{P_t^C, P_t^N, P_t^T, P_t^R, \Pi_t, w_t^s, w_t^u\}$ and chooses $\{B_t, N_t, C_t^d, N_t, T_t\}$ to maximize lifetime utility (equation (6)) subject to the budget constraint (equation (8)). Finally, we use the Euler equation for bonds to define the household's stochastic discount factor $\Lambda_{t,t+1}$:

$$\Lambda_{t,t+1} = \beta \frac{\lambda_{t+1}}{\lambda_t} = \frac{1}{1+r} \,. \tag{9}$$

Here λ_t is the marginal utility of consumption in period t. When equation (9) holds, $\Lambda_{t,t+1}$ is constant. When the household is surprised by a new sequence of commodity prices revealed at period $\tilde{t}$, the Euler equation does not hold in that period and the marginal utility experiences a discrete jump before reaching a new constant level in $\tilde{t}+1$. In other words, the permanent income hypothesis holds implying that the household decides a new constant level of consumption immediately after a change in the present value of its income. This optimal re-balancing of consumption is critical for characterizing the wealth channel.

3.2 Production

Labor Input. Firm f in each sector $k \in \{C, N, T\}$ uses a composite labor input that combines skilled and unskilled labor according to the following CES function:

$$l_f^k\left(l_{f,t}^{u,k}, l_{f,t}^{s,k}\right) \equiv \left[(\phi^k)^{\frac{1}{\gamma}}(l_{f,t}^{u,k})^{\frac{\gamma-1}{\gamma}} + (1-\phi^k)^{\frac{1}{\gamma}}(l_{f,t}^{s,k})^{\frac{\gamma-1}{\gamma}}\right]^{\frac{\gamma}{\gamma-1}} , \tag{10}$$

where $l_{f,t}^{s,k}$ and $l_{f,t}^{u,k}$ are the skilled and unskilled labor inputs, respectively. We assume that the relative use of these two types of labor varies across sectors (ϕ^k), while the elasticity of substitution between skill types is constant across them. We can define the effective labor cost of sector k as:

$$\hat{w}_t^k = \left[\frac{\phi^k}{(w_t^u)^{\gamma-1}} + \frac{1-\phi^k}{(w_t^s)^{\gamma-1}}\right]^{-\frac{1}{\gamma-1}} . \tag{11}$$

Note that as the skill premium falls sectors with higher skill intensity (lower ϕ^k) see a decrease in their effective labor cost relative to low skill intensity sectors.

Commodity Sector. A representative firm produces the commodity good that can be sold domestically and exported internationally. To produce the commodity good, the representative firm combines the composite labor input L_t^C and a resource input R_t according to the following production function:

$$Y_t^C = \left(L_t^C\right)^{\eta} \cdot (R_t)^{1-\eta} .$$

The price of the commodity good, P_t^C, is determined exogenously in the international market. After the domestic demand is met, the remaining output is exported internationally.

Nontradable Sector. The nontradable sector is composed of a continuum of firms heterogeneous in productivity z each producing a differentiated variety φ under monopolistic competition. Each firm's permanent productivity type z is drawn upon entry from a log-normal distribution g with mean μ_N and standard deviation σ_N. A type-z firm produces its variety using the composite labor as the only input according to the following production function:

$$q_{Nt}\left(z\right) = z \cdot \left(l_{Nt}\left(z\right) - f_N\right) \ ,$$

where f_N is a fixed per-period production cost in terms of the labor input. In every period, all continuing firms are subject to an i.i.d. exit shock with probability δ_N. Potential entrants can enter the market by paying an entry cost consisting of a fixed and a convex component in terms of the composite labor input:

$$f_{Net} = f_{Ne} + \xi_{Ne}\left[\exp\left(\frac{\mathcal{M}_{Net} - \overline{\mathcal{M}}_{Ne}}{\overline{\mathcal{M}}_{Ne}}\right) - 1\right] \ ,$$

where $\mathcal{M}_{Net}$ is the mass of potential entrants that pay the entry cost to observe their permanent individual productivity. The convex component of the entry cost captures congestion externalities or competition for a fixed resource at entry.[25] After paying the entry cost, the productivity type of the firm is revealed.

Tradable Sector. The production technology of the tradable sector is identical to that of the nontradable sector. There is a continuum of firms heterogeneous in productivity z, each producing a differentiated variety ζ using the following production function:

$$q_{Tt}(z) = z \cdot \left(l_{Tt}(z) - f_d\right) \ ,$$

where f_d is a fixed per-period production cost in terms of the composite labor input. Akin to the nontradable sector, each firm's productivity type z is permanent and is drawn from a log-normal distribution g with mean μ_T and standard deviation Σ_T upon entry.

Unlike the nontradable sector, firms in the tradable sector can export subject to a per-period

---

[25] This assumption is standard in dynamic models of heterogeneous firms (see for example Jaef and Lopez [2014]). It is used to avoid corner solutions in the transition, and the size of the convex component of the cost is set to be small relative to the fixed component so its effects on dynamics are minimal.

fixed cost f_x expressed in terms of the composite labor input. The foreign demand for exports is:

$$q_{Tt}^{x}(z) = \Gamma \left[p_{Tt}(z)\right]^{-\sigma} ,$$

where Γ is exogenous and fixed throughout the commodity super-cycle. In each period, all continuing firms face an exogenous exit risk that occurs with probability δ_T that is i.i.d. across firms and time. Meanwhile, potential entrants can enter the market subject to a fixed entry cost:

$$f_{Tet} = f_{Te} + \xi_{Te} \left[\exp\left(\frac{\mathcal{M}_{Tet} - \overline{\mathcal{M}}_{Te}}{\overline{\mathcal{M}}_{Te}} \right) - 1 \right] .$$

3.3 Equilibrium

We now describe some key aspects of the equilibrium. The full system of equations and the definition of the equilibrium are deferred to Appendix D.1 and the steady state is described in Appendix D.2.

Entry and Exit in the Nontradable Sector. Firms are subject to an exogenous exit probability δ_N in every period. In addition, firms exit endogenously when the present discounted value of profits is negative. Letting optimal profits of a nontradable firm i be denoted by π_t, the value function of a firm with productivity z is:

$$W_t(z) = \max\left\{0, \pi_t(z) + \beta(1-\delta_N)W_{t+1}(z)\right\} .$$

The domestic productivity cutoff $\bar{z}_t$ is implicitly defined by the productivity level that solves $W_t(\bar{z}_t) = 0$.

In each period, potential entrants may pay the fixed entry cost to draw their permanent productivity levels z. If the realized draw is higher than the domestic cutoff, the firm will enter the market and produce. Potential entrants will pay to get their productivity draws until the present discounted value of entry is equal to the fixed entry cost:

$$\hat{w}_{Nt} \cdot f_{Net} = \int_{\bar{z}_t}^{\infty} W_t(z) g(z) dz \, .$$

Entry, Exit, and Exporting in the Tradable Sector. Entry and exit decisions in the tradable sector are identical to those in the nontradable sector. However, in the tradable sector, firms make the additional decision of whether or not to export.

Let $\pi_{m,t}$ denote the static profits associated with each sales mode $m \in \{d, x\}$ (selling domestically and exporting internationally, respectively). The value functions corresponding to each of these modes are:

$$\begin{aligned} W_{d,t}(z) &= \max\left\{0, \pi_{d,t}(z) + \beta(1-\delta_N)W_{t+1}(z)\right\} \text{ , and} \\ W_{x,t}(z) &= \max\left\{0, \pi_{d,t}(z) + \pi_{x,t}(z) + \beta(1-\delta_N)W_{t+1}(z)\right\} \, . \end{aligned}$$

The continuation value is given by: $W_{t+1}(z) = \max\{W_{d,t}(z), W_{x,t}(z)\}$. The domestic ($\bar{z}_{d,t}$) and export cutoffs ($\bar{z}_{x,t}$) are implicitly defined by $W_{d,t}(\bar{z}_{d,t}) = 0$ and $\pi_{x,t}(\bar{z}_{x,t}) = 0$. Finally, the free entry condition is:

$$\hat{w}_{Tt} \cdot f_{Tet} = \int_{\bar{z}_t}^{\infty} W_t(z) g(z) dz \, .$$

4 Quantitative Analysis

We calibrate the model with the goal of quantifying the mechanisms that propagate commodity price super-cycles throughout the economy. We use counterfactual simulations to determine the extent to which the labor market is a relevant channel of transmission to other sectors of the economy. This is relevant given the common assumption in the literature that commodity booms are purely endowment shocks. We proceed as follows. First, we describe the calibration of the model to the Brazilian economy. We then validate the calibration by simulating data from the model and replicating regressions results in Section 2.3 for outcomes which were not targeted in the calibration. Finally, we use the calibrated model to quantify transmission mechanisms.

4.1 Calibration

We calibrate two sets of parameters. The calibration is done at an annual frequency and uses data for the year 2001.[26]

---

[26] We choose 2001 because it is the year right before the start of the boom. The model is calibrated to the exact same sample used in the empirical analysis.

Externally Calibrated Parameters. Table 4 shows a first set of parameters that are either normalized or externally calibrated to values in the literature and to regional outcomes computed using the data described in Section 2.1. The discount factor (β) is consistent with an annual interest rate of 4%. The parameters governing risk aversion (υ), substitution between sectoral goods (θ's), substitution between tradable varieties (σ), and substitution between skilled and unskilled labor (γ) are taken from the international trade and international finance literatures. The exogenous sectoral exit rates (δ^N and δ^T) are computed directly from RAIS.

Because we study a small open economy that faces an exogenous interest rate, an initial condition for bond holdings is needed to solve the model. Without loss of generality, we assume that the economy starts with zero net foreign assets.[27] We normalize the total population of the economy to be 100. In addition, we set $\alpha_T = \alpha_M = 1$ and $P^l_T = P^l_M = 1$ such that importables and domestic tradables are symmetric in the initial steady state.

Table 4: Externally–Calibrated Parameters

Parameter	Description	Value	Source
r^*	Interest rate	0.04	Macroeconomic data
ν	Discount factor	1.5	Literature
σ	Elast. of substitution between varieties in T and N	4	Simonovska and Waugh [2014]
ρ	Inverse of markup in T and N	0.75	$\sigma = \frac{1}{1-\rho}$
γ	Elast. of subs. between skilled and unskilled labor	1.41	Katz and Murphy [1992]
η	Labor share in C	0.5	Gollin et al. [2014]
δ_T	Exogenous exit probability in T	0.14	Firm-level data (RAIS)
δ_N	Exogenous exit probability in N	0.18	Firm-level data (RAIS)
θ_1	Elast. of subs. between importables and domestic tradables	$\frac{1}{1-0.85}$	Corsetti et al. [2008]
θ_2	Elast. of subs. between tradables and nontradables	$\frac{1}{1-0.74}$	Mendoza [1995]
θ_3	Elast. of subs. between commodities and non-commodities	$\frac{1}{1-0.50}$	Buera and Kaboski [2009]
μ_T	Mean of productivity distribution in T	0	Normalization
$\bar{B}$	Steady state debt	0	Normalization
$\bar{H}$	Total inelastic labor supply	100	Normalization
α_T	Weight on T good in utility function	1	Normalization
P_T	Steady state price of T good	1	Normalization

Note: This table shows the values and sources of the externally-calibrated parameters. The abbreviations C, T and N stand for the commodity, tradable and nontradable sectors. Note that the elasticity of substitution between regional commodity goods is assumed to be the same than that between importables and domestic tradables.

---

[27] The model is solved using a shooting algorithm, in the spirit of Alberola and Benigno [2017], extended to our setting with heterogeneous firms. We show in Appendix B.3 that the quantitative results of the model are robust to alternative initial conditions or to including a bond–holding cost.

Internally Calibrated Parameters. Table 5 displays the set of 16 parameters that are internally calibrated. Twelve of these parameters are calibrated in the steady state. The initial price level of the commodity good (P_C) is used to match the commodity output share, while the weight on the nontradable good in the utility function (α_N) is used to target the nontradable output share. Because all commodity output in excess of domestic demand is exported, the weight on the commodity good (α_C) can be identified by the share of commodities in total exports. The fixed cost of exporting in the tradable sector (f_x) is calibrated to match the fraction of exporting firms. The firm size distributions in the tradable and nontradable sectors are used to discipline the fixed operating costs (f_d,f_N) and the variance of the productivity distributions (Σ_T,Σ_N). Having normalized μ_T (the mean of the productivity distribution in the tradable sector) to 0, the mean of the productivity distribution in the nontradable sector (μ_N) captures the relative differences in average size between the tradable and nontradable sectors. As such, we calibrate this parameter to the relative mass of firms between the two sectors. The size of foreign demand for exports (Γ) targets the export to GDP ratio. The endowment of the natural resource ($\bar{R}$) is used to match the commodity sector's share of total employment. We observe a close match between the model's moments and the target moments.[28]

The next set of parameters is calibrated to match perfectly macroeconomic ratios of the Brazilian economy. The share of unskilled labor in the model economy (κ) is set to generate the skill premium observed in Brazil. The skill intensities in each sector (ϕ_C, ϕ_T, ϕ_N) are chosen to replicate the share of unskilled labor in each sector.

The entry congestion elasticity (ξ_e) is calibrated using the transition over the commodity super-cycle. It is chosen such that at the peak of the boom, the convex component of the entry cost is 20% of the fixed component.

As Table 5 shows, the commodity sector in the Brazilian economy concentrates a modest share of firms and workers (commodity production represents only 12% of output). This means powerful internal amplification and propagation mechanisms must be present to generate the aggregate effects of a commodity price super-cycles documented in the literature.

Simulation. In the following exercises, we assume that the economy starts in a steady state in period 0. In period 1, a commodity price super-cycle is revealed to the agents. As Figure

---

[28] The only exception is the share of nontradable employment by the largest firms which can be attributed to the fact that the log normal productivity distribution used does not have fat tails.

Table 5: Internally–Calibrated Parameters

Parameter	Description	Moment Targeted	Moment (Data)	Moment (Model)	Param. Value
P_C	Commodity price	C share in GDP	0.14	0.13	26.86
α_N	Weight on N good in utility	N share in GDP	0.55	0.56	8.40
α_C	Weight on C good in utility	Share of C in exports	0.32	0.34	0.54
f_d	Production fixed cost in T	Share of T emp. in 20% smallest firms	0.02	0.03	0.06
f_x	Fixed cost of exporting in T	Fraction of exporting firms in T	0.09	0.08	1.40
Σ_T	St. dev. of productivity distribution in T	Share of T emp. in 10% largest firms	0.68	0.71	0.91
Γ	Foreign demand	Export to GDP ratio	0.15	0.19	281.16
$\bar{R}$	Resource endowment	C share in employment	0.08	0.08	3.63
f_N	Production fixed cost in N	Share of N emp. in 20% smallest firms	0.04	0.03	0.03
μ_N	Mean of productivity distribution in N	Relative mass of firms: M_T/M_N	0.27	0.27	0.25
Σ_N	St. dev. of productivity distribution in N	Share of N emp. in 10% largest firms	0.66	0.56	0.79
κ	Aggregate share of unskilled labor	Skill premium (w^s/w^u)	3.16	3.16	0.83
ϕ^C	Skill intensity in C	Shared of unskilled labor in C $(L^{u,C}/L^C)$	0.92	0.92	0.69
ϕ^N	Skill intensity in N	Shared of unskilled labor in N $(L^{u,N}/L^N)$	0.79	0.79	0.43
ϕ^T	Skill intensity in T	Shared of unskilled labor in T $(L^{u,T}/L^T)$	0.88	0.88	0.59
ξ_e	Entry congestion elasticity	Fraction of entry cost: $\max\{\frac{\text{cong. term}}{f_e}\}$	0.20	0.19	1

NOTE: This table shows the values of internally-calibrated parameters. The abbreviations C, T and N stand for the commodity, tradable and nontradable sectors.

3 shows, the price of the commodity good increases steadily for ten years (period identified by the darker gray area in each graph) and reaches a cumulative increase slightly above 70%. This trough-to-peak amplitude is obtained from the increase in commodity prices observed in Brazil between 2002 and 2011, as measured with employment-weighted prices at the national level. The commodity price then decreases steadily after the peak of the boom, reverting to the initial level after ten years. The price dynamic during the bust is extrapolated from the data in a symmetric fashion.

4.2 Model Validation

Prior to quantifying the mechanisms of the model, we perform a validation exercise focusing on key nontargeted moments. Using simulated data based on the model's response to the price sequence used in our quantitative analysis (shown in Figure 3), we estimate regressions that are analogous

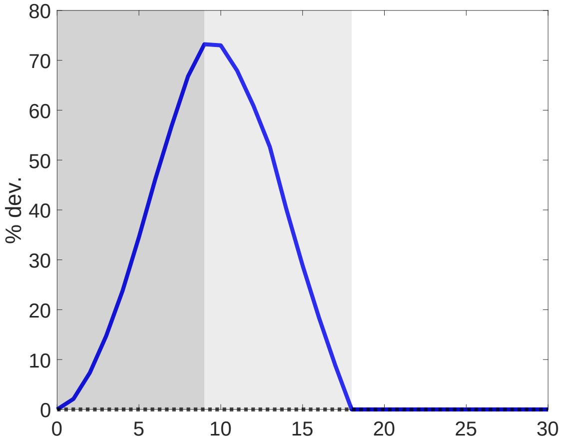

Figure 3: Commodity Price Super-Cycle

Note: This figure shows the path of the price of the commodity good used in the quantitative analysis.

to those estimated in the empirical section.

First, we show that the model can replicate the evolution of the skill premium. We estimate the following regression of the (log) skill premium on the (log) commodity price sequence using the simulated data:

$$\text{SP}_t = \beta_1 \cdot \text{Price}_t + \epsilon_t \,. \tag{12}$$

Since our model has a single region, we exclude the regional fixed effects from equation (3) when using the simulated data. The results are displayed in column 1 of Table 6. In column 2 we report equivalent results for the skill premium estimated using the actual data. To match the model calibration, we use annual data. In addition, we use state–level data because some of the data used to calibrate the model varies at the state–level.[29] This comparison indicates that the model matches very closely the skill premium dynamics triggered by the commodity price super-cycle.[30]

The intersectoral reallocation of labor between the tradable and nontradable sectors lies at the core of the transmission of commodity price super-cycles in our framework. Therefore, this is another important dimension along which to test the calibrated model. To this end, we estimate the following equation of sectoral employment shares (Emp. Share$_{rst}$) on an interaction term

---

[29] Specifically, in Appendix B.1 we develop a multiregion version of the model which we calibrate using state–level domestic trade flows.

[30] The simulated price series fed into the model (shown in Figure 3) is based on the bust to boom commodity price increase in each state from 2002 to 2013, with 2002 corresponding to the initial steady state in the model. In this subsection, we contrast data and model using in both cases the "transition years" 2003-2013.

Table 6: Model Validation: Skill Premium and Regional Employment

	(1)	(2)	(3)	(4)
	Skill Premium		Regional Employment	
	Model	Data	Model	Data
Price	-0.085***	-0.221***		
	(0.003)	(0.087)		
Price $\cdot T_s$			-0.053***	-0.199***
			(0.002)	(0.026)
Observations	11	267	22	594

Note: The first two columns report the results of the estimation of equation (12). Column 1 includes year fixed effects; column 2 includes state and year fixed effects. The last two columns report the results of the estimation of equation (13). Column 3 includes sector and year fixed effects; column 4 includes sector and state × year fixed effects. Columns 1 and 3 use simulated data. Columns 2 and 4 use the actual data. ***, **, and * denote statistical significance at a 1, 5 and 10 percent confidence level, respectively.

between the commodity price and a tradable sector dummy (T_s), and sector and time fixed effects:

$$\text{Emp. Share}_{st} = \beta_1 \cdot \text{Price}_t \cdot T_s + \gamma_s + \delta_t + \epsilon_{st}\,. \tag{13}$$

In the regression estimated with the actual data, we replace the sector and time fixed effects by a set of region × time and sector fixed effects. The coefficient β_1 measures the differential employment response in the tradable sector relative to the nontradable sector. As before, to match the model closely, we use data at an annual frequency with each observation corresponding to a state, although we find similar patterns using data at the microregion level. We compare the estimated parameters obtained using the simulated data with those obtained based on the actual data. The results are reported in columns 3 and 4 in Table 6. We observe a close match between model and data in the reallocation of employment across sectors.

Finally, we compare the model-generated firm-level employment responses to the commodity price super-cycle to those estimated from the actual data to distinguish between the impact of prices on exporting and nonexporting firms in the tradable sector. Specifically, we estimate the following equation:

$$\log(\text{Emp})_{ft} = \beta_1 \cdot \text{Exporter}_{ft} + \beta_2 \cdot \text{Exporter}_{ft} \cdot \text{Price}_t + \gamma_f + \delta_t + \epsilon_{ft}\,. \tag{14}$$

This equation is analogous to equation (5) estimated in the empirical section. Once again, we use actual data with prices defined at the state level at an annual frequency to be consistent with the model calibration. The results are displayed in Table 7, and again show a close parallel between the behavior of the model and the empirical results.

Table 7: Model Validation: Firm-Level Employment

	(1)	(2)
	Tradable	
	Model	Data
Exporter_{ft} x Price_{rt}	-0.264***	-0.094***
	(0.000)	(0.024)
Observations	1,315,872	1,538,219

Note: Columns 1 and 2 in this table reports the results of the estimation of equation (14) using the simulated and actual data respectively. ***, **, and * denote statistical significance at a 1, 5 and 10 percent confidence level, respectively.

It is worth noting that when using the actual data we estimate cross-regional regressions. In contrast, the simulated data are based on a single–region model. Appendix B.1 presents the extension to the multiregion setting, which includes interregional trade and migration. Given that the multiregion model is closer to the empirical design, it is not surprising that with this framework we find an even closer fit between the regression results obtained using model–simulated and actual data. Nonetheless, to a large extent our more parsimonious model captures the key features of the data.

4.3 Transmission Mechanisms

We first use the calibrated model to analyze the extent to which the labor market is a relevant channel of transmission to other sectors of the economy. We first focus on labor reallocation, including both intersectoral reallocation between the tradable (i.e., manufacturing) and nontradable (i.e., service) sectors as well as on reallocation within the tradable sector. Having a calibrated model allows us to document general equilibrium effects, furthering our understanding of the role of labor markets in the transmission of a commodity boom.

For this purpose, we compare the baseline economy to a counterfactual economy in which the commodity sector's labor demand does not change in response to commodity price fluctuations, thus shutting down transmission through the labor market. Therefore, in the counterfactual economy, a commodity boom becomes purely an endowment shock. To construct this counterfactual economy, we fix the commodity sector's employment at the initial steady-state level. We calculate and feed to it a path for the natural resource stock such that commodity output follows the same trajectory in the counterfactual and baseline economies. Thus, commodity prices, output, and revenue are also identical in both economies. In the counterfactual economy, by construction, the commodity super-cycle has no direct impact on wages. Second–order wage movements in the counterfactual economy are the result of general equilibrium effects due to increases in wealth.

Consider first the reallocation of labor between the tradable and nontradable sectors. In the baseline economy, the commodity boom raises the demand for labor in the commodity sector, leading to an increase in wages during the boom (see Figure 4a), which implies a cost increase in both the tradable and nontradable sectors. Given that the tradable sector is unskilled intensive relative to the nontradable sector, the skill premium falls (see Figure 4b) and the cost increase in the baseline economy is larger in the tradable sector. In addition, the increase in wages during the boom raises the prices of all goods, lowering the domestic demand for output, which has a negative impact on sales for the nontradable sector and on domestic sales for the tradable sector. Overlapping these mechanisms, the commodity price boom leads to a decline in employment in tradables relative to nontradables, as shown in Figure 5a. In contrast, in the counterfactual economy, the expansion of the commodity sector leads to changes in wages only through general equilibrium effects, so the reallocation of labor is driven primarily by changes in spending. It is worth noting that in the counterfactual economy the increase in wealth leads to more demand for both tradable and nontradable goods. Given that only tradables can be imported, there is an increase in the production of nontradables relative to tradables. Because the nontradable sector is relatively skilled intensive, this leads to a (small) increase in the skill premium, in contrast to the large decline in the skill premium seen in the baseline economy.

By comparing the baseline and counterfactual economies, we find that about 70% (measured at the peak of the boom) of the intersectoral reallocation in labor between tradables and nontradables is due to changes in wages resulting from the increased labor demand from the commodity sector. The remaining reallocation can thus be attributed to the overall wealth increase in the economy due to the commodity boom, which causes a disproportionate increase in the demand for nontradables. It is worth noting that in the first periods, both the baseline and counterfactual economies behave similarly, and the difference appears gradually over time. The reason is that both economies immediately adjust to the increase in permanent wealth resulting from the anticipated boom, while only the baseline economy faces a gradual reallocation of labor as a result of the transmission occuring through the labor market as the boom develops.

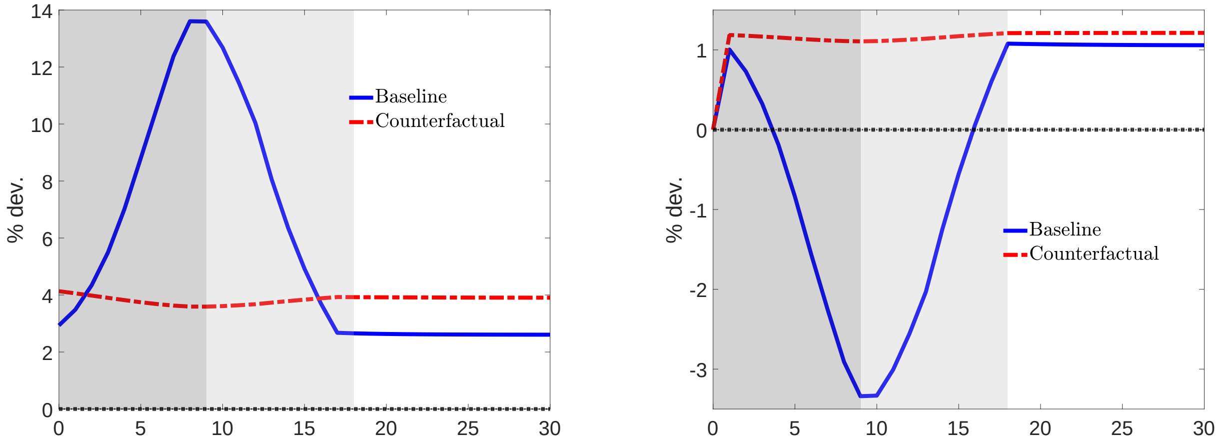

a) Unskilled Wage

b) Skill Premium

Figure 4: Wage Dynamics in the Baseline and Counterfactual Economies

NOTE: This figure shows the response of the unskilled wages and the skill premium to the commodity price super-cycle shown in Figure 3. The blue solid line corresponds to the baseline economy and the red dashed line corresponds to the counterfactual economy described in Section 4.3.

The second outcome on which we focus is the intrasectoral reallocation of labor within the tradable sector, distinguishing between workers producing output for domestic sales versus exports. In the baseline economy, the commodity boom raises local spending, which benefits goods sold locally relative to goods that are exported. In the tradable sector, this leads to an increase in employment oriented toward domestic relative to export sales during the boom.

We focus on a single firm that sells domestically and exports in every period, and compute the ratio between employment oriented to domestic versus export sales.[31] In Figure 5b, we plot this ratio in the baseline and counterfactual economies. Comparing both economies, we find that about 60% (measured at the peak of the boom) of the reallocation of labor within this marginal exporter in the tradable sector is due to changes in wages resulting from the increased labor demand from the commodity sector. Note that wages impact changes in the proportion of employment producing output for domestic versus export sales through general equilibrium

---

[31] We choose a "marginal exporter" that is the least productive exporter at the peak of the boom (i.e., the least productive firm that exports in every period). In the counterfactual economy, we focus on a firm with the same productivity than this marginal exporter in the baseline economy.

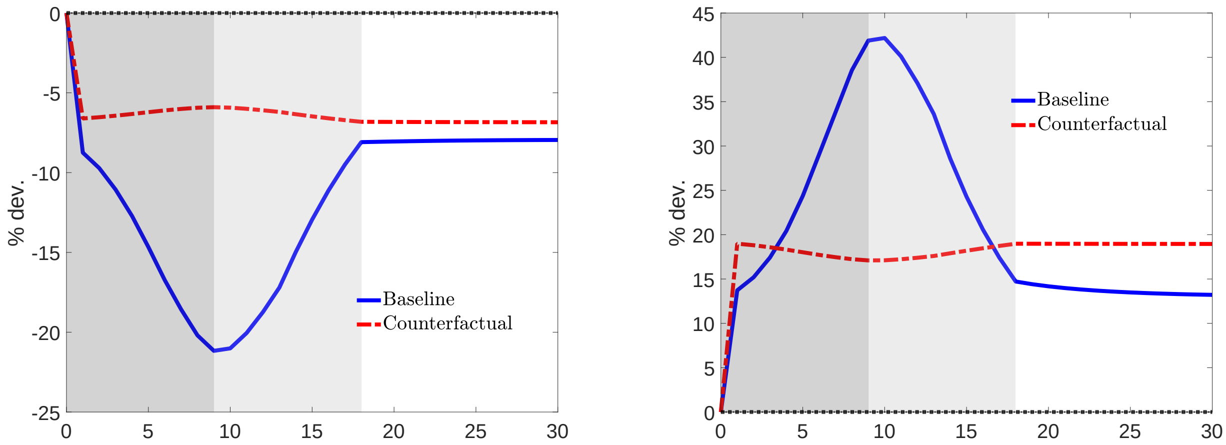

a) Relative T to N Employment (L^T/L^N)

b) Relative Employment in Domestic to Exported Production for Marginal Exporter in T (L^T_d/L^T_x)

Figure 5: Employment Reallocation in the Baseline and Counterfactual Economies

Note: This figure shows the response of the relative tradable to nontradable sector employment (left) and the ratio of labor employed in the production of tradables for the domestic market to labor employed in the production of tradables for the foreign market in the marginal exporter defined in the text (right) to the commodity price super-cycle shown in Figure 3. The blue solid line corresponds to the baseline economy and the red dashed line corresponds to the counterfactual economy described in Section 4.3.

effects. Higher wages in the domestic economy lead to an increase in prices, which, aggregated in the price index, lead to changes in demand for domestic relative to exported output. The remaining reallocation is caused by the wealth increase in the economy due to the commodity boom, which causes a disproportionate increase in the demand for local sales. As in the case of intersectoral reallocation, we document a similar behavior of the baseline and counterfactual economies in the initial period and a subsequent divergence for the same reason discussed earlier.

Overall, the findings on labor reallocation highlight the first order role played by the labor market in propagating commodity price cycles throughout the economy.

4.4 Aggregate Impact and Welfare

We now evaluate the macroeconomic effects of a commodity price super-cycle. We motivated this section arguing that standard models in the literature view commodity booms as endowment shocks, abstracting from the transmission through the labor market. We showed earlier that labor markets are quantitatively important when studying the impact of a commodity price cycle on the reallocation of labor across and within sectors. We now conclude by asking what are the aggregate impacts of acknowledging labor markets as a transmission mechanism.

In Figure 6a, we plot GDP (measured in units of the consumption basket). In the counterfactual economy, the increase in GDP is more than twice as large, foreshadowing larger welfare gains as well. The reason is that in this case the expansion of the commodity sector is driven by an increase in the resource endowment instead of by absorbing labor from the tradable and nontradable sectors. Thus, the other sectors do not contract to fuel the commodity boom.

Figure 6b plots the path of the real exchange rate. Consistent with the empirical cross-country literature, the real exchange rate appreciates persistently (domestic prices increase as the economy is permanently wealthier).[32] Interestingly, the counterfactual economy has a weaker appreciation of the real exchange rate. The domestic consumption basket faces the same increase in commodity prices in both economies. However, in the baseline economy, the increase in wages following rising labor demand in the commodity sector leads to an increase in the cost of tradable and nontradable production. This supply–side effect dominates differences in the demand for these goods between both economies, and leads to a larger increase in the price of tradable and nontradable output in the baseline case.

Next, Figure 6c describes the evolution of the trade balance. In line with the permanent income hypothesis, the representative household adjusts its permanent consumption level on impact and chooses a new and higher level of consumption. Thus, the trade balance follows a countercyclical path. The difference between the baseline and counterfactual economies is due to two opposite forces. First, the larger increase in wealth in the counterfactual economy leads to a larger rise in consumption, which worsens the trade balance relative to the baseline economy. Second, in the counterfactual economy the commodity sector does not draw labor from the tradable sector. In addition, the smaller adjustment in wages leads to a smaller increase in the price of tradable output. Thus, tradable exports fall less in the counterfactual economy, improving the trade balance relative to the baseline economy.

In addition, Figure 6d plots the path of net foreign assets, showing larger wealth accumulation in the counterfactual economy. The reason is that the expansion of the commodity sector in counterfactual economy is driven by an expansion of the natural resource, without drawing labor from other sectors.

To conclude, we analyze the impact of the commodity cycle on welfare and on the long–

---

[32] The real exchange rate is defined as the inverse of the ideal price index of the consumption bundle.

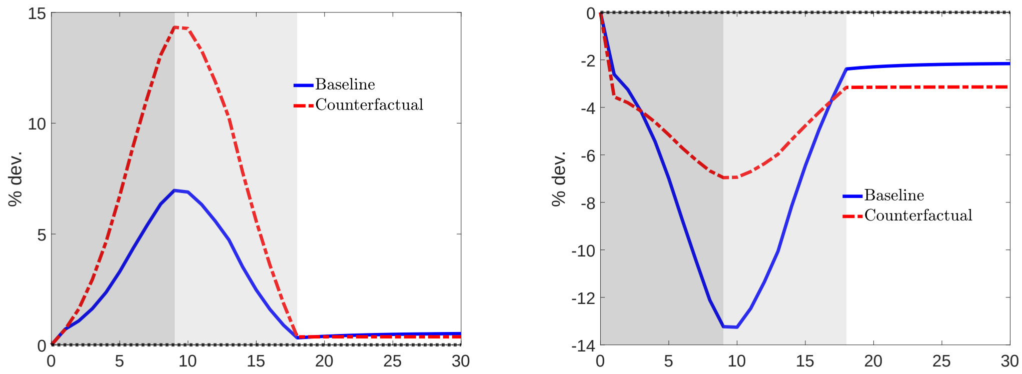

a) GDP

b) Real Exchange Rate

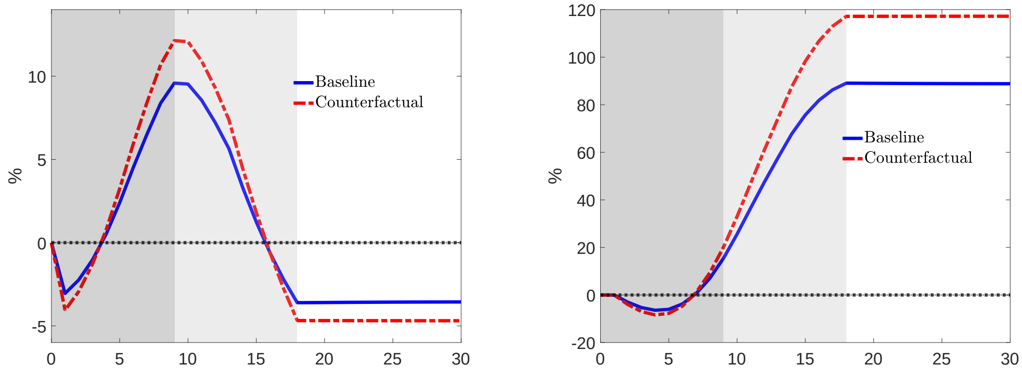

c) Trade Balance to GDP ratio

d) Net Foreign Assets to GDP ratio

Figure 6: Response of Macroeconomic Outcomes to a Commodity Price Super-Cycle

NOTE: This figure shows the response of consumption-deflated GDP (top left), the real exchange rate (top right), the trade balance to GDP ratio (bottom left), and the net foreign assets to GDP ratio (bottom right) to the commodity price super-cycle shown in Figure 3. The blue solid line corresponds to the baseline economy and the red dashed line corresponds to the counterfactual economy described in Section 4.3.

run steady state of net foreign assets.[33,34] The first row in column 1 in Table 8 indicates that

---

[33] To assess the impact of the commodity price super-cycle on welfare we calculate the consumption-equivalent welfare gain from the episode. This is the difference between the welfare gain of an economy exposed to the cycle (W_{cycle}) versus that of one not exposed (W_0), measured as the amount of consumption the representative household would have to receive to be indifferent between experiencing the cycle and remaining in the initial steady state. This is:

$$\text{Consumption-Equivalent Welfare Gain} = \left(\frac{W_{cycle}}{W_0}\right)^{\frac{1}{1-\upsilon}} - 1\,. \tag{15}$$

[34] There are two sources of inefficiency in the model. First, the standard inefficiency of Melitz [2003] type models arising from markups, which does not vary during the commodity price cycle as markups are constant. Second, the cycle ends with net savings and a permanently appreciated real exchange rate (relative to the original steady-state level). This long-run effect on prices leads to a permanent change in production allocation between sectors, which is not internalized by the firms or the representative household. This is akin to the long-run effect of a financial liberalization analyzed by Saffie et al. [2020] and can be traced back to the immiserating growth result in Bhagwati [1958]. Note that a bond–holding cost (an extension discussed later in this section) eliminates this

consumption–equivalent welfare in the baseline economy is 2.27%, implying that the representative household should receive at least a 2.27% permanent increase in steady–state consumption to be willing to forgo the commodity boom–bust cycle. We compare this to the counterfactual economy (shown in the second row), which features nearly twice the amount of consumption–equivalent welfare from the commodity cycle and which accumulates 30% more wealth than the baseline economy. The reason is that, in the counterfactual economy, the lack of transmission through the labor market eliminates the contraction in tradable and nontradable output.

Next, we explore whether some of our key modeling assumptions also impact welfare and wealth accumulation. To this end, we refer to the extensions of our model presented in Appendix B.[35] In the case of the multiregion model developed in Appendix B.1 (see the third row), consumption–equivalent welfare from the commodity cycle is 0.41 percentage points lower than in the baseline economy, given that it is costly to move goods and workers across regions. At the same time, wealth accumulation is only slightly below the baseline economy. We also consider a model with a representative firm, which is much closer to the model of Alberola and Benigno [2017] (see Appendix B.2). As seen in the fourth row, this model has a 0.26 percentage point lower consumption–equivalent welfare from the commodity cycle than the baseline economy. We also see in rows five and six that assumptions about initial household debt have a small impact on welfare and wealth accumulation. Similarly, introducing a bond–holding cost that forces the economy to converge to the original steady state in the long run does not significantly impact welfare, as seen in the seventh row (see Appendix B.4). Finally, given the role played by the trade balance, it is also worth assessing whether changes in trade costs impact welfare. However, changes in export fixed costs only have minor effects, as seen in the last two rows.[36]

Summing up, the most important determinant of consumption–equivalent welfare and wealth accumulation is whether the labor market is allowed to play a role in the transmission of the commodity cycle. To a smaller extent, modeling assumptions such as interregional trade and migration (featured in the multiregion model) as well as firm heterogeneity also have an effect on welfare calculations.

---

long–run inefficiency by forcing the economy to return to its original steady state after the commodity cycle.

[35] In Appendix B we also report the robustness of other outcomes to the alternative modeling assumptions and parametrizations discussed in these paragraphs.

[36] We explore further model simulations in the Appendix. Specifically, given that the literature has established that commodity booms are often accompanied by reductions in interest rates [Shousha, 2016, Fernández et al., 2017a] in Appendix B.6 we use the calibrated model to contrast the response of the economy to a commodity cycle versus a temporary reduction in the real interest rate.

Table 8: Welfare and Long–Run Net Foreign Assets

	CEQ Welfare (%)	Long Run B/Y
Baseline	2.27	0.89
GE Counterfactual	4.38	1.17
Multiregion Model (Appendix B.1)	1.86	0.84
Rep. Firm Model (Appendix B.2)	2.02	1.06
$B_0/Y = 0.20$ (Appendix B.3)	2.19	1.07
$B_0/Y = -0.20$ (Appendix B.3)	2.37	0.71
Bond-holding cost (Appendix B.4)	2.25	0
High f_x (Appendix B.5)	2.10	0.92
Low f_x (Appendix B.5)	2.49	0.85

NOTE: This table reports consumption–equivalent welfare (column 1) and the long–run value of the net foreign assets to GDP ratio (column 2) to the commodity price super-cycle shown in Figure 3. The rows correspond to the baseline economy (first row), the counterfactual economy described in Section 4.3 (second row), the model with a representative firm in the tradable and nontradable sectors (third row), the multiregion model (fourth row), the baseline economy with an initial value of net foreign assets to GDP equal to 0.2 or -0.2 (fifth and sixth rows), the baseline economy adding a bond–holding cost (seventh row), and the baseline economy with a higher or lower value of the fixed cost of exporting (eighth and ninth rows).

5 Conclusions

In this paper, we study the transmission of commodity price super-cycles throughout the economy. We combine rich administrative microdata, a novel empirical strategy, and a quantitative multi-sector model of firm dynamics to identify and quantify key transmission channels in the context of a large commodity boom in Brazil.

We focus on two key transmission channels that were originally proposed by Corden and Neary [1982]'s analysis of the Dutch disease. In response to a commodity boom, an increase in wealth boosts local demand. In addition, an increase in the commodity sector's labor demand generates a cost increase for all sectors of the economy. We extend this framework by studying the heterogeneous impact of each of these mechanisms across firms and sectors, depending on their exposure to local demand (in the case of the wealth channel) and on their skill intensity (in the case of the cost channel). This heterogeneity is the key to the empirical strategy we design to identify the existence of these channels in the data. Our empirical work exploits variation across Brazilian regions in terms of their exposure to the large commodity boom in the 2000s. Using unique microdata on the universe of formal-sector firms across the commodity, manufacturing, and services sectors, we provide robust evidence of the existence of these two transmission channels. We then build and calibrate a multisector model that fits the empirical setting tightly. This model allows us to quantify the transmission mechanisms, study counterfactuals, and assess welfare. We find that transmission through the labor market is crucial. A counterfactual economy which abstracts from it experiences only 30% of the intersectoral labor reallocation between tradables and nontradables, and 40% of the labor reallocation between domestic and exported production within the tradable sector. From an aggregate perspective, the consumption–equivalent welfare gain obtained from a commodity super–cycle is twice as large in the counterfactual economy.

Given the importance of the labor market in the transmission of commodity price super–cycles, one can imagine that labor market frictions can lead to recessions and unemployment during the bust. Future work could extend our framework to study the optimal policy response to commodity cycles under such frictions.

6 Data Availability Statement

Part of the data used in this article consist of confidential administrative records from Brazil, which cannot be shared publicly. A replication package including codes, instructions on how to access the confidential data, and the part of the data that can be shared publicly, are available on Zenodo at `https://doi.org/10.5281/zenodo.8118050`

References

R. Adão. Worker heterogeneity, wage inequality, and international trade: Theory and evidence from Brazil. *Working Paper*, 2015.

R. Adão, M. Kolesár, and E. Morales. Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4):1949–2010, 2019.

E. Alberola and G. Benigno. Revisiting the commodity curse: A financial perspective. *Journal of International Economics*, 108:S87–S106, 2017.

G. Alessandria and H. Choi. Do sunk costs of exporting matter for net export dynamics? *The Quarterly Journal of Economics*, 122(1):289–336, 2007.

G. Alessandria, S. Pratap, and V. Z. Yue. Export dynamics in large devaluations. *Working paper*, 2013.

H. Allcott and D. Keniston. Dutch disease or agglomeration? The local economic effects of natural resource booms in modern America. *The Review of Economic Studies*, 85(2):695–731, 2018.

J. Alvarez, F. Benguria, N. Engbom, and C. Moser. Firms and the decline in earnings inequality in Brazil. *American Economic Journal: Macroeconomics*, 10(1):149–89, 2018.

R. Arezki, V. A. Ramey, and L. Sheng. News shocks in open economies: Evidence from giant oil discoveries. *The Quarterly Journal of Economics*, 132(1):103–155, 2017.

S. T. Ates and F. E. Saffie. Fewer but better: Sudden stops, firm entry, and financial selection. *American Economic Journal: Macroeconomics*, 13(3):304–56, 2021.

D. Autor, D. Dorn, and G. H. Hanson. The China syndrome: Local labor market effects of import competition in the US. *The American Economic Review*, 103(6):2121–68, 2013.

J. Bhagwati. Immiserizing growth: A geometrical note. *The Review of Economic Studies*, 25(3): 201–205, 1958.

F. J. Buera and J. P. Kaboski. Can traditional theories of structural change fit the data? *Journal of the European Economic Association*, 7(2-3):469–477, 2009.

F. Caselli and G. Michaels. Do oil windfalls improve living standards? Evidence from Brazil. *American Economic Journal: Applied Economics*, 5(1):208–238, 2013.

W. M. Corden and J. P. Neary. Booming sector and de-industrialisation in a small open economy. *The Economic Journal*, 92(368):825–848, 1982.

G. Corsetti, L. Dedola, and S. Leduc. International risk sharing and the transmission of productivity shocks. *The Review of Economic Studies*, 75(2):443–473, 2008.

F. Costa, J. Garred, and J. P. Pessoa. Winners and losers from a commodities-for-manufactures trade boom. *Journal of International Economics*, 102:50–69, 2016.

R. Dix-Carneiro and B. K. Kovak. Trade liberalization and the skill premium: A local labor markets approach. *The American Economic Review*, 105(5):551–57, 2015.

R. Dix-Carneiro and B. K. Kovak. Trade liberalization and regional dynamics. *The American Economic Review*, 107(10):2908–46, 2017.

T. Drechsel and S. Tenreyro. Commodity booms and busts in emerging economies. *Journal of International Economics*, 2017.

N. Engbom and C. Moser. Earnings inequality and the minimum wage: Evidence from Brazil. *The American Economic Review*, 112(12):3803–3347, 2022.

B. Faber and C. Gaubert. Tourism and economic development: Evidence from Mexico's coastline. *The American Economic Review*, 109(6):2245–93, 2019.

A. Fernández, A. González, and D. Rodriguez. Sharing a ride on the commodities roller coaster: Common factors in business cycles of emerging economies. *Journal of International Economics*, 111:99–121, 2017a.

A. Fernández, S. Schmitt-Grohé, and M. Uribe. World shocks, world prices, and business cycles: An empirical investigation. *Journal of International Economics*, 108:S2–S14, 2017b.

A. Fernández, S. Schmitt-Grohé, and M. Uribe. Does the commodity super cycle matter? *NBER Working Paper 27589*, 2020.

F. Ghironi and M. J. Melitz. International trade and macroeconomic dynamics with heterogeneous firms. *The Quarterly Journal of Economics*, 120(3):865–915, 2005.

D. Gollin, D. Lagakos, and M. E. Waugh. The agricultural productivity gap. *The Quarterly Journal of Economics*, 129(2):939–993, 2014.

G. Gopinath, Ş. Kalemli-Özcan, L. Karabarbounis, and C. Villegas-Sanchez. Capital allocation and productivity in South Europe. *The Quarterly Journal of Economics*, 132(4):1915–1967, 2017.

T. Harding and A. J. Venables. The implications of natural resource exports for nonresource trade. *IMF Economic Review*, 64(2):268–302, 2016.

E. Helpman, O. Itskhoki, M.-A. Muendler, and S. J. Redding. Trade and inequality: From theory to estimation. *The Review of Economic Studies*, 84(1):357–405, 2017.

T. J. Holmes and J. J. Stevens. An alternative theory of the plant size distribution, with geography and intra-and international trade. *Journal of Political Economy*, 122(2):369–421, 2014.

R. N. F. Jaef and J. I. Lopez. Entry, trade costs, and international business cycles. *Journal of International Economics*, 94(2):224–238, 2014.

R. W. Jones. Income distribution and effective protection in a multicommodity trade model. *Journal of Economic Theory*, 11(1):1–15, 1975.

L. F. Katz and K. M. Murphy. Changes in relative wages, 1963–1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78, 1992.

D. Kohn, F. Leibovici, and H. Tretvoll. Trade in commodities and emerging market business cycles. *American Economic Journal: Macroeconomics*, 13(3):173–208, 2021.

M. A. Kose. Explaining business cycles in small open economies: How much do world prices matter? *Journal of International Economics*, 56(2):299–327, 2002.

B. K. Kovak. Regional effects of trade reform: What is the correct measure of liberalization? *The American Economic Review*, 103(5):1960–1976, 2013.

M. J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003.

E. G. Mendoza. The terms of trade, the real exchange rate, and economic fluctuations. *International Economic Review*, pages 101–137, 1995.

C. M. Reinhart, V. Reinhart, and C. Trebesch. Global cycles: Capital flows, commodities, and sovereign defaults, 1815–2015. *The American Economic Review*, 106(5):574–580, 2016.

J. D. Sachs and A. M. Warner. Natural resource abundance and economic growth. *NBER Working Paper*, 1995.

F. Saffie, L. Varela, and K.-M. Yi. The micro and macro dynamics of capital flows. *NBER Working Paper # 27371*, 2020.

S. Shousha. Macroeconomic effects of commodity booms and busts: The role of financial frictions. *Working Paper*, 2016.

I. Simonovska and M. E. Waugh. The elasticity of trade: Estimates and evidence. *Journal of International Economics*, 92(1):34–50, 2014.

L. Varela. Reallocation, competition and productivity: Evidence from a financial liberalization episode. *The Review of Economic Studies*, 2015.