Economic Growth and the Evolution of Comparative Advantage in an Occupation-Based Network of Industries*

Will Johnson†

July 24, 2018

Abstract

Recent evidence suggests significant changes over time in the pattern of comparative advantage across countries and industries. What drives such dynamics remains an open question. A mechanism suggested by theoretical literature, but not yet brought to bear on this evidence, is learning-by-doing. In this paper I develop a quantitative model of trade and growth with the goal of characterizing the relationship between learning-by-doing and the dynamics of comparative advantage. The model features an occupational dimension to learning, which endogenously generates a particular network structure of inter-industry learning diffusion, based on occupational similarity. The model predicts that countries with a comparative advantage in industries more central in this network will grow more in the aggregate. I use the model-implied dynamics of comparative advantage to quantitatively discipline the amount of occupational learning and the extent to which learning diffuses across industries. Compared to intra-industry learning, I find that cross-industry learning diffusion explains at least four times as much of the dynamics of comparative advantage, as well as forty percent of an industry's contribution to aggregate growth.

JEL Codes: O41, F11, D85

Keywords: economic growth, comparative advantage, industrial networks, occupations, learning-by-doing

---

*First and foremost I would like to thank Dilip Mookherjee, Stefania Garetto, and Martin Fiszbein for their intellectual guidance and support. Lidong Yang and Bingchang Zeng provided excellent research assistance. This research has also benefited from comments from Daron Acemoglu, Kehinde Ajayi, Amrit Amirapu, Diego Anzoategui, Dany Bahar, Samuel Bazzi, Sebastian Bustos, Arnaud Costinot, Simon Gilchrist, Ricardo Hausmann, Danial Lashkari, Frank Neffke, Michael Peters, John van Reenen, Pascual Restrepo, Esteban Rossi-Hansberg, Ulrich Schetter, and Stephen Terry; as well as participants at the 2014 CIREQ Ph.D. conference; the 2015 Advanced Graduate Workshop on Poverty, Development and Globalization; DEGIT 2015; the Fall 2015 Midwest Trade Meetings; NEUDC 2015; SEA 2015; the 2016 Workshop on Networks in Trade and Finance; the 2017 Summer European Meetings of the Econometric Society; the 2018 New Faces Conference; and seminar participants at Boston University, Columbia University's Initiative for Policy Dialogue, and Harvard Kennedy School.

†Department of Economics, Dartmouth College. Email: william.johnson-2@dartmouth.edu

1 Introduction

The principle of comparative advantage is two centuries old, but only recently have Ricardian models of trade been developed that are capable of guiding empirical estimation of the pattern of comparative advantage and how it changes over time.[1] Recent contributions in this literature have shown that there *are* in fact significant changes in comparative advantage over time, and in particular, a country's *level* of comparative advantage in an industry has a significant effect on future *growth* of that country's comparative advantage in that industry. What drives these dynamics is an open question.[2] A leading theory, heavily emphasized in prior theoretical literature but not yet brought to bear on these new findings, is the theory of *learning-by-doing* − that countries get better over time at what they already produce.[3]

In this paper I develop a tractable model of trade and growth that allows me to analytically and quantitatively characterize the relationship between learning-by-doing and the evolution of comparative advantage. Intuitively, the essence of learning-by-doing is that the amount one learns something (and hence the amount one gets relatively better or worse at it, i.e., the increase or decrease in one's comparative advantage in it) depends on how much one does it, and how much one does something depends, in equilibrium, on how much of a comparative advantage one has in it in the first place. The model allows for not only intra-industry learning, but also diffusion of learning across industries. The idea captured by the model is that the amount of learning diffusion between two industries is higher when the two industries are more *similar* to each other, in the tasks and skills needed to produce in those industries.

I operationalize this idea by focusing on one particular, easily quantifiable dimension of similarity in skills and tasks, namely, the *occupational* dimension. The source of growth in the model is occupation-specific learning-by-doing, which partially diffuses across workers in an occupation working in different industries. For example, since the automobile and airplane manufacturing industries employ more engineers than economists, but vice versa for the finance industry, then extra production of automobiles − and the resulting large increase in learning-by-doing among engineers, but smaller increase in learning-by-doing among economists − lowers future costs of production in the airplane manufacturing industry more than the finance industry. This is captured in the model through a parsimonious two-parameter formulation, with one parameter governing the overall amount of learning, and the other parameter governing the extent to which this learning diffuses across industries. These parameters, in combination with the intensity with which each industry uses each occupation, govern the extent to which learning in each particular industry diffuses into each other industry.

[1]Eaton and Kortum (2012) provide a survey of this literature.

[2]See Levchenko and Zhang (2016) and Hanson, Lind, and Muendler (2016). Quoting from Levchenko and Zhang (2016), *"A theoretical and quantitative framework with endogenous sectoral productivity that can be used for understanding the empirical patterns uncovered here has not yet been developed, and remains a potentially fruitful direction for future research."*

[3]See Acemoglu (2008) for a survey of the theoretical literature on learning-by-doing.

The cross-sectional trade structure of the model follows Costinot, Donaldson, and Komunjer (2012), a multi-industry extension of the multi-country model of Eaton and Kortum (2002). This structure overcomes a crucial limitation of previous Ricardian models of trade, by allowing for a country's productivity in an industry to be characterized not just by a single number but by an entire distribution, with the mean of this distribution varying at the country-industry level. Hence, the model-implied pattern of comparative advantage across countries and industries is not − as in a classic Ricardian model − simply a specification of which country has a comparative advantage in which industry, but also *how much* of a comparative advantage each country has in each industry. This allows one to examine gradual changes over time in the pattern of comparative advantage.

Given this structure, I analytically derive the particular evolution of comparative advantage implied by learning-by-doing. In the case of purely intra-industry learning, this evolution takes a simple form: the change over time in a country's comparative advantage in an industry is only a function of its current level of comparative advantage in that industry itself. When learning has an occupational dimension, growth in a country's comparative advantage in industry i is a function of its level of employment of each occupation, which in equilibrium is a function of the country's current level of comparative advantage not only in industry i but also each other industry h, to the extent that industry h uses the same occupations as i.

The dynamics of comparative advantage are examined empirically using data on bilateral, industry-specific trade flows. In line with previous literature, I find a significant association between a country's future growth in comparative advantage in an industry and its level of comparative advantage in that industry itself. In addition, I provide a novel empirical finding: a country's growth in comparative advantage in an industry is positively correlated with its initial level of comparative advantage in occupationally similar industries, holding fixed the country's initial comparative advantage in the industry itself. Occupational similarity is measured using US data on the relative intensities with which each industry employs each occupation. This finding is in line with the mechanism of inter-industry occupational learning diffusion highlighted by the theory.

I then take a more structural approach, using the observed dynamics in comparative advantage, in combination with the structure of the model, to back out the model-implied amount of occupational learning and the extent to which it spills over across industries. I find that a one percentage point increase in the share of a country's labor force in an occupation is associated with 10-13% higher growth in productivity in that occupation in that country from one decade to the next, with at least 70% of this higher occupational productivity diffusing across industries. Compared to intra-industry learning, I find that cross-industry learning diffusion explains at least four times as much of the observed changes over time in comparative advantage.

Occupational learning has important implications not only for the dynamics of comparative advantage, but also for aggregate growth. In particular, my model endogenously generates a particular *network structure* of inter-industry learning spillovers, as a function of occupational similarity across industries. The model predicts that countries with comparative advantages in industries that are *more central* in this network will grow more in the aggregate. According to the calibrated model, on average, 38% of an industry's contribution to aggregate growth is through the inter-industry learning spillovers that it generates.

In sum, this paper shows that a significant fraction of the observed changes over time in the pattern of comparative advantage across countries and industries can be rationalized by occupational learning-by-doing, and that this has important implications for aggregate growth. The paper does this by exploiting a particularly salient implication of learning-by-doing, namely, a specific relationship between the cross-sectional pattern of comparative advantage at a given point in time, and changes in comparative advantage from that period to the next. It is worth noting, however, that learning-by-doing is not the only mechanism that can generate such a relationship. For example, dynamic occupational economies of scale on the firms' side − i.e., when a larger number of people are employed in an occupation, employers better learn over time how to efficiently hire and make use of this occupation − can result in a similar relationship. This paper makes no claim of specifically isolating the particular mechanism of workers learning over time in their occupation. But given the heavy emphasis that learning-by-doing has received in prior literature, it is used to guide our thinking throughout this paper.

This paper contributes to several strands of literature. The mechanism underlying the model is related to the spatial economics literature that finds that occupational similarity plays a significant role in explaining the geographic co-agglomeration of industries. This was first explored in a static setting by Ellison, Glaeser, and Kerr (2010), using cross-sectional data from the US at the metropolitan, county, and state levels. Hanlon and Miscio (2016) examine the dynamics of these industrial co-agglomeration patterns using city-level panel data from the UK. These papers focus on estimating the reduced-form effect of occupational similarity on co-agglomeration across pairs of industries. In this paper I develop a model of economic growth that offers a theoretical rationale for why occupational similarity is an important channel through which growth in one industry spurs growth in another. The model allows me to expand on previous contributions by examining not just the effects of growth in one industry on another industry, but the equilibrium effects on aggregate growth, in both a closed and an open economy.

At the heart of my model is the notion of economic growth through learning-by-doing, which dates back to Arrow (1962). The first to analyze this within a multi-industry framework were Clemhout and Wan (1970) and Bardhan (1971), who showed that if certain industries exhibit more learning-by-doing than others, and if learning is external to individual firms, then this gives theoretical (although not necessarily practical) justification for subsidizing the industries with more learning. Lucas (1988), Young (1991), Matsuyama (1992), and Galor and Mountford (2008) theoretically show that these considerations are further amplified by international trade: if certain countries have a comparative advantage in high-growth industries (i.e., industries with large learning-by-doing externalities) while other countries have a comparative advantage in low-growth industries, then the dynamic gains from trade will be higher for the former countries than for the latter countries. I draw on this literature to inform the choice of functional form for the learning-by-doing function. To the best of my knowledge, I am the first to analyze occupation-specific learning-by-doing and its implications for cross-industry learning spillovers.

The open-economy version of my model draws from the recent literature on quantitative, multi-country models of Ricardian trade. In addition to Costinot, Donaldson, and Komunjer (2012), my paper is closely related to Levchenko and Zhang (2016) and Hanson, Lind, and Muendler (2016) (henceforth "L&Z" and "HLM," respectively), both of whom examine the dynamics of country-industry-level comparative advantage. In particular, both studies find a negative effect of a country's initial level of comparative advantage in an industry on that country's subsequent growth in comparative advantage in that industry. My results are in line with theirs, despite various differences in the data and methodology, which suggests the finding is quite robust. L&Z informally argue (but do not formally demonstrate) that their results go against the theory of learning-by-doing, at least at their data's particular level of aggregation. In this paper, however, I formally show that L&Z's argument only holds under certain strong assumptions about the learning process; I show that, in general, learning-by-doing is, in fact, consistent with the empirical evidence.

The inter-industry learning spillovers that endogenously result from my model connect this paper with a recently flourishing literature on the macroeconomic implications of network structures among industries. This literature, which has given particularly extensive attention to the input-output structure of the economy, dates back to Hirschman (1958), who influentially argued that economic development in one sector induces development in other sectors that either use, or are used by, that sector as an input ("forward" and "backward" linkages, respectively). This view of the development process was formalized by Rodriguez-Clare (1996a and 1996b). Jones (2011) develops a static model in which forward and backward linkages amplify the effects of exogenous sector-specific distortions on aggregate total factor productivity, using US input-output data for illustration. Oberfield (2013) develops a model in which the input-output structure of the economy arises endogenously from firms searching for the lowest-cost suppliers of inputs; low-cost suppliers endogenously emerge as "star suppliers," providing inputs for many other firms and playing an important role in propagating cost savings throughout the economy. The implications of the input-output structure of the economy for volatility and business cycles were explored in Long and Plosser (1983), and more recently Carvalho (2010) and Acemoglu et al (2012).

Note, however, that the network of industries in this paper is *not* based on input-output linkages. While input-output linkages are an important channel through which shocks and distortions are transmitted across industries, they are not, in and of themselves, a source of growth.[4] In contrast, the occupational learning spillovers that are the source of the inter-industry network structure in this paper are indeed, in and of themselves, a source of endogenous long-run growth. The focus of this paper on similarity across industries in their required skills and knowhow complements the work of Hidalgo, Klinger, Barabasi, and Hausmann (2007), who use observed overlap in the countries that export each product to estimate an underlying network of products (which they call the "product space") that

[4]This same point is made by Hanlon and Miscio (2016).

represents overlap in the capabilities required to produce products. They offer evidence that when countries branch out into new products, they do so by moving into products that are near their old products in this network.

The difference between the network of input-output linkages and the network of industries based on occupational similarity can be seen in Figures 1 and 2, respectively. As these figures illustrate, the topology of the network of industries based on occupational similarity (the subject of this paper) is significantly different from that of the network of input-output linkages.

The rest of this paper is organized as follows. In section 2, I introduce and analyze the model in the context of a closed economy, while in section 3 I extend the analysis into a multi-country framework. Section 4 analytically characterizes the link between learning-by-doing and the evolution of comparative advantage, and provides reduced-form evidence of such a link. In Section 5 I put quantitative discipline on the theory by taking the model to data on trade and employment. I then use the calibrated model to quantify the importance of inter-industry learning diffusion to the dynamics of comparative advantage and aggregate growth. Section 6 concludes.

Figure 1: The input-output structure of the US manufacturing sector

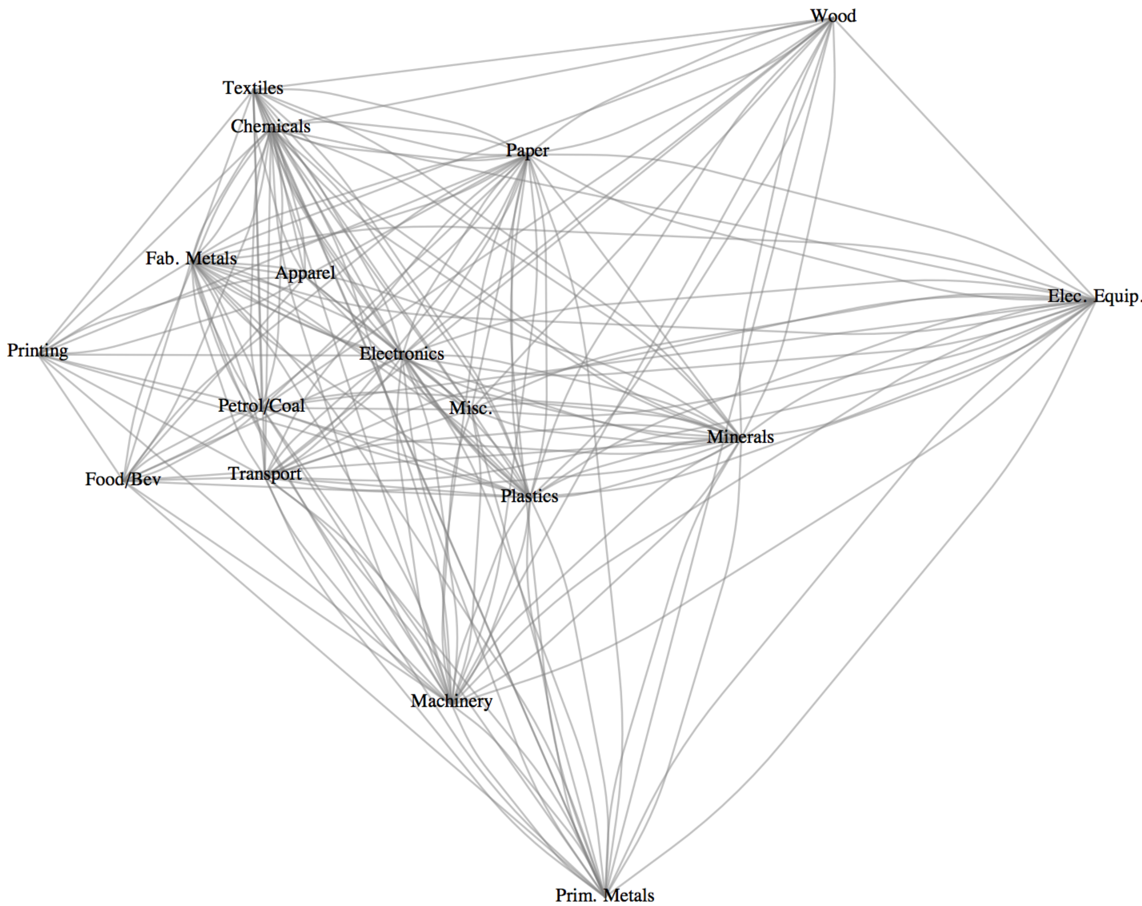

Note: this figure was generated using the 2013 US Direct Requirements input-output table from the Bureau of Economic Analysis, at the three-digit NAICS level of aggregation (comprising 20 industries). The (h, i) element of the Direct Requirements table is the number of dollars worth of intermediate input h used to make one dollar's worth of output in industry i. For each pair of industries h and i, if the (h, i) element in the Direct Requirements table is above the 60th percentile of entries, then a line is drawn in this figure between industries h and i.

Figure 2: Network structure of US manufacturing sector based on occupational similarity

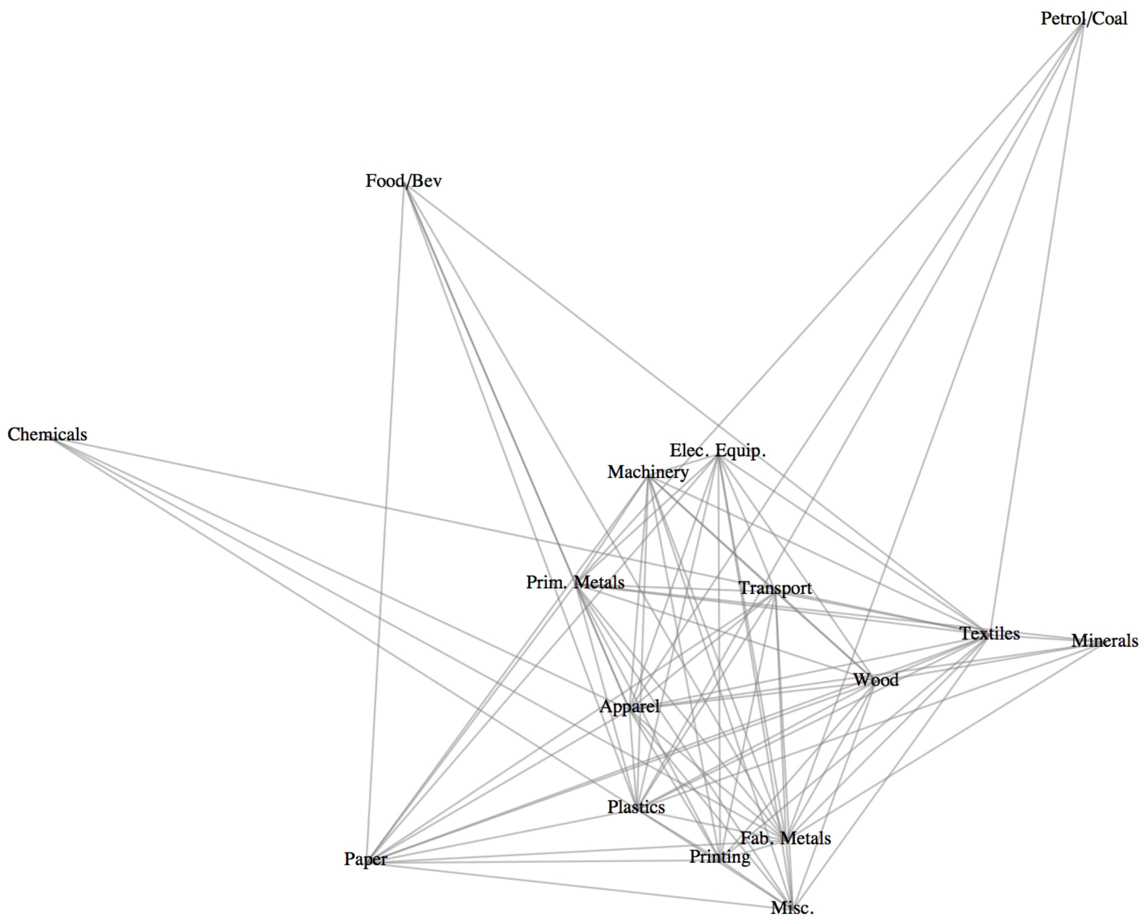

Note: this figure was generated using 2013 US data on each industry's employment of each occupation, from the Occupational Employment Statistics program at the Bureau of Labor Statistics. Industries are classified by their three-digit NAICS codes, and occupations are classified by their two-digit SOC (Standard Occupational Classification) codes. The figure is based on the matrix whose (h, i) element is the correlation coefficient between the industry-h vector of employment of each occupation (as a fraction of industry h's total employment) and the corresponding industry-i vector. For each pair of industries h and i, if the (h, i) element in this matrix is above the 60th percentile of entries, then a line is drawn in this figure between industries h and i.

2 Closed-economy model

I model an economy with multiple industries and multiple occupations. To fix ideas, consider the following three industries and two occupations: automobile manufacturing, airplane manufacturing, and finance; and engineers and economists. There is learning-by-doing within each occupation, which spills over to everyone in the occupation regardless of the industry for which they are working.

Consider, then, what happens if production increases in the car industry. Since the car industry employs a large number of engineers but only a small number of economists, this will cause a significant increase in learning-by-doing among engineers, not so much among economists. The extent to which this benefits another industry corresponds to how much that other industry is engineer-intensive rather than economist-intensive − in particular, it will lower the cost of production in the airplane manufacturing industry more than the finance industry, since the former is engineer-intensive while the latter is economist-intensive.

We can then think of industries as forming a network, where, for any two industries, the strength of the link between them corresponds to how similar they are in their intensity of usage of different occupations. As we will see in the analysis that follows, an industry that is more central in this network will generate more learning spillovers and thereby contribute more to aggregate economic growth.

2.1 The economic environment

Consider a closed economy with I industries[5], indexed by i; J occupations, indexed by j; and an arbitrary number of discrete time periods, indexed by t.

The representative household in this economy has a CES utility function over its consumption of each good:

$$U_t = (\sum_i \beta_i C_{it}^{\frac{\gamma-1}{\gamma}})^{\frac{\gamma}{\gamma-1}} \tag{1}$$

where U_t is the representative household's utility at date t, C_{it} is its consumption of good i at date t, β_i is an exogenous parameter governing the intensity of the household's preference for good i, and γ is the hosehold's elasticity of substitution across the different goods.

At each date t, the representative household chooses $\{C_{it}\}_i$ to maximize U_t subject to its budget constraint:

$$\sum_i P_{it} C_{it} = w_t L_t \tag{2}$$

where P_{it} is the price of good i at date t; w_t is the wage at date t; and L_t is the household's exogenous endowment of labor at date t, which it supplies inelastically.

The representative firm in industry i produces its good using labor from each occupation, according to a production function that is CES with respect to the different occupations:

$$Y_{it} = [\sum_j \alpha_{ij} (\phi_{ijt} L_{ijt})^{\frac{\epsilon-1}{\epsilon}}]^{\frac{\epsilon}{\epsilon-1}} \tag{3}$$

where Y_{it} is output in industry i at date t, L_{ijt} is employment of occupation j in industry i at date t (in units of raw labor), ϕ_{ijt} is productivity in occupation j in industry i at date t (which evolves over time from learning-by-doing, as explained below), α_{ij} is an

[5] I will use the terms "industry" and "good" interchangeably.

exogenous parameter governing the intensity with which industry i uses occupation j, and ϵ is the elasticity of substitution across occupations (which, for tractability, is assumed to be constant across industries).

At each date t and in each industry i, the industry-i representative firm chooses $\{L_{ijt}\}_j$ to maximize its profit π_{it}:

$$\pi_{it} = P_{it}Y_{it} - \sum_j w_t L_{ijt} \tag{4}$$

At each date t, labor markets must clear:

$$\sum_i \sum_j L_{ijt} = L_t \tag{5}$$

as well as goods markets, so for each industry i at each date t:

$$C_{it} = Y_{it} \tag{6}$$

Productivity in each occupation j and industry i evolves over time from learning-by-doing $-$ which is not internalized by individual agents[6] $-$ according to the following learning-by-doing function:

$$\phi_{i,j,t+1} = \phi_{ijt}[(1 + \tilde{L}_{ijt})^{1-\sigma}(1 + \tilde{L}_{jt})^{\sigma}]^{\rho} \tag{7}$$

where ϕ_{ijt} is productivity in occupation j in industry i at date t; $\tilde{L}_{ijt}$ is the share of the economy's total effective labor in occupation j in industry i at date t; $\tilde{L}_{jt}$ is the share of the economy's total effective labor in occupation j at date t, summed across all industries[7]; and $\rho > 0$ is an exogenous parameter governing the rate of learning-by-doing, while $\sigma \in (0, 1)$ governs the extent to which this learning spills over across industries. ϕ_{ij0} is exogenously given for each occupation j and industry i.

Equation (7) is saying, for example, that the higher the fraction of the work force working as engineers, the more learning-by-doing there will be among engineers. The elasticity of learning with respect to labor usage is given by ρ. A fraction $\sigma \in (0, 1)$ of this learning spills over across all engineers, regardless of which industry they are working

[6] The assumption that learning-by-doing is in the form of externalities is made for tractability, and is common in the macroeconomic literature on learning-by-doing, a brief survey of which is given by Acemoglu (2008). Furthermore, the functional form in equation (7) $-$ in particular, the assumption of a constant elasticity of learning with respect to labor usage $-$ is widely used in this same literature; the novelty here is in the occupational dimension of learning.

[7] That is, $\tilde{L}_{ijt} \equiv \frac{\phi_{ijt}L_{ijt}}{\sum_j \sum_i \phi_{ijt}L_{ijt}}$ and $\tilde{L}_{jt} \equiv \frac{\sum_i \phi_{ijt}L_{ijt}}{\sum_j \sum_i \phi_{ijt}L_{ijt}}$. This formulation (in terms of shares rather than levels) is chosen in order to avoid country-level scale effects $-$ i.e., a doubling of the total size of the labor force causing a doubling of the rate of per capita economic growth $-$ which are at odds with the data (see Rose (2006)). Equivalently, one could model the learning process as a function of levels rather than shares, thereby exhibiting scale effects, but with the learning happening at a local level (e.g., the city level), with learning spillovers across localities weak enough that the scale effects do not operate at the country level $-$ this echoes a similar point made by Ramondo et al (2016).

in, while the other fraction $1 - \sigma$ of the engineers' learning is industry-specific. Note that implicit in equation (7) is an assumption, made for tractability, that ρ and σ are constant across industries and occupations. In results that are available upon request, I relax this assumption. However, it is worth emphasizing that even without making any assumptions about occupations or industries exogenously varying from one another in their rates of learning-by-doing, the model still results in heterogeneity across industries in the amount of learning spillovers they generate, due to their different levels of centrality in the occupational network of industries, as we will see below.

2.2 Equilibrium

Given the parameters $\{\beta_i, \gamma, L_t, \alpha_{ij}, \epsilon, \rho, \sigma, \phi_{ij0}\}_{i,j,t}$, an *equilibrium* of the economy is defined as a path $\{L_{ijt}, Y_i, C_i\}_{i,j,t}$ such that at each date t,

1. The household's consumption of each good $\{C_{it}\}_i$ maximizes its utility subject to its budget constraint, given prices $\{P_{it}\}_i$ and the wage w_t,
2. In each industry i, the industry-i representative firm's employment of each occupation $\{L_{ijt}\}_j$ maximizes its profit, given the price of its output P_{it} and the wage w_t,
3. The labor market clears: $\sum_i \sum_j L_{ijt} = L_t$,
4. The goods markets clear: $C_{it} = Y_{it}$ for every industry i, and
5. Productivity in each occupation j and industry i evolves over time according to the learning-by-doing equation (7).

The equilibrium of this economy, which for any set of parameters always exists and is unique, is characterized as follows. The equations below are derived in the usual way from the CES structure of the production and utility functions in equations (1) and (3). In what follows, I normalize the nominal wage w_t at each date t to 1, with all other prices expressed relative to this.

The equilibrium price of output in industry i is

$$P_{it} = [\sum_j (\alpha_{ij})^\epsilon (\frac{1}{\phi_{ijt}})^{1-\epsilon}]^{\frac{1}{1-\epsilon}} \tag{8}$$

Equilibrium consumption and production in industry i are

$$Y_{it} = C_{it} = \frac{(\beta_i)^\gamma (P_{it})^{-\gamma} L_t}{\sum_i (\beta_i)^\gamma (P_{it})^{1-\gamma}} \tag{9}$$

The equilibrium usage of occupation-j labor in industry i is

$$L_{ijt} = \frac{(\alpha_{ij})^\epsilon (\phi_{ijt})^{\epsilon-1} P_{it} Y_{it}}{\sum_j (\alpha_{ij})^\epsilon (\phi_{ijt})^{\epsilon-1}} \tag{10}$$

As described in the previous section, productivity in each industry and occupation ϕ_{ijt} evolves from date t to $t+1$ according to the learning-by-doing equation (7), with ϕ_{ij0} given for each occupation j and industry i. For a given specification of parameters, equations (7) through (10) characterize the equilibrium of the economy.

2.3 Inter-industry learning spillovers

To see how the occupational learning-by-doing induces a network structure among industries, consider an increase in production in some arbitrary industry i and what effect this has on growth in some other arbitrary industry h.

Given an increase in production in industry i, this induces more usage of each occupation j, by an amount governed by α_{ij}, which is the parameter in industry i's production function (3) which governs how intensely industry i uses occupation j.

This extra usage of occupation j induces more learning-by-doing in occupation j. This lowers the cost of production in industry h, by an amount governed by α_{hj}, since α_{hj} governs how intensely industry h uses occupation j.

Thus, the extent to which growth in industry i lowers the cost of production in industry h is a function of how *similar* industries i and h are, in terms of how much they use each occupation.

We can gain further intuition into the network structure of learning spillovers among industries by considering the simple case where the representative household's utility function is Cobb-Douglas, the representative firms' production functions are Leontief, and $\sigma = 1$ (that is, occupational learning perfectly spills over across industries). In this case, we get the following intuitive result, as a first-order log-linear approximation, for how much an exogenous increase in production in industry i at date t induces extra output growth in industry h in equilibrium between date t and date $t+1$:

$$\frac{d \log Y_{h,t+1}}{d \log Y_{it}} \approx \rho \sum_j \alpha_{hj} \tilde{\alpha_{ij}} \tag{11}$$

where $\tilde{\alpha_{ij}}$ is the *relative* intensity with which industry i uses occupation j – that is, $\tilde{\alpha_{ij}} \equiv \frac{\alpha_{ij}}{\sum_i \alpha_{ij}} - \frac{1}{I}$, where I is the number of industries. The derivation of this result is given in Appendix (A).

The result in equation (11) makes it explicit that if industries i and h are more *similar* to each other in their occupational usage – that is, if α_{hj} is high for the same occupations for which $\tilde{\alpha_{ij}}$ is high – then industries i and h will have a larger amount of learning spillovers between them.

Note that equation (11) only describes how growth in industry i induces growth in industry h one period ahead. But also note that, under this first-order approximation, this relationship does not change over time. If we carry this approximation forward, then we get the intuitive result that industry i's importance to long-run aggregate growth (in a sense that will be made precise below) is a function of the *Bonacich centrality* of industry

i in the network of learning spillovers.[8]

Formally, let A denote the matrix whose (i,h) element is $\frac{d\log Y_{h,t+1}}{d\log Y_{it}}$ in equation (11) above − that is, A is the network matrix for the network of inter-industry learning spillovers. Then an industry's importance to long-run aggregate growth is captured, under this first-order approximation, by its Bonacich centrality in this network − that is:

$$W \approx \beta + \delta(\mathbb{I} - \delta A)^{-1} A\beta \tag{12}$$

where β is the vector of each industry's exponent in the representative household's Cobb-Douglas utility function, δ is the representative household's discount factor, $\mathbb{I}$ is the identity matrix, and W is the vector whose ith element is the percentage increase in the total discounted utility of the representative household (from date t onward) from a one percent increase in production in industry i at date t. The derivation of this result is given in Appendix (A).

3 Open-economy model

In the previous section I showed how, in the presence of occupational learning-by-doing, industries that are more central in the network of inter-industry occupational learning spillovers will contribute more to aggregate growth. In this section I extend the analysis to incorporate multiple countries trading with each other, in order to address the question of how a country's amount of aggregate growth depends on which industries it produces in equilibrium, i.e., on which industries it has a comparative advantage in. Moreover, the dynamics of comparative advantage that endogenously arise from this open-economy model will allow me in Section 5 to identify learning-by-doing in the data.

In this section I combine the model from section 2 with the static model of Ricardian trade from Costinot, Donaldson, and Komunjer (2012) (henceforth "CDK").[9] At each date t, the model in this section is essentially the CDK model − the only difference is that there are multiple occupations, but this only matters for dynamics. The dynamics of the model, as in section 2, are governed by learning-by-doing − as before, this learning-by-doing is within each occupation, which spills over to everyone in the occupation regardless of the industry for which they are working, generating network effects among industries. The important thing to note here, which was a moot point in the single closed economy case, is that these spillovers are *within* countries, not across countries.[10]

[8]Bonacich centrality is a measure of how important a node is in a network − e.g., in a network of friends, the Bonacich centrality of an individual is her number of friends, plus a discount factor times the number of friends her friends have, plus the discount factor squared times the number of friends her friends of friends have, ad infinitum.

[9]The CDK model is an extension of Eaton and Kortum (2002) that allows for multiple industries. More precisely, it allows for asymmetries in the production function parameters across industries; this allowance for asymmetries across industries (in my case, asymmetries in how intensely each industry uses each occupation) is what makes the CDK model useful for my purposes.

[10]In an extension to the analysis that is available upon request, I allow for learning to partially spill over across countries. Note that if learning *perfectly* spills over across countries, the model is trivial; the amount

3.1 Economic environment and equilibrium

As before, time is discrete and indexed by t. There are now N countries, indexed by m and n. As before, there are I goods[11], but now each good i comes in a countably infinite number of varieties indexed by $\omega \in \Omega \equiv \{1, ..., +\infty\}$.[12]

As before, labor is the only factor of production; workers can work in J different occupations, indexed by j. Labor is perfectly mobile across occupations but immobile across countries. Country m is endowed exogenously with L_{mt} workers at date t; each worker in country m at date t is paid wage w_{mt}, which will be determined in equilibrium.

The production structure of the economy is analogous to the closed-economy version of the model, except with the addition of total factor productivity terms, which will be discussed below. Specifically, the production function for variety ω of final good i in country m is as follows:

$$y_{imt}(\omega) = z_{im}(\omega)[\sum_j \alpha_{ij}(\phi_{ijmt}L_{ijmt}(\omega))^{\frac{\epsilon-1}{\epsilon}}]^{\frac{\epsilon}{\epsilon-1}} \tag{13}$$

where $y_{imt}(\omega)$ is the quantity of variety ω of final good i produced in country m at date t; $L_{ijmt}(\omega)$ is the amount of raw labor in occupation j used in the production of variety ω of final good i in country m at time t; ϕ_{jmt} is productivity in occupation j in industry i in country m at date t (which evolves over time from learning-by-doing, as explained below); α_{ij} is an exogenous parameter governing the intensity with which industry i uses occupation j; ϵ is the elasticity of substitution across occupations; and $z_{im}(\omega)$ is the total factor productivity of variety ω of final good i in country m, to be discussed below.

The TFP term $z_{im}(\omega)$ is a random variable drawn independently for each triplet (i, m, ω) from a Fréchet distribution $F_{im}(\cdot)$ such that[13]

of productivity growth in each industry is then equalized across countries rather than being a function of which industries are produced in each particular country. The model is only interesting when there is at least *some* degree of localization of learning, and in the baseline analysis I explore the simplest possible case in which learning is entirely localized.

[11]I use the terms "goods" and "industries" interchangeably throughout this section.

[12]This infinite-variety structure is standard in the Ricardian trade literature, for the following reason. If there were only one variety of each good, then the equilibrium would be riddled with corner solutions – that is, for each good j and each country i, one hundred percent of country i's consumption of good j would be sourced by whichever country n could produce and deliver good j to country i most cheaply (or, in knife-edge cases in which two or more countries could do so equally cheaply, there would be multiple equilibria). With the infinite-variety structure, this is exactly what happens at the variety level, but when we aggregate up to the good level − as we will do when taking the model to data − we have interior solutions.

[13]This Fréchet distributional assumption is widely used in the quantitative Ricardian trade literature. Kortum (1997) and Eaton and Kortum (1999) show how this structure can arise endogenously from a process of technological innovation and diffusion − the intuition is that when each variety of each good is produced according to the best known technique for producing that variety in that country, then the distribution of productivity across varieties in a country will be an extreme value distribution such as the Fréchet. Moreover, as discussed in Eaton and Kortum (2002), among extreme value distributions, the Fréchet distribution uniquely gives us a tractable framework in which the equilibrium distributions of labor requirements, costs of production, and prices are all in the same family.

$$F_{im}(z) = \exp[-(\frac{z}{z_{im}})^{-\theta}] \tag{14}$$

for all $z \geq 0$, where $z_{im} > 0$ $\forall i, m$ and $\theta > 1$. z_{im} is the total factor productivity of country m in good i when averaged across good i's infinite varieties, while θ governs the dispersion of productivity, which is an important parameter in the Ricardian trade literature, because the more that productivity varies, the more important is the force of comparative advantage. Note that, for my purposes, these parameters are fixed over time, while the occupational productivity terms evolve over time from learning-by-doing, as will be discussed below.

Now let us consider trade between countries. I will make the standard assumption of iceberg trade costs, meaning that for each unit of good i shipped from country m to country n, only $\frac{1}{d_{imn}} < 1$ units arrive, with d_{imn} such that $d_{imm} = 1$ $\forall m$ and $d_{imn} \leq d_{iml}d_{iln}$ for any third country l.

It follows from the CES production structure in (13), combined with the assumption of iceberg trade costs, that

$$c_{imnt}(\omega) = \frac{d_{imn}w_{mt}}{z_{im}(\omega)}[\sum_j (\alpha_{ij})^{\epsilon}(\phi_{ijmt})^{\epsilon-1}]^{\frac{1}{1-\epsilon}} \tag{15}$$

is the cost of producing and delivering one unit of variety ω of good i from country m to country n at date t. Aggregating up to the good level, define c_{imnt} as follows:

$$c_{imnt} \equiv \frac{d_{imn}w_{mt}}{z_{im}}[\sum_j (\alpha_{ij})^{\epsilon}(\phi_{ijmt})^{\epsilon-1}]^{\frac{1}{1-\epsilon}} \tag{16}$$

Markets are perfectly competitive, and therefore the price $p_{int}(\omega)$ paid by buyers in country n for variety ω of good i at date t is

$$p_{int}(\omega) = \min_{1\leq m\leq N}[c_{imnt}(\omega)] \tag{17}$$

and the set of varieties of good i that are exported by country m to country n at date t is given by

$$\Omega_{imnt} \equiv \{\omega \in \Omega | c_{imnt}(\omega) = \min_{1\leq l\leq N} c_{ilnt}(\omega)\} \tag{18}$$

Each country has a representative consumer whose utility function is a Cobb-Douglas function of the composite goods, where each composite good is a CES function of its infinite varieties. Let β_{im} be country m's Cobb-Douglas exponent on good i, and let σ_{im} be country m's elasticity of substitution among the infinite varieties of good i. (As in CDK (2012), I assume $\sigma_{im} < 1 + \theta$ $\forall i, m$.) Accordingly, define p_{imt} as follows:

$$p_{imt} \equiv [\sum_{\omega\in\Omega} p_{imt}(\omega)^{1-\sigma_{im}}]^{\frac{1}{1-\sigma_{im}}} \tag{19}$$

Then, defining $e_{imt}(\omega)$ as total expenditure by country m on variety ω of good i at date t, we have

$$e_{imt}(\omega) = (\frac{p_{imt}(\omega)}{p_{imt}})^{1-\sigma_{im}} \beta_{im} w_{mt} L_{mt} \tag{20}$$

Furthermore, define e_{imnt} as the value (in dollar terms) of total exports of good i from country m to country n at date t; that is,

$$e_{imnt} \equiv \sum_{\omega \in \Omega_{imnt}} e_{int}(\omega) \tag{21}$$

Then we get the result that

$$e_{imnt} = \frac{(c_{imnt})^{-\theta}}{\sum_{l=1}^{N} (c_{ilnt})^{-\theta}} \beta_{in} w_{nt} L_{nt} \tag{22}$$

The date-t equilibrium of the world economy is pinned down by a balanced trade condition. Let π_{imnt} be country m's share of the world exports (in dollar terms) of good i to country n at date t; that is,

$$\pi_{imnt} \equiv \frac{e_{imnt}}{\sum_{l=1}^{N} e_{ilnt}} \tag{23}$$

Then, for a given wage vector $w_t = (w_{mt})_m$,

$$Z_{mt} = (\sum_{n=1}^{N} \sum_{i=1}^{I} \pi_{imnt} \beta_{in} w_{nt} L_{nt}) - w_{mt} L_m \tag{24}$$

is the excess demand for country m's labor at date t. The equilibrium at date t is pinned down by specifying that $Z_{mt} = 0$ for every country m.

As is typical in the Ricardian trade literature, there is no closed form solution for this date-t equilibrium, but it can be computed using an algorithm from Alvarez and Lucas (2007). The basic idea behind their algorithm is simple: start with an arbitrary guess for the equilibrium wage vector $w_t = (w_{mt})_i$, calculate each country's excess demand for labor Z_{mt}, and then raise the wage of any country m for which $Z_{mt} > 0$ while lowering the wage of any country m for which $Z_{mt} < 0$. Keep doing this, and, under regularity conditions discussed by Alvarez and Lucas, the algorithm will converge to a unique equilibrium wage (from which the equilibrium values of all other variables can be straightforwardly computed).

That completes the description of the economy at date t. Each country's productivity in each occupation evolves over time from learning-by-doing − which is not internalized by individual agents − according to the following learning-by-doing function:

$$\phi_{i,j,m,t+1} = \phi_{ijmt}[(1 + \tilde{L}_{ijmt})^{1-\sigma}(1 + \tilde{L}_{jmt})^{\sigma}]^{\rho} \tag{25}$$

where $\tilde{L}_{ijmt}$ is the share of country m's total effective labor in occupation j in industry i at date t, summed across all varieties of industry i; $\tilde{L}_{jmt}$ is the share of country m's total effective labor in occupation j at date t, summed across all varieties of all industries[14]; $\rho > 0$ is an exogenous parameter governing the rate of occupational learning-by-doing, while $\sigma \in (0, 1)$ governs the extent to which this learning spills over across industries; and ϕ_{ijm0} is exogenously given for each occupation j, industry i, and country m.

3.2 Inter-industry spillovers in the open-economy model

In section 2.3 we asked, in the context of the closed-economy version of the model, when we give an exogenous positive shock to production in a specific industry, what are the effects on every other industry? In this section we ask, using the open-economy version of the model, when we give an exogenous positive shock to production in a specific industry in a specific country, what are the effects on every other industry in every other country? Furthermore, what are the effects on each country's welfare?

To start with, the inter-industry learning spillovers that were already present in the closed-economy model carry over to the open-economy model. Now, however, thanks to international trade, the learning-by-doing induced in a country by extra production in an industry will affect other industries in that country not only through direct learning spillovers, but also through indirect general equilibrium effects; an increase in learning-by-doing in a country pushes up the country's equilibrium wage, which – all else being equal – makes industries in that country less competitive, and furthermore, consumers in all countries benefit from the fall in the costs of production (and hence prices) induced by learning-by-doing, not just the learning country.

If, as in section 2.3, we consider the particular case in which representative firms' production functions are Leontief, then the network structure of inter-industry spillovers takes a particularly simple, intuitive form, as we will see below. Furthermore, in what follows, purely to simplify the expressions, I set the size of each country's labor force equal to one another (normalized to 1), I set each industry's exponent in each country's representative consumer's Cobb-Douglas utility function equal to one another (namely, $1/I$, where I is the number of industries), and I set $\sigma = 1$ (i.e., occupational learning perfectly spills over across industries). Lastly, in order for the model to be analytically tractable, I assume in this section that trade costs are zero.[15]

Under the above simplifying assumptions, the first-order approximation of country m's wage at date t is[16]

$$w_{mt} \approx [\sum_i (\frac{z_{im}}{\sum_j \alpha_{ij}\psi_{jmt}})^{\theta}]^{\frac{1}{1+\theta}} \tag{26}$$

Note that this is a weighted average of country m's date-t productivity in industry i across all i, as one would intuitively expect. This result holds with exact equality when

[14]That is, $\tilde{L}_{ijmt} \equiv \frac{\sum_\omega \phi_{ijmt} L_{ijmt}(\omega)}{\sum_j \sum_i \sum_\omega \phi_{ijmt} L_{ijmt}(\omega)}$ and $\tilde{L}_{jmt} \equiv \frac{\sum_i \sum_\omega \phi_{ijmt} L_{ijmt}(\omega)}{\sum_j \sum_i \sum_\omega \phi_{ijmt} L_{ijmt}(\omega)}$

[15]Generalizing this analysis to non-zero trade costs is straightforward, so long as the trade costs are symmetric across countries and industries – as discussed in Alvarez and Lucas (2007), asymmetric trade costs make it difficult to get any analytical traction in a Ricardian trade model like this one.

[16]The derivation of this result is available upon request, along with all the other results of this section.

the productivity terms are symmetric across countries and industries, but it is only an approximation otherwise. The results in this section, which are first-order approximations, are derived by plugging (26) into the equations of the model and then log-linearizing the resulting system of equations.

For the purposes of this section, let $\hat{y}_{imt}$ denote the logarithm of production in country m in industry i at date t, and let $\hat{W}_{mt}$ denote the logarithm of country m's welfare at date t. Further, let

$$\tilde{\alpha}_{ijm} \equiv \frac{\frac{\alpha_{ij}}{z_{im}}}{\sum_i \frac{\alpha_{ij}}{z_{im}}} - \frac{\frac{1}{z_{im}}}{\sum_i \frac{1}{z_{im}}} \tag{27}$$

which is the *relative* intensity with which industry i uses occupation j in country m.

3.2.1 Effects of production in one industry on next-period production in another industry

Results (28) and (29) below answer the question, given a positive shock at date t to production in industry i in country m, what effect does this have on production at date $t+1$ in industry h in country n?

For any country m and any pair of industries i and h:

$$\frac{d\hat{y}_{h,m,t+1}}{d\hat{y}_{imt}} \approx [1 + (\frac{N-1}{N})\theta]\rho \sum_j \alpha_{hj}\tilde{\alpha}_{ijm} - [(\frac{1}{I})(\frac{N-1}{N})\theta]\rho \sum_{i'}\sum_j \alpha_{i'j}\tilde{\alpha}_{ijm} \tag{28}$$

and for any pair of countries m and $n \neq m$ and any pair of industries i and h:

$$\frac{d\hat{y}_{h,n,t+1}}{d\hat{y}_{imt}} \approx -[(\frac{1}{N})\theta]\rho \sum_j \alpha_{hj}\tilde{\alpha}_{ijm} + [(\frac{1}{I})(\frac{1}{N})\theta]\rho \sum_{i'}\sum_j \alpha_{i'j}\tilde{\alpha}_{ijm} \tag{29}$$

The intuition behind (28) and (29) is as follows. The increase in production in industry i in country m has a *direct effect* and an *indirect effect*.

The direct effect is as follows. For each occupation j, the extent to which an increase in production in industry i in country m corresponds to an increase in usage of occupation j relative to other occupations (and hence an increase in learning-by-doing in occupation j) is given by $\tilde{\alpha}_{ijm}$ (which, examining (27), can be positive or negative, since learning-by-doing is based on *relative* occupational usage). The extent to which this extra learning-by-doing in occupation j benefits industry h is given by α_{hj}. Hence, the size of learning spillovers between industries i and h is $\sum_j \alpha_{hj}\tilde{\alpha}_{ijm}$. If industries i and h are similar (dissimilar) enough to each other in their occupational usage, then $\sum_j \alpha_{hj}\tilde{\alpha}_{ijm}$ is greater (less) than zero, and the direct effect on industry h within country m is positive (negative), while it is negative (positive) in every other country, because in every other country industry h becomes relatively less (more) competitive compared to country m.

The direct effect is scaled by ρ, since ρ is the rate of learning-by-doing. Furthermore, the direct effect on each other country is scaled by $\frac{1}{N}$, where N is the number of countries, as well as θ, since θ is the trade elasticity. This is balanced by the fact that the direct effect on country m itself is scaled by $[1+(\frac{N-1}{N})\theta]$; note that $[1+(\frac{N-1}{N})\theta]-(N-1)(\frac{1}{N})\theta=1$, i.e., the scale factors on the direct effects across the world sum to one.

The indirect effect on industry h is as follows. Industry h is, of course, not the only industry directly affected by industry i. Summing the term $\sum_j \alpha_{hj}\tilde{\alpha}_{ijm}$ across all industries gives us $\sum_{i'}\sum_j \alpha_{i'j}\tilde{\alpha}_{ijm}$, which is the size of the total learning spillovers from industry i to all other industries – or, using network terminology, it is the first-degree centrality of industry i in the network of industries. If industry i is sufficiently central (sufficiently peripheral), then $\sum_{i'}\sum_j \alpha_{i'j}\tilde{\alpha}_{ijm}$ is greater (less) than zero, and the high (low) amount of learning-by-doing induced by the increase in production in industry i in country m raises (lowers) country m's equilibrium wage, which (all else being equal) makes each industry in country m less (more) competitive and makes each industry in every other country more (less) competitive.

As with the direct effect, the indirect effect is scaled by ρ, since ρ is the rate of learning-by-doing. Moreover, the indirect effect (which, bear in mind, is capturing an individual industry's effect on the entire economy) is scaled by $\frac{1}{I}$, where I is the number of industries. As with the direct effect, the indirect effect on each other country is scaled by $\frac{1}{N}$, where N is the number of countries, as well as θ, since θ is the trade elasticity. This is balanced by the fact that the indirect effect on country m itself is scaled by $(\frac{N-1}{N})\theta$; note that $-(\frac{N-1}{N})\theta+(N-1)(\frac{1}{N})\theta=0$, i.e., the scale factors on the indirect effects across the world sum to zero.

Note that if $N=1$ (i.e., there is only one country in the world), then the indirect effect is zero, and the total effect of the industry-i shock on industry h is $\rho\sum_j \alpha_{hj}\tilde{\alpha}_{ijm}$, which is exactly the same as the closed-economy results from section 2.3.[17]

3.2.2 Effects on each country's next-period welfare

Results (30) and (31) below answer the question, given the aforementioned positive shock at date t to production in industry i in country m, what effect does this have on the date $t+1$ welfare of country n?

For any country m and industry i:

$$\frac{d\hat{W}_{m,t+1}}{d\hat{y}_{imt}} \approx (\frac{1}{I})[1-(\frac{N-1}{N})(\frac{1}{1+\theta})]\rho\sum_{i'}\sum_j \alpha_{i'j}\tilde{\alpha}_{ijm} \tag{30}$$

and for any pair of countries m and $n \neq m$ and any industry i:

$$\frac{d\hat{W}_{n,t+1}}{d\hat{y}_{imt}} \approx (\frac{1}{I})[(\frac{1}{N})(\frac{1}{1+\theta})]\rho\sum_{i'}\sum_j \alpha_{i'j}\tilde{\alpha}_{ijm} \tag{31}$$

The intuition behind (30) and (31) partly carries over from the intuition above for the indirect effects in Results (28) and (29) – the effects of the date t shock to industry

[17]There is a trivial difference, namely, the $\tilde{\alpha}$ terms now (by (27)) include z terms, which made no appearance in the closed-economy results, but that was just because there were no z terms in the closed-economy model. If we were to add them in, we would get exactly the same result as here.

i on countries' welfare at date $t+1$ is a function of industry i's first-degree centrality $\sum_{i'}\sum_j \alpha_{i'j}\tilde{\alpha}_{ijm}$, and again this is scaled by ρ and $\frac{1}{I}$ (and by $\frac{1}{N}$ for countries other than m) for the same reasons as above.

Note, though, that the right-hand sides of (30) and (31) have the same sign rather than opposite signs; $\sum_{i'}\sum_j \alpha_{i'j}\tilde{\alpha}_{ijm}$ is greater (less) than zero when industry i is sufficiently central (peripheral) in the network that an increase in production in industry i in country m induces more (less) learning in the aggregate economy of country m, in which case other countries benefit (are hurt) as well, due to buying products from country m at a lower (higher) cost.

Furthermore, note that the effect on other countries is scaled by $\frac{1}{1+\theta}$ rather than θ; a higher θ dampens the effect on other countries rather than exacerbating it $-$ a higher θ means less heterogeneity in intra-industry productivity, meaning (all else being equal) international trade is less important for a country's welfare, meaning extra economy-wide learning in country i benefits other countries less. (This is in contrast with Results (28) and (29), which were looking at the effects on a specific industry h, which are exacerbated when intra-industry productivity varies less.)

Given that the effect on other countries is scaled by $(\frac{1}{N})(\frac{1}{1+\theta})$, this is balanced by the effect on country m itself being scaled by $[1-(\frac{N-1}{N})(\frac{1}{1+\theta})]$; note that $[1-(\frac{N-1}{N})(\frac{1}{1+\theta})] + (N-1)(\frac{1}{N})(\frac{1}{1+\theta}) = 1$, i.e., the scale effects on welfare across the world sum to one.

3.2.3 Effects on production and welfare more than one period ahead

Results (28) through (31) are only telling us the *next-period* effects of a shock to production in industry i in country m; now let us consider the effects arbitrarily far into the future. First we will consider the effects over time on production in each industry in each country. Let A_{mn} denote the I X I matrix whose (i,h) element is $\frac{d\hat{y}_{h,n,t+1}}{d\hat{y}_{imt}}$ (which we found an approximation for above, which does not depend on t). Let A be the (NI) X (NI) matrix formed by appending the A_{mn} matrices to each other, so that the (m,n) block of A is A_{mn}.

Start from an arbitrary equilibrium path $\{\hat{y}^{\star}_{imt}\}_{i,m,t}$ and consider an arbitrary vector of shocks to production in each industry in each country at date t: let y_t be the (NI)-dimensional vector whose first I elements are $\hat{y}_{i1t}-\hat{y}^{\star}_{i1t}$ for each industry i in country 1; the next I elements of y_t are $\hat{y}_{i2t}-\hat{y}^{\star}_{i2t}$ for each industry i in country 2; and so on.

If we take the first-order approximations that we found above and suppose that they approximately hold at any arbitrary point, then we have the result that for any date t and any length of time τ beyond t:

$$y_{t+\tau} \approx (A')^{\tau} y_t \tag{32}$$

While the single-closed-economy model involved a network of industries, (32) is saying that we can think of this multi-country, open-economy model as involving a network of countries and industries, where each node in the network is a country-industry pair, and the network matrix A (whose entries we found above) gives us the effect of an increase in production in industry i in country m on every other industry in every other country, with these effects being the aforementioned sum of direct learning spillovers and general equilibrium effects via international trade.

Now let us consider the effects of this shock to production in industry i in country m on each country's discounted sum of welfare, summing from date t to ∞. Let w_n be the (NI)-dimensional vector whose first I elements are $\frac{d\hat{W}_{n,t+1}}{d\hat{y}_{i1t}}$ for each industry i in country 1 (which we found an approximation for above, which does not depend on t), whose second I elements are $\frac{d\hat{W}_{n,t+1}}{d\hat{y}_{i2t}}$ for each industry i in country 2, and so on. Let $\bar{W}_n$ denote the discounted sum of country n's logarithm of welfare over time, discounted at the rate δ − that is, $\bar{W}_n \equiv \sum_{t=0}^{\infty} \delta^t \hat{W}_{nt}$.

Note, then, that $\frac{d\overline{W}_n}{dy_t}$ is the (NI)-dimensional vector whose first I entries are $\frac{d\overline{W}_n}{d\hat{y}_{i1t}}$ for each industry i in country 1, whose second I entries are $\frac{d\overline{W}_n}{d\hat{y}_{i2t}}$ for each industry i in country 2, and so on. In other words, $\frac{d\overline{W}_n}{dy_t}$ is the vector telling us how much a shock at date t to each country-industry pair affects country n's discounted infinite sum of welfare from date t onward. For any arbitrary country n, we have the following result:

$$\frac{d\overline{W}_n}{dy_t} \approx \delta w_n + \delta^2 A w_n + \delta^3 A^2 w_n + ... \tag{33}$$

Letting $\mathbb{I}$ denote the identity matrix, we can write this as

$$\frac{d\overline{W}_n}{dy_t} \approx \delta(\mathbb{I} + \delta A + \delta^2 A^2 + ...) w_n \tag{34}$$

And so we have, for any arbitrary country n:

$$\frac{d\overline{W}_n}{dy_t} \approx \delta(\mathbb{I} - \delta A)^{-1} w_n \tag{35}$$

The right-hand side of (35) is the vector of each country-industry pair's *Bonacich centrality* (from country n's perspective) in the network of country-industry pairs. Recall from section 2.3 that for any arbitrary network, a node's Bonacich centrality is equal to the sum of its first-degree links with every other node discounted by a discount factor δ, plus the sum of its second-degree links with other nodes discounted by δ^2, and so on. In this case the links are weighted by the vector w_n, which tells us how much a shock to production in a given country-industry pair in a given time period affects country n's next-period welfare − which, as we found above, relates to each country-industry pair's first-degree centrality in the network.

4 The evolution of comparative advantage

The theoretical analysis showed how, in the presence of occupational learning, countries with a comparative advantage in industries that are more central in the network of occupational learning spillovers will grow more in the aggregate. How does the pattern of comparative advantage itself evolve over time, as workers in different countries are learning different things? In this section, I answer this question theoretically and empirically.

I start in Section 4.1 by deriving the model-implied dynamics of comparative advantage. To aid intuition, I analyze in section 4.2 a stripped-down version of the model with only industry-level learning-by-doing (i.e., without any occupational dimension to learning), in which case the evolution of comparative advantage takes a particularly simple form. In Section 4.3 I provide evidence in support of the mechanism underlying the model. In particular, a country's growth in comparative advantage in an industry is positively correlated with its initial level of comparative advantage in occupationally similar industries.

4.1 Model-implied dynamics of comparative advantage

I use the open-economy model of Section 3 to provide an analytical characterization of the relationship between learning-by-doing and the evolution of comparative advantage. I will use this characterization in Section 5 to calibrate the learning-by-doing parameters to match the observed changes over time in trade data.

Let e_{imnt} denote the date-t dollar value of exports in industry i from country m to country n. In the CDK model, the logarithm of e_{imnt} is equal to a sum of exporter-importer, importer-industry, and exporter-industry dummies, plus an orthogonal error term:

$$\ln e_{imnt} = \delta_{mnt} + \delta_{int} + \delta_{imt} + \varepsilon_{imnt} \tag{36}$$

The pattern of comparative advantage (i.e., relative productivity differences across country-industry pairs) can be identified off of the exporter-industry dummies in the above regression.[18] In particular, note that according to the model[19],

$$\ln e_{imnt} = \delta_{mnt} + \delta_{int} + \theta \ln z_{im}[\sum_j (\alpha_{ij})^{\epsilon} (\phi_{ijmt})^{\epsilon-1}]^{\frac{1}{\epsilon-1}} \tag{37}$$

Combining equations (36) and (37), we have

$$e^{\frac{\delta_{imt}}{\theta}} = z_{im}[\sum_j (\alpha_{ij})^{\epsilon} (\phi_{ijmt})^{\epsilon-1}]^{\frac{1}{\epsilon-1}} \tag{38}$$

where the right-hand side of (38) is country m's productivity in industry i at date t.

Thus, given an estimate of the trade elasticity θ, running regression (36) and plugging the exporter-industry dummy coefficients δ_{imt} into equation (38) gives us estimates of country-industry-level productivity. Note that the degrees of freedom in regression (36) are such that δ_{imt} is only identified up to a double-normalization of $\delta_{im^{\star}t} = 1\ \forall i$ for some baseline country $m^{\star}$ and $\delta_{i^{\star}mt} = 1\ \forall m$ for some baseline industry $i^{\star}$ – in other words, δ_{imt} captures country m's *comparative advantage* in industry i at date t. The proposition below

[18] Meanwhile, the exporter-importer dummies account for bilateral trade costs (e.g., distance between countries), and the importer-industry dummies account for demand-side factors.

[19] Equation (37) can be derived by combining equations (16) and (22).

captures the model's predictions for how the pattern of comparative advantage evolves over time.

Proposition 4.1. *In the economic environment of section 3.1, with learning-by-doing governed by equation (25), country m's comparative advantage in industry i at date t, δ_{imt}, evolves from one period to the next in the following way:*

$$\delta_{i,m,t+1} - \delta_{imt} = \theta(\frac{1}{\epsilon - 1}) \ln(\sum_j (\alpha_{ij})^{\epsilon}[(1 + \tilde{L}_{ijmt})^{1-\sigma}(1 + \tilde{L}_{jmt})^{\sigma}]^{\rho(\epsilon-1)}) \tag{39}$$

Proof. See Appendix A.3. □

Proposition 4.1 shows that a country's change over time in its comparative advantage in industry i is a weighted average of the size of each occupation in that country, both within industry i and summing across industries, with the weight on an occupation given by how intensely industry i uses the occupation. Note that the size of each occupation in a country is itself, in equilibrium, determined by that country's comparative advantage in each industry, meaning that the *change* in a country's comparative advantage in an industry is a function of its initial *level* of comparative advantage in each industry.

This point − that in the presence of learning-by-doing, changes in comparative advantage are a function of levels of comparative advantage − is particularly easy to see when learning-by-doing is only at the industry level, without any occupational component. To that end, in the following section I analyze a stripped-down version of the model, in which there are no occupations and learning only happens at the industry level.

4.2 Dynamics of comparative advantage under purely industry-level learning-by-doing

Empirically, we observe that countries with a larger comparative advantage in an industry tend on average to experience *less* future growth in comparative advantage in that industry. This was first documented by Levchenko and Zhang (2016) and Hanson, Lind, and Muendler (2016), and is replicated in Section 4.3 of this paper. One might think that this empirical finding goes against the theory of learning-by-doing, at least at the industry level. Quoting from p. 106 of Levchenko and Zhang, "A strong implication of [learning-by-doing] is that relative productivity differences *increase* over time − comparative advantage strengthens. This is because learning is faster in sectors that produce more, and comparative advantage sectors are the ones that produce more."

In this section, I formally show that this argument only holds under certain strong assumptions about the learning process. I show that, in general, learning-by-doing is, in fact, consistent with the empirical evidence. In order to make the argument as simple as possible, in this section I drop the occupational dimension from the analysis − that is, the model in this section is the same as the model in Section 3, except there are no occupations, and learning-by-doing is at the industry level. More specifically, learning-by-doing is a function of industry-level output, and this learning-by-doing affects industry-level TFP.

Given that comparative advantage is estimated at the industry level (and not the occupation level), this makes the link between learning-by-doing and the evolution of comparative advantage particularly simple. The aim of this section is purely to crystalize our thinking on the relationship between learning-by-doing and the dynamics of comparative advantage; I will not be making use of the results in this section in the quantitative analysis in Section 5.

For the sake of simplicity, in this section I consider the case in which trade costs are zero, the size of the labor force in each country is equal to 1, and each country's representative household's utility function puts equal weight on each industry.

Given the above discussion, the learning-by-doing equation I will consider in this subsection takes the following form:

$$\frac{z_{i,m,t+1}}{z_{imt}} = (y_{imt})^{\rho} \tag{40}$$

where ρ is the rate of learning-by-doing, y_{imt} is country m's industry-i output at date t, and z_{imt} is country m's industry-i productivity[20] at date t, with z_{im0} given for each country m and industry i.

For each country m and each industry i, define country m's comparative advantage in industry i at date t (which I will denote by CA_{imt}) in terms relative to country M's productivity in industry I at date t – that is:

$$CA_{imt} \equiv \frac{\left(\frac{z_{imt}}{z_{Imt}}\right)}{\left(\frac{z_{iMt}}{z_{IMt}}\right)} = \frac{z_{imt} z_{IMt}}{z_{Imt} z_{iMt}} \tag{41}$$

Proposition 4.2. *In the economic environment described above, with industry-level learning-by-doing governed by equation (40) and comparative advantage defined by (41), the* ***change*** *over time in a country's comparative advantage in an industry is a function of that country's* ***level*** *of comparative advantage in that industry. In particular, as a first-order approximation:*

$$\frac{d\ln(CA_{i,m,t+1})}{d\ln(CA_{imt})} \approx 1 + \rho(1 + [\frac{I-1}{I}]\theta) \tag{42}$$

Proof. See Appendix A.4. □

Note from equation (42) that if $\rho > 0$, then $\frac{d\ln(CA_{i,m,t+1})}{d\ln(CA_{imt})} > 1$, i.e., learning-by-doing (under this particular formulation of learning) induces *divergence* over time in comparative advantage: if country 1 has a greater comparative advantage in an industry than country 2, then over time country 1 will have an *even greater* comparative advantage in that industry compared to country 2. Proposition 4.2 thereby formalizes the aforementioned intuition from p. 106 of Levchenko and Zhang (2016).

It turns out, however, that this theoretical prediction of divergence in comparative advantage is, in part, an artifact of the particular formulation of learning assumed in

[20]More precisely, z_{imt} is the parameter of the Fréchet distribution that governs country m's average productivity in industry i at date t, averaged across the infinitely many varieties of industry i.

equation (40). There are two aspects of this formulation that are key to generating the divergence result: (1) the learning-by-doing is purely at the industry level, and (2) the elasticity of learning with respect to output is constant (given by the single parameter ρ). If we relax either one of these two assumptions, then learning-by-doing does not necessarily generate divergence in comparative advantage.

For the sake of the argument, I illustrate below the implications of relaxing assumption (2). In particular, suppose that learning-by-doing exhibits decreasing returns – that is, for a given amount of date-t output in industry i, a country will learn more in industry i at date t if its date-t productivity in industry was low to start with. We can think of this as there being "low-hanging fruit" when one is just starting to learn something, while the more advanced one becomes, the harder and harder it is to become yet more advanced. Then, whether learning-by-doing induces convergence or divergence in comparative advantage depends on the relative sizes of two opposing forces.

Imagine comparing two countries, one of which is more productive than the other in industry i at date t. Will the backward country converge to the advanced one, or will the two countries further diverge? On the one hand, due to the decreasing returns in the learning-by-doing function, if the two countries were to produce at date t the same amount of industry i output, the more backward country would learn more – the size of this effect is captured below by a parameter λ. On the other hand, the two countries will obviously *not* produce the same amount of industry i output; the more advanced country will produce more, due to having a comparative advantage in it – the size of this effect is captured by the trade elasticity θ. Hence, whether there is convergence or divergence depends on the relative sizes of λ vs. θ. The tension between these two forces is captured by the proposition below.

Proposition 4.3. *Suppose we have the same economic environment as in Proposition 4.2, except there are diminishing returns in the learning-by-doing function. Specifically, suppose the learning-by-doing equation is as follows:*

$$\frac{z_{i,m,t+1}}{z_{imt}} = (z_{imt})^{-\lambda}(y_{imt})^{\rho} \tag{43}$$

where $\lambda \geq 0$ governs the rate of diminshing returns to learning-by-doing.

Then learning-by-doing induces convergence or divergence in comparative advantage, depending on how quickly returns to learning-by-doing are diminishing vs. the size of the trade elasticity (scaled by the number of industries). Specifically, as a first-order approximation, there is convergence in comparative advantage if $\frac{\lambda}{\rho} > 1 + (\frac{I-1}{I})\theta$, while there is divergence if $\frac{\lambda}{\rho} < 1 + (\frac{I-1}{I})\theta$.

Proof. Through the same steps as in the proof of Proposition 4.2,

$$\frac{d\ln(CA_{i,m,t+1})}{d\ln(CA_{imt})} \approx 1 - \lambda + \rho(1 + [\frac{I-1}{I}]\theta) \tag{44}$$

Whether there is convergence or divergence in comparative advantage corresponds to whether this expression is greater than or less than one. Hence, there is convergence if $\frac{\lambda}{\rho} > 1 + (\frac{I-1}{I})\theta$, while there is divergence if $\frac{\lambda}{\rho} < 1 + (\frac{I-1}{I})\theta$. $\square$

This section and the previous section have shown theoretically how learning-by-doing induces a particular relationship between the cross-sectional pattern of comparative advantage, and the evolution of comparative advantage from one period to the next. Given these theoretical findings, in the following section I provide evidence of such a relationship.

4.3 Evidence on the dynamics of comparative advantage

Section 4.2 has shown how, in the presence of industry-level learning-by-doing, the *change* in a country's comparative advantage in an industry is a function of its *level* of comparative advantage in that industry. Section 4.1 showed how, in the presence of occupational learning, the change in country m's comparative advantage in industry i is a function of country m's level of usage of each occupation, which is itself a function of country m's level of comparative advantage in each industry h, to the extent that industry h uses similar occupations as i. In this section I provide evidence of these relationships between levels of comparative advantage and changes in comparative advantage, both within industries and across industries, as a function of their occupational similarity. I do so through the following regression:

$$\ln(CA_{i,m,t+\Delta}) - \ln(CA_{imt}) = \beta_0 + \beta_1 \ln(CA_{imt}) + \beta_2 \ln(OccCA_{imt}) + \epsilon_{imt} \quad (45)$$

where CA_{imt} is country m's comparative advantage in industry i at date t, and $OccCA_{imt}$ is a weighted average of country m's date-t comparative advantage in every industry h other than industry i, with each industry h weighted by its occupational similarity to industry i. The details of the construction of these variables are provided in Section 4.3.1.

4.3.1 Data

In order to estimate comparative advantage using the CDK method, I use international trade data from Feenstra et al (2005). The data report bilateral exports among 72 countries, at the 4-digit SITC Revision 2 product level, annually over the years 1962-2000. Hence, an example of an observation in this dataset is that, in 1978, Japan exported to Italy \$2,447,000 (in 1978 nominal US dollars) worth of silk worm cocoons and silk waste (SITC Rev. 2 code 2614). The number of products that appear in the data gradually increase over time from 696 in 1962 to 1288 in 2000.[21]

In order to merge these data with the industry-occupation table described below, I reclassify exports from 4-digit SITC Rev. 2 product codes into 3-digit 1997 NAICS industry codes, using a concordance table from Feenstra and Lipsey (n.d.). I restrict the sample

[21] These data include the dollar value reported by the importing country as well as the value reported by the exporting country; I follow the standard practice of using the value reported by the importing country, which is generally seen as more reliable, since countries have more of an incentive to carefully keep track of goods entering their borders than leaving them.

to manufacturing industries that appear in the industry-occupation table, and I further restrict the sample to countries and industries that appear across the sample period. We are left, then, with 44 countries and 20 industries, which are listed in Appendix B, in Tables 6 and 7, respectively.

Regression (45) requires cross-sectional trade data from two different points in time, t and $t+\Delta$. Rather than estimating each country's comparative advantage in each industry each year, and then examining year-by-year fluctuations (which could be confounded by business cycle phenomena far removed from the subject of this paper, as well as a large amount of measurement error), I follow Levchenko and Zhang (2016) and Hanson, Lind, and Muendler (2016) ("L&Z" and "HLM") by averaging e_{imnt} (exports of industry i from country m to country n in year t) by decade, and then using these averaged export data to estimate each country's average comparative advantage in each industry each decade. Comparative advantage is estimated following CDK; that is, $\ln(CA_{imt}) = \delta_{imt}/\theta$, following equation (38).

In order to construct the variable $OccCA_{imt}$, I use 2013 US data on how many people each industry employs in each occupation, from the Occupational Employment Statistics (OES) program at the Bureau of Labor Statistics (BLS). The industries, as explained above, are classified at the 3-digit NAICS level, while the occupations are classified at the 2-digit Standard Occupational Classification (SOC) level; the industries and occupations are reported in Appendix B, in Tables 7 and 8, respectively. From these data I construct a simple reduced-form measure of occupational similarity $OccSim_{ih}$ between industries i and h, namely, the correlation coefficient between the industry-h vector of employment of each occupation (as a fraction of industry h's total employment) and the corresponding industry-i vector. $OccCA_{imt}$ is then defined as $\sum_h OccSim_{ih} CA_{imt}$.

4.3.2 Results

Table 1 reports the results of running regression (45) for each pair of decades within the span of the data. The first result worth noting is that the within-industry effect is always (with one exception) negative and significant − that is, if a country has a larger comparative advantage in an industry, then on average its future growth in comparative advantage in that industry is smaller.[22] This finding is in line with L&Z and HLM. As discussed in section 4.2, L&Z argue that this finding contradicts the theory of learning-by-doing (at least at the industry level), but as demonstrated in section 4.2, industry-level learning-by-doing is in fact consistent with this finding, if the within-industry returns to learning are decreasing at a rate fast enough in comparison with the size of the trade elasticity.

The second result worth noting in Table 1 is a significant and positive cross-industry effect, as a function of industries' occupational similarity. That is, holding fixed a country's comparative advantage in industry i, if that country has a higher comparative advantage in occupationally similar industries, then that country will on average have higher future growth in its comparative advantage in industry i.

[22] The one exception is from the 1960's to 1970's, in which the within-industry effect is statistically insignificant.

It is worth emphasizing that this is not simply regressing changes in comparative advantage on concurrent changes in comparative advantage in occupationally similar industries, which could easily be explained by any arbitrary occupation-specific shocks (e.g., shocks to occupation-specific education policy) that affect different industries to the extent they use those occupations. Instead, Table 1 is showing that future *changes* in comparative advantage are positively related to previous *levels* of comparative advantage in occupationally similar industries. This finding is suggestive of significant *dynamic occupation-based agglomeration economies*, with occupational learning being one possible mechanism behind this.

Table 1: Effect of the *level* of comparative advantage on *growth* in comparative advantage (within industries and across industries, as a function of their occupational similarity)

	$\ln(\frac{CA\_70's}{CA\_60's})$	$\ln(\frac{CA\_80's}{CA\_70's})$	$\ln(\frac{CA\_90's}{CA\_80's})$	$\ln(\frac{CA\_80's}{CA\_60's})$	$\ln(\frac{CA\_90's}{CA\_70's})$	$\ln(\frac{CA\_90's}{CA\_60's})$
ln(CA_60's)	0.0243			-0.859***		-0.249***
	(0.0359)			(0.0253)		(0.0626)
ln(OccCA_60's)	**0.113****			**0.139*****		**0.233*****
	(0.0471)			**(0.0281)**		**(0.0858)**
ln(CA_70's)		-0.856***			-0.290***	
		(0.0201)			(0.0455)	
ln(OccCA_70's)		**0.103*****			**0.181****	
		(0.0252)			**(0.0725)**	
ln(CA_80's)			-0.347***			
			(0.122)			
ln(OccCA_80's)			**0.736*****			
			(0.132)			
Observations	880	880	880	880	880	880
R-squared	0.033	0.789	0.051	0.699	0.068	0.029

Note: each observation is at the country-by-industry level. CA_t is country m's comparative advantage in industry i at date t, estimated following Costinot, Donaldson, and Komunjer (2012). OccCA_t is a weighted average of country m's date-t comparative advantage in every industry h other than industry i, with each industry h weighted by its occupational similarity with industry i. All variables averaged by decade. Constant terms not reported in this table. Standard errors (in parentheses) clustered by country (44 countries). *** Significant at 1% level, ** 5% level, * 10% level.

5 Quantitative analysis

In the theoretical analysis of this paper I have shown how, in the presence of occupational learning-by-doing, a country grows more when it has comparative advantages in industries that are more central in the network of inter-industry occupational learning diffusion, and the size of a country's comparative advantage in an industry itself evolves over time as a function of the size of different occupations in that country. This raises the question of how quantitatively important is this occupation-based inter-industry network structure in explaining the evolution of comparative advantage and the importance of different industries to aggregate growth.

The key challenge in answering this question is in estimating the two key parameters that govern the dynamics of the model, namely, ρ and σ − the rate of occupational learning-by-doing, and the extent to which this learning spills over across industries, respectively. I use the model-predicted dynamics of comparative advantage, as characterized in section 4.1, to quantitatively discipline these two parameters. I describe below the calibration of the static parameters of the model and the subsequent estimation of the learning-by-doing parameters. I then use the calibrated model for two quantitative exercises. First, I quantify the importance of learning-by-doing to the dynamics of comparative advantage. Second, I assess how important each industry is to aggregate growth − in particular, how much an increase in a country's productivity in an industry boosts that country's equilibrium real income.

5.1 Calibration

In this section I describe the calibration of the static parameters of the model: the intensity α_{ij} with which industry i uses occupation j, the industry-i exponent β_i in the representative household's Cobb-Douglas utility function, the elasticity of substitution across occupations ϵ, the size of country m L_m, and the trade elasticity θ. The calibration of these parameters is summarized in Table 2.

I calibrate α_{ij} using 2013 US data on each industry's employment of each occupation from the Bureau of Labor Statistics. For each industry i and occupation j, I set α_{ij} equal to the number of people in occupation j that industry i employs, divided by the total number of people that industry i employs.

Under the assumption that preferences are constant over time and across countries, I calibrate β_i using the 2012 Use Table from the Bureau of Economic Analysis (BEA) in the US. For each industry i, I set β_i equal to personal consumption expenditures on industry i divided by total personal consumption expenditures.

As described in Section 4.3.1, industries are classified at the 3-digit NAICS level, while occupations are classified at the 2-digit Standard Occupational Classification (SOC) level. The industries and occupations are reported in Appendix B, in Tables 7 and 8, respectively.

To calibrate the size of each country, I use the World Development Indicators (WDI) from the World Bank.[23] Specifically, for each country m, I set L_m so that $w_m L_m$ equals country m's GDP.

There is a substantial literature that provides a range of estimates of the trade elasticity θ. I borrow the value 6.53 from CDK, as their trade model has the same cross-sectional

[23] Taiwan is not included in the World Development Indicators, so for Taiwan I use the Monthly Bulletin of Statistics of the Republic of China.

structure as mine, and they estimate the trade elasticity using data from the same set of manufacturing industries as in my paper. For the elasticity of substitution across occupations, ϵ, I borrow the estimate 0.9 from Goos et al (2014), which is also used by Lee (2017). Goos et al. estimate this elasticity off of the observed correlation between the level of demand for an occupation within an industry and a measure of industry marginal costs, using industry-occupation employment data at the same level of aggregation as my paper.

Table 2: Summary of calibration of static parameters

Parameter	Meaning	Value	Target / source
α_{ij}	Occupation j's weight in industry i's prod. fun.	–	US employment data (BLS)
β_i	Industry i's weight in utility function	–	US output data (BEA)
ϵ	Elasticity of substitution across occupations	0.9	Goos et al. (2014)
L_m	Size of country m	–	World Development Indicators
θ	Trade elasticity	6.53	CDK (2012)

5.2 Estimating learning-by-doing and quantifying its relevance to dynamics of comparative advantage

Following CDK, I estimate each country's comparative advantage in each industry for each time period. (See Section 4.1.) I jointly estimate the learning-by-doing parameters ρ and σ to minimize the sum of squared errors between the model's predictions and the data, with regard to the moment condition from (39) linking changes in comparative advantage to previous levels of employment − that is:

$$(\hat{\rho}, \hat{\sigma}) = \underset{\rho,\sigma}{\operatorname{argmin}} \sum_i \sum_m \{\varepsilon_{imt}(\rho, \sigma)\}^2 \tag{46}$$

where

$$\varepsilon_{imt}(\rho, \sigma) \equiv \delta_{i,m,t+1} - \delta_{imt} - \theta(\frac{1}{\epsilon - 1}) \ln(\sum_j (\alpha_{ij})^{\epsilon}[(1 + \tilde{L}_{ijmt})^{1-\sigma}(1 + \tilde{L}_{jmt})^{\sigma}]^{\rho(\epsilon-1)}) \tag{47}$$

This estimation strategy exploits the fact that in the presence of occupational learning-by-dong, a country's *change* over time in its comparative advantage in an industry is a weighted average of the *size* of each occupation in that country, both within that industry and across other industries, with the weight on an occupation determined by how intensely that industry uses each occupation.

Note that this estimation strategy requires two different cross sections of trade data. I average the export data into ten-year bins, estimating each country m's average comparative advantage in each industry i over ten-year periods, plugging these estimates δ_{imt} into (47) in order to calibrate ρ and σ. Table 3 below shows the resulting values of ρ and σ for each possible combination of pairs of time periods, along with bootstrapped 95% confidence intervals.[24]

Table 3: Best-fit values of ρ and σ, and 95% confidence intervals

Time period	ρ (rate of learning)	σ (extent of inter-industry diffusion)
60's to 70's	12.5 [8.4, 18.1]	0.74 [0.49, 1.02]
70's to 80's	10.4 [4.8, 14.6]	0.96 [0.68, 1.97]
80's to 90's	10.2 [7.4, 16.7]	0.72 [0.46, 1.08]
60's to 80's	10.8 [6.8, 14.7]	0.93 [0.66, 1.50]
70's to 90's	10.9 [7.5, 15.1]	0.78 [0.54, 1.10]
60's to 90's	10.7 [8.1, 15.4]	0.81 [0.52, 1.08]

The estimates of ρ range from 10.2 to 12.5. A value of 12.5 for ρ means that a one percentage point increase in the share of a country's labor force in an occupation is associated with a 12.5% higher productivity in that occupation in that country from one decade to the next. Estimates of σ range from 0.72 to 0.96. A value of 0.74 for σ means that 74% of occupational learning spills over across industries. The fact that, in each of the above estimates, $\sigma > 0.5$ means that the occupational dimension to learning is more important the industry dimension.

How well does the calibrated model match the observed dynamics of comparative advantage? Table 4 shows the correlation ("*Corr*") between the calibrated model's predictions and the data, with regard to changes over time in comparative advantage (that is, the correlation between the left-hand and right-hand sides of (39) given the calibrated values

[24]Confidence intervals were obtained by randomly drawing a subsample of 22 countries and 10 industries, finding the $\hat{\rho}$ and $\hat{\sigma}$ that solve minimization problem (46) within this subsample of countries and industries, and doing this 1000 times to obtain a distribution of 1000 values of $\hat{\rho}$ and $\hat{\sigma}$ for each time period. The resulting histograms for $\hat{\rho}$ and $\hat{\sigma}$ for the 1960's to 1970's time period are shown in Figures 5 and 6 in Appendix B. The histograms for the other time periods are available upon request.

of ρ and σ). This is shown for each possible time horizon, namely, changes in comparative advantage over ten, twenty, and thirty years. The correlation between the model's predictions and the data is positive over each time horizon: 0.09 over a ten-year time horizon, 0.05 over twenty years, and 0.07 over thirty years.

For comparison, Table 4 also shows the corresponding correlation under the null model, in which $\sigma = 0$, i.e., in which learning is only within-industry. (Under this null model, I re-calibrate ρ so as to best fit the data under the condition $\sigma = 0$. The resulting value for ρ is also reported in Table 4.) The null model performs significantly worse, meaning that learning-by-doing helps explain the dynamics of comparative advantage significantly better when one accounts for inter-industry learning diffusion. More precisely, by this metric, accounting for inter-industry learning diffusion improves the fit of the model to the data by a factor of 4.5 when looking at changes in comparative advantage from the 60's to the 70's, and even more so when looking further out.

Table 4: Goodness of fit with regard to changes in comparative advantage

	60's to 70's	60's to 80's	60's to 90's
Best-fit ρ	12.5	10.8	10.7
Best-fit σ	0.74	0.93	0.81
$Corr$	**0.09**	**0.05**	**0.07**
ρ when $\sigma = 0$	19.1	18.0	16.8
$Corr$ when $\sigma = 0$	**0.02**	**-0.16**	**-0.04**

Note that the method in this section for estimating ρ and σ targets the average extent of *dynamic occupational agglomeration economies*. Among the possible mechanisms that can generate such dynamics, learning-by-doing has received particularly heavy emphasis in prior literature (as surveyed in Acemoglu (2008), and as discussed in Levchenko and Zhang (2016)), and hence is used to guide our thinking throughout this paper. However, other possible mechanisms include dynamic occupational economies of scale on the firms' side – i.e., when a larger number of people are employed in an occupation, employers better learn over time how to efficiently hire and make use of this occupation. The estimation strategy in (46) is not equipped to quantify the relative amounts of occupational learning-by-doing per se vs. other sources of dynamic occupation-based agglomeration economies, but rather it is intended to quantify the total importance of such occupation-based scale effects.

5.3 Quantifying each industry's importance to aggregate growth

The theoretical analysis showed how an industry contributes more to aggregate growth when it is more central in the network of occupational learning spillovers. The goal of this section is to quantify the size of this effect.

Consider a counterfactual 10% increase in a country m's productivity in some particular industry i, and hence an increase in m's comparative advantage in industry i. We can think of this shock as a simple analytical stand-in for a policy in country m that disproportionately benefits industry i – for example, a subsidy to research and development in that industry. How much does this raise the country's next-period real income? How does this depend on the particular industry i? In this section I carry out this exercise for each country and each industry in the sample – note that this is a separate counterfactual exercise for each country and industry; i.e., this is not meant to capture any interactions in the effects of shocks to different industries or countries.

As an illustration, Figure 3 plots the results for China, starting from the initial equilibrium of the 1960's. The dark blue bars in the graph show the equilibrium increase in China's GDP predicted by the calibrated model (with $\rho = 12.5$ and $\sigma = 0.74$), while the light grey bars show the increase in the case of the null model without any inter-industry learning diffusion, i.e., with $\sigma = 0$ (still with $\rho = 12.5$).[25] The ratio of the length of industry i's dark blue bar to the length of its light grey bar represents the extent to which industry i's contribution to aggregate growth is through the diffusion of learning in that industry into other industries, rather than the direct effect of the increase in industry i's productivity itself. Figure 4 plots the average ratio for each industry, averaging across each of the countries in the sample.

[25] Another natural benchmark is with no learning at all, i.e., with $\rho = 0$, the results of which are plotted in Figure 7 in Appendix B. The results are similar to the case in which $\sigma = 0$. Given that the rate of learning ρ is the same for each industry, the main driver of asymmetries across industries in their contribution to aggregate growth in this model is not learning in itself (i.e., $\rho > 0$), but rather the diffusion of occupational learning across industries (i.e., $\sigma > 0$) combined with the asymmetries in how intensely each industry uses each occupation.

Figure 3: Percent increase in China's GDP, given 10% increase in industry-i productivity

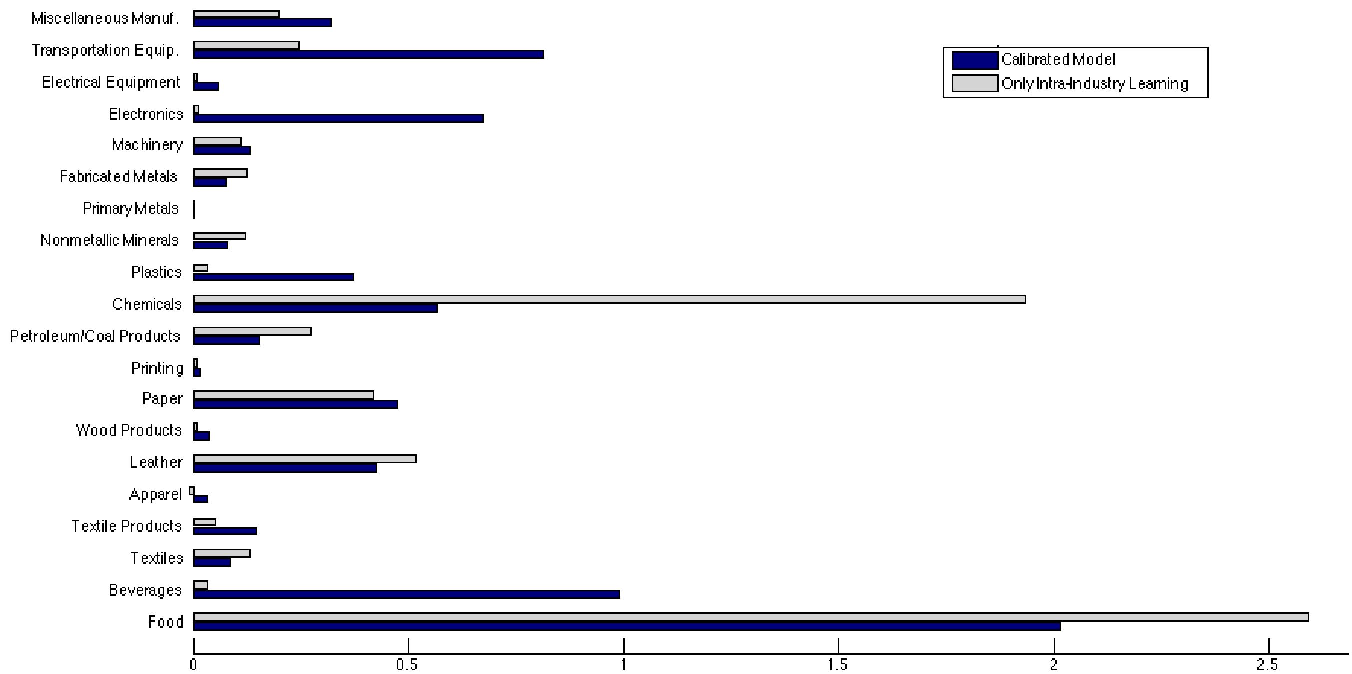

Figure 4: Effect of industry-i productivity shock on GDP: average ratio, calibrated model vs. only intra-industry learning

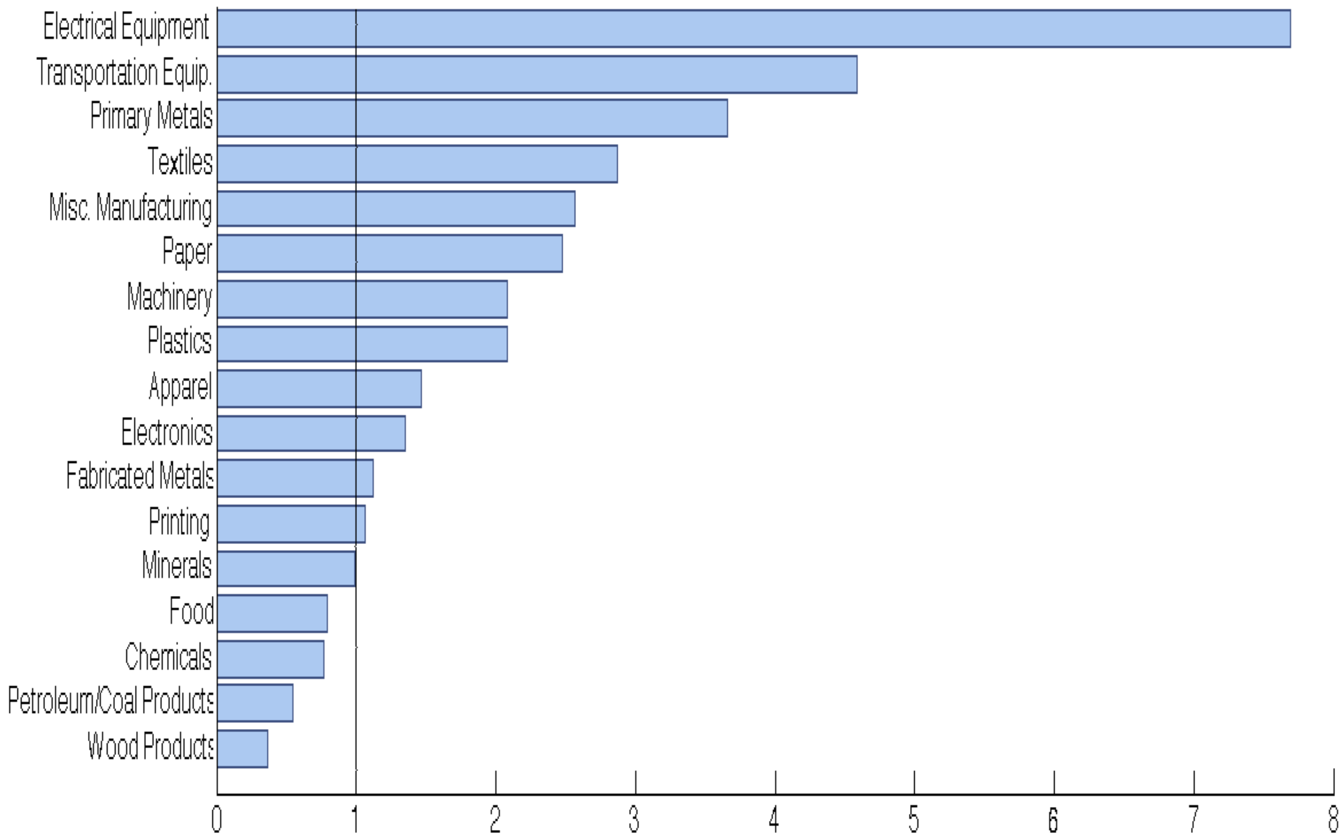

Note: this graph plots the geometric average (across each country m in the sample) of $\frac{\uparrow GDP_m(i)}{Null \uparrow GDP_m(i)}$ for each industry i, where $\uparrow GDP_m(i)$ is the prediction of the calibrated model (with $\rho = 12.5$ and $\sigma = 0.74$) for the percent increase in country m's next-period real income, given a 10% increase in country m's productivity in industry i, while $Null \uparrow GDP_m(i)$ is the corresponding prediction of the null model with no inter-industry learning diffusion, i.e., with $\sigma = 0$.

We can get an intuitive sense of what drives the differences across industries in Figures 3 and 4 by examining the diagrammatic representation of the occupation-based network of industries in Figure 2 from the introduction. For industries that are highly central in the network, such as primary metals, the null model with no inter-industry learning diffusion significantly understates their contribution to aggregate growth. Meanwhile, for industries that are peripheral in the network, such as petroleum and chemicals, an increase in productivity in these industries − and hence an increase in the country's comparative advantage in these industries − draws workers into these industries and away from industries that generate more inter-industry learning diffusion, meaning that the null model overstates their contribution to aggregate growth.

More precisely, recall the result from Equation (12) that, as a first-order approximation, the total impact of the inter-industry learning spillovers generated by an industry on a country's aggregate real income is given by that industry's Bonacich centrality in the network of occupational learning diffusion.[26] Table 5 below ranks industries by their Bonacich centrality, and also ranks industries by the ratios reported in Figure 4, i.e., by the model-predicted extent to which an industry's contribution to aggregate growth is through the diffusion of learning in that industry into other industries.

The two rankings in Table 5 do not perfectly match, as Equation (12) is only an approximation of the model derived in the special case of Leontief production, but the two rankings have a significantly positive correlation of 0.43. Hence, the intuitive notion of Bonacich centrality does a good job of capturing much of the model-predicted variation in how much inter-industry learning diffusion each industry generates.

The average discrepancy between the calibrated model and the model with no inter-industry learning diffusion (i.e., the average length of the bars in Figure 4) is a factor of 1.6, meaning that according to these estimates of the learning parameters, $\frac{1}{1.6}$ (62%) of the average industry's contribution to aggregate growth is through the direct effect of growth in that industry itself. This means that according to these estimates, a sizable fraction − 38% − of the average industry's contribution to aggregate growth is through the inter-industry learning spillovers that it generates.

Table 5: Industries ranked by amount of learning diffusion and by their centrality

Industry	Ranking by Inter-Industry Learning Diffusion	Ranking by Bonacich Centrality
Electrical Equipment	1	9
Transportation Equipment	2	11
Primary Metals	3	5
Textiles	4	2
Misc. Manufacturing	5	12
Paper	6	6
Machinery	7	7
Plastics	8	4
Apparel	9	1
Electronics	10	17
Fabricated Metals	11	3
Printing	12	8
Minerals	13	14
Food	14	13
Chemicals	15	16
Petroleum/Coal Products	16	15
Wood Products	17	10

Rank Correlation = 0.43

[26]Table 9 in Appendix B reports each industry's Bonacich centrality, as calculated using Equation (12).

6 Conclusions

This paper provides a theoretical and quantitative characterization of the link between occupational learning-by-doing and the evolution of comparative advantage, shedding new light on why the pattern of comparative advantage across countries and industries has changed over time. In so doing, it also provides a novel mechanism for why different industries contribute differently to long-run economic growth. In particular, this paper shows how occupational learning-by-doing induces an endogenous network of inter-industry learning spillovers, in which industries that are more central in this network generate more aggregate growth, and countries with comparative advantages in these more central industries grow more in equilibrium. These effects turn out to be quantitatively significant, with more than a third of the average industry's contribution to aggregate growth being in the form of inter-industry learning spillovers.

Although this paper's focus is on occupational learning at the country-industry level, there is scope for future research at finer levels of aggregation. Employer-employee matched data, for example, allow one to track the movement over time of individual workers across firms and industries, as well as the variation over time in firm-level productivity, allowing one to estimate how much workers' previous knowhow carries over to their new jobs, and to what extent their knowledge spills over to their new coworkers and employers.

Moreover, this paper focuses on one of possibly many components to learning, in particular the occupational component, as occupations are a particularly salient dimension through which productive knowledge can diffuse. However, there are other potentially important channels of learning as well; for example, firms might learn over time how to more efficiently use intermediate inputs. A quantitative examination of the relative importance of these and other learning channels to the evolution of comparative advantage is beyond the scope of this paper, but is an important area for future research.

Given the learning externalities that play a front-and-center role in the model, the model implies wide scope for welfare-enhancing policy. In some respects, this echoes a previous generation of theoretical literature that argued that, if some industries generate more learning spillovers than others, then there are grounds for governmental prioritization of the more learning-intensive industries. The novelty here is that, in the presence of occupational learning spillovers, even if industries are symmetric to one another in their rate of learning, there are grounds for prioritizing industries that are more central in the network of occupational learning diffusion. Such prioritization can take the form of industry-specific tariff policy, or credit subsidies to specific industries, which can be rationalized by this model even in the absence of financial frictions.

References

[1] Acemoglu, D. (2008). "Learning-by-Doing, Trade, and Growth." Section 19.7 of *Introduction to Modern Economic Growth*, Princeton University Press.

[2] Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). "The network origins of aggregate fluctuations." *Econometrica* 80 (5), 19772016.

[3] Alvarez and Lucas (2007), "General equilibrium analysis of the Eaton-Kortum model of international trade." *Journal of Monetary Economics*, Vol. 54, pp. 17261768.

[4] Arrow, Kenneth (1962). "The Economic Implications of Learning by Doing." *The Review of Economic Studies*, Vol. 29, No. 3, pp. 155-173.

[5] Autor, Dorn, and Hanson (2013), "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103(6).

[6] Bardhan (1971). "On Optimum Subsidy to a Learning Industry: An Aspect of the Theory of Infant-Industry Protection." *International Economic Review*, Vol. 12, No. 1, pp. 54-70.

[7] Carvalho, V. M. (2010). "Aggregate fluctuations and the network structure of intersectoral trade." Mimeo, CREI, Barcelona.

[8] Cavalcanti and Giannitsarou (2013). "Growth and Human Capital: A Network Approach." Working paper.

[9] Clemhout, S. and H. Y. Wan, Jr. (1970). "Learning-by-Doing and Infant Industry Protection." *The Review of Economic Studies*, Vol. 37, No. 1, pp. 33-56.

[10] Costinot, Donaldson, and Komunjer (2012), "What Goods Do Countries Trade? A Quantitative Exploration of Ricardo's Ideas." *Review of Economic Studies*, Vol. 79, pp. 581-608.

[11] Dehejia, Vivek (1993). "Optimal Endogenous Growth in a Two-Sector Model with Learning-by-Doing and Spillovers." *Journal of Economic Integration.*

[12] Eaton, Jonathan and Kortum, Samuel (1999), "International Technology Diffusion: Theory and Measurement," *International Economic Review*, 40, 537-570.

[13] Eaton, Jonathan and Kortum, Samuel (2002), "Technology, Geography, and Trade." *Econometrica*, Vol. 70, No. 5.

[14] Eaton, Jonathan and Kortum, Samuel (2012). "Putting Ricardo to Work." *Journal of Economic Perspectives*, Volume 26, Number 2, pp. 65?90.

[15] Ellison, Glaeser, and Kerr (2010). "What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns." *American Economic Review* 100 (June 2010): 1195?1213.

[16] Feenstra and Lipsey (n.d.). "Concordance between SITC Rev 2 at the 4-digit level and NAICS (1997)." Downloaded from < http://www.nber.org/lipsey/sitc22naics97/>.

[17] Feenstra, Lipsey, Deng, Ma, and Mo (2005). "World Trade Flows: 1962-2000." NBER Working Paper 11040.

[18] Fogli, A. and L. Veldkamp (2012). "Germs, social networks and growth." NBER Working Paper 18470.

[19] Galor, Oded and Mountford, Andrew (2008). "Trading Population for Productivity: Theory and Evidence." *Review of Economic Studies* 75, pp. 1143-1179.

[20] Goos, Manning, and Salomons (2014), "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 104, 2509-26.

[21] Hamilton, Alexander (1791). "Report on Manufactures." Communicated to the United States House of Representatives on December 5, 1791.

[22] Hanlon and Miscio (2016), "Agglomeration: A Long-Run Panel Data Approach." Working paper.

[23] Hanson, Lind, and Muendler (2016), "The Dynamics of Comparative Advantage." NBER Working Paper No. 21753.

[24] Hausmann, Hwang, and Rodrik (2007), "What you export matters," *Journal of Economic Growth* 12:125.

[25] Hidalgo, C. A.; Klinger, B.; Barabasi, A. L.; and Hausmann, R (2007). "The Product Space Conditions the Development of Nations." *Science*. Vol. 317, Iss. 482.

[26] Hirschman, Albert O. (1958). *The Strategy of Economic Development*. Yale University Press: New Haven, CT.

[27] Jones, Charles (2011). "Misallocation, Economic Growth, and Input-Output Economics." NBER Working Paper 16742.

[28] Jorgenson, Dale (2007), "35 Sector KLEM." Downloaded from ¡ http://scholar.harvard.edu/jorgenson/data ¿

[29] Kortum, Samuel (1997), "Research, Patenting, and Technological Change," *Econometrica*, 65, 1389-1419.

[30] Lee, Eunhee (2017), "Trade, Inequality, and the Endogenous Sorting of Heterogeneous Workers," Working Paper.

[31] Levchenko and Zhang (2016), "The evolution of comparative advantage: Measurement and welfare implications." *Journal of Monetary Economics*, Vol. 78, pp. 96-111.

[32] Long, John; and Plosser, Charles (1983), "Real Business Cycles." *Journal of Political Economy*, vol. 91, no. 1.

[33] Lucas, Robert E. Jr. (1988). "On the Mechanics of Economic Development." *Journal of Monetary Economics* 22, pp. 3-42.

[34] Matsuyama, Kiminori (1992). "Agricultural Productivity, Comparative Advantage, and Economic Growth." *Journal of Economic Theory*. Vol. 58, Iss. 2.

[35] Matsuyama, Kiminori (2008). "Structural Change." From *The New Palgrave Dictionary of Economics*, Second Edition, edited by Steven Durlauf and Lawrence Blume.

[36] Oberfield, Ezra (2013). "Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture." Working Paper.

[37] Ramey, Valerie (2011). "Identifying Government Spending Shocks: It's all in the Timing." *Quarterly Journal of Economics*, Vol. 126 (1), pp. 1-50.

[38] Ramey, Valerie; and Matthew Shapiro (1998). "Costly capital reallocation and the effects of government spending." *Carnegie-Rochester Conference Series on Public Policy*, Vol. 48, pp. 145-194.

[39] Ramondo, Rodriguez-Clare, and Saborio-Rodriguez (2016). "Trade, Domestic Frictions, and Scale Effects." *American Economic Review*, Vol. 106 (10), pp. 3159-3184.

[40] Rodriguez-Clare, Andres (1996a). "Multinationals, Linkages and Economic Development." *American Economic Review*. Vol. 86, No. 4.

[41] Rodriguez-Clare, Andres (1996b). "The Division of Labor and Economic Development." *Journal of Development Economics*. Vol. 49, p. 3-32.

[42] Rodrik (2006). "Industrial development: stylized facts and policies." Chapter prepared for the U.N.-DESA publication *Industrial Development for the 21st Century*.

[43] Rose, Andrew K. (2006). "Size Really Doesn't Matter: In Search of a National Scale Effect." *Journal of the Japanese and International Economies* 20 (4): 482-507.

[44] Succar, Patricia (1987). "The Need For Industrial Policy In LDC's - A Re-statement Of The Infant Industry Argument." *International Economic Review*.

[45] Young (1991). "Learning by Doing and the Dynamic Effects of International Trade." *The Quarterly Journal of Economics*, Vol. 106, No. 2, pp. 369-405.

A Derivation of theoretical results

A.1 Derivation of Equation 11

We will log-linearize around the symmetric case in which, at some date t, all occupational productivities are the same and production is equal across industries, that is, $\phi_{jt} = \phi_{kt}$

$\forall j, k$ and $Y_{it} = Y_{ht}$ $\forall i, h$. Combining the equilibrium equations (7) through (10) (with Cobb-Douglas utility and Leontief production, i.e., $\gamma \to 1$ and $\epsilon \to 0$), we have

$$\begin{aligned} log(Y_{h,t+1}) \approx log(\beta_h L_{t+1}) - \sum_j \alpha_{hj} log(\phi_{jt}) \\ + \rho \sum_j \alpha_{hj} log(\sum_i \alpha_{ij}) + \rho \sum_j \alpha_{hj} \sum_i (\frac{\alpha_{ij}}{\sum_m \alpha_{mj}}) log(Y_{it}) \\ - \rho \sum_j \alpha_{hj} log(I) - \rho \sum_j \alpha_{hj} \sum_i (\frac{1}{I}) log(Y_{it}) \end{aligned} \quad (48)$$

By taking the derivative of (48) with respect to $log(Y_{it})$, we see that a 1% increase in Y_{it} causes approximately an A_{ih}% increase in $Y_{h,t+1}$, where

$$A_{ih} \equiv \rho \sum_j \alpha_{hj} (\frac{\alpha_{ij}}{\sum_m \alpha_{mj}} - \frac{1}{I}) \quad (49)$$

A.2 Derivation of Equation 12

Consider an arbitrary equilibrium path $\{Y^{\star}_{it}\}_{i,t}$ and corresponding $\{U^{\star}_t\}_t$. Suppose, starting from this equilibrium path, we increase Y_{it} by 1% for some industry i at some date t. Let us now calculate the effect that this has on the total discounted utility of the representative household from date t onward.

Let lowercase letters denote log-deviations from the previous equilibrium path. Specifically, let $y_{it} \equiv log(Y_{it}) - log(Y^{\star}_{it})$ and $u_t \equiv log(U_t) - log(U^{\star}_t)$. Let $\vec{y_t}$ denote the I-dimensional vector specifying y_{it} for each industry i. Let u denote the discounted sum of log-deviations in utility over time, i.e., let $u \equiv \sum_t \delta^t u_t$, where $0 < \delta < 1$ is the representative household's discount factor. Let β denote the I-dimensional vector specifying the industry-i exponent in the representative household's Cobb-Douglas utility function, and let $\mathbb{I}$ denote the IxI identity matrix.

Using the approximation from above that a 1% increase in Y_{it} causes an A_{ih}% increase in $Y_{h,t+1}$, where A_{ih} is defined by (49), then if we let A denote the IxI matrix whose (i, h) element is A_{ih}, then $\vec{y}_{t+1} \approx A\vec{y_t}$. Let us now calculate $\frac{du}{d\vec{y_t}}$, which is the I-dimensional vector whose ith element gives the total discounted sum of percent increases in the representative household's utility from a 1% increase in Y_{it}:

$$\begin{aligned} \frac{du}{d\vec{y_t}} \approx \beta + \delta A\beta + \delta^2 A^2 \beta + \delta^3 A^3 \beta + \dots \\ = \beta + \delta(\mathbb{I} + \delta A + \delta^2 A^2 + \dots) A\beta \\ = \beta + \delta(\mathbb{I} - \delta A)^{-1} A\beta \end{aligned} \quad (50)$$

The right-hand side of (50) is the I-dimensional vector whose ith element is the Bonacich centrality of industry i in the network of industries, with each industry i weighted by its exponent β_i in the representative household's Cobb-Douglas utility function.

A.3 Proof of Proposition 4.1

Combining equations (25) and (38), we have, at date $t+1$,

$$e^{\frac{\delta_{i,m,t+1}}{\theta}} = z_{im}\{\sum_j (\alpha_{ij})^{\epsilon}(\phi_{ijmt}[(1+\tilde{L}_{ijmt})^{1-\sigma}(1+\tilde{\tilde{L}}_{jmt})^{\sigma}]^{\rho})^{\epsilon-1}\}^{\frac{1}{\epsilon-1}} \tag{51}$$

Given the log-linearity of the learning-by-doing equation[27], we can normalize ϕ_{ijmt} to 1 for every industry-i-occupation-j-country-m. We then have, at date t,

$$e^{\frac{\delta_{imt}}{\theta}} = z_{im} \tag{52}$$

while at date $t+1$ we have

$$e^{\frac{\delta_{i,m,t+1}}{\theta}} = z_{im}\{\sum_j (\alpha_{ij})^{\epsilon}[(1+\tilde{L}_{ijmt})^{1-\sigma}(1+\tilde{\tilde{L}}_{jmt})^{\sigma}]^{\rho(\epsilon-1)}\}^{\frac{1}{\epsilon-1}} \tag{53}$$

Dividing equation (53) by equation (52) and taking logs on both sides[28], we have our result:

$$\delta_{i,m,t+1} - \delta_{imt} = \theta(\frac{1}{\epsilon-1})\ln\{\sum_j (\alpha_{ij})^{\epsilon}[(1+\tilde{L}_{ijmt})^{1-\sigma}(1+\tilde{\tilde{L}}_{jmt})^{\sigma}]^{\rho(\epsilon-1)}\} \tag{54}$$

This proves the proposition. ■

A.4 Proof of Proposition 4.2

A change in CA_{imt} changes $CA_{i,m,t+1}$ through z_{imt}'s effect (via learning-by-doing) on $z_{i,m,t+1}$, which operates both directly and via changes in y_{imt}. Formally,

$$\frac{d\ln(CA_{i,m,t+1})}{d\ln(CA_{imt})} = \frac{d\ln(z_{imt})}{d\ln(CA_{imt})}(\frac{d\ln(z_{i,m,t+1})}{d\ln(z_{imt})}\frac{d\ln(CA_{i,m,t+1})}{d\ln(z_{i,m,t+1})} + \frac{d\ln(y_{imt})}{d\ln(z_{imt})}\frac{d\ln(z_{i,m,t+1})}{d\ln(y_{imt})}\frac{d\ln(CA_{i,m,t+1})}{d\ln(z_{i,m,t+1})}) \tag{55}$$

[27]That is, given the fact that, for any given specification of date-t labor usage in each industry and occupation, the percent change in productivity from date t to $t+1$ from learning-by-doing does not depend on the date-t level of producivity ϕ_{ijmt}.

[28]Note that I am assuming here that z_{im} − the parameter governing country m's average TFP in industry i − does not change over time. For the purpose of this characterization of the link between learning-by-doing and dynamics of comparative advantage, this assumption is WLOG: suppose, to the contrary, that z_{im} were changing over time, and suppose, to take the most extreme example, that ϕ_{ijmt} were *not* changing over time. This would be isomorphic to z_{im} staying the same over time while ϕ_{ijmt} changes, with ϕ_{ijmt} changing at the same rate for each occupation within each industry.

We will compute each of these derivatives in turn. In order to do so, first note that, under the simplifying assumptions made in Section **??**[29], the equilibrium value (in quantity terms) of country m's total production of industry i at date t is equal to

$$y_{imt} = (\frac{1}{I})\frac{(w_{mt})^{-\theta-1}(z_{imt})^{\theta+1}}{\sum_{m'=1}^{N}(w_{m't})^{-\theta}(z_{im't})^{\theta}} \tag{56}$$

Taking a first-order log-linear approximation of y_{imt}:

$$\ln(y_{imt}) \approx \ln(\frac{1}{I})+(-\theta-1)\ln(w_{mt})+(\theta+1)\ln(z_{imt})-\ln(N)+\sum_{m'=1}^{N}(\frac{1}{N})[\theta\ln(w_{m't})-\theta\ln(z_{im't})] \tag{57}$$

And a first-order log-linear approximation of country m's equilibrium wage[30] is

$$\ln(w_{mt}) \approx (\frac{1}{1+\theta})\ln(I) + (\frac{1}{1+\theta})\sum_{i=1}^{I}(\frac{1}{I})\theta\ln(z_{imt}) \tag{58}$$

We are now ready to compute the derivatives in equation (55). From the definition of comparative advantage (equation (41)), we have

$$\frac{d\ln(z_{imt})}{d\ln(CA_{imt})} = \frac{d\ln(CA_{i,m,t+1})}{d\ln(z_{i,m,t+1})} = 1 \tag{59}$$

From the learning-by-doing equation (equation (40)) we have

$$\frac{d\ln(z_{i,m,t+1})}{d\ln(z_{imt})} = 1 \tag{60}$$

and

$$\frac{d\ln(z_{i,m,t+1})}{d\ln(y_{imt})} = \rho \tag{61}$$

From equation (57) we have

$$\frac{d\ln(y_{imt})}{d\ln(z_{imt})} \approx (1+\theta)(1-\frac{d\ln(w_{mt})}{d\ln(z_{imt})}) \tag{62}$$

From equation (58) we have

[29] Recall that these simplifying assumptions are that trade costs are zero, the size of the labor force in each country is equal to 1, and each country's representative household's utility function puts equal weight on each industry.

[30] The derivation of this approximation is available upon request. This expression for the equilibrium wage holds exactly in the case of complete symmetry across countries and industries (and, as mentioned above, zero trade costs), but only holds as a first-order approximation outside of it. In the case of only one industry, this reduces to the closed-form solution to the Eaton and Kortum (2002) model with zero trade costs provided by Alvarez and Lucas (2007).

$$\frac{d\ln(w_{mt})}{d\ln(z_{imt})} \approx (\frac{1}{I})(\frac{\theta}{1+\theta}) \tag{63}$$

Combining equations (59), (60), (61), (62), and (63) gives us

$$\frac{d\ln(CA_{i,m,t+1})}{d\ln(CA_{imt})} \approx 1 + \rho(1 + [\frac{I-1}{I}]\theta) \tag{64}$$

This proves the proposition. ■

B Tables and figures

Table 6: List of countries in sample

Morocco	Malaysia
Tunisia	Pakistan
Nigeria	Philippines
Canada	Singapore
USA	Thailand
Argentina	Taiwan
Brazil	Belgium and Luxembourg
Chile	Denmark
Colombia	France and Monaco
Ecuador	Greece
Mexico	Ireland
Peru	Italy
Venezuela	Netherlands
Israel	Portugal
Japan	Spain
Turkey	UK
China, Hong Kong, and S.A.R.'s	Norway
South Korea	Sweden
	Switzerland and Lichtenstein

Table 7: List of industries in sample

NAICS	Industry	NAICS	Industry
311	Food Manufacturing	325	Chemical Manufacturing
312	Beverages and Tobacco Products	326	Plastics and Rubber Products
313	Textile Mills	327	Nonmetallic Mineral Products
314	Textile Product Mills	331	Primary Metal Manufacturing
315	Apparel Manufacturing	332	Fabricated Metal Products
316	Leather and Allied Products	333	Machinery
321	Wood Products	334	Computers and Electronic Products
322	Paper Manufacturing	335	Electrical Equipment and Appliances
323	Printing and Related Activities	336	Transportation Equipment
324	Petroleum and Coal Products	339	Miscellaneous Manufacturing

Table 8: List of occupations in sample

SOC code	Occupation
11	Management Occupations
13	Business and Financial Operations Occupations
15	Computer and Mathematical Occupations
17	Architecture and Engineering Occupations
19	Life, Physical, and Social Science Occupations
21	Community and Social Services Occupations
23	Legal Occupations
25	Education, Training, and Library Occupations
27	Arts, Design, Entertainment, Sports, and Media Occupations
29	Healthcare Practitioners and Technical Occupations
31	Healthcare Support Occupations
33	Protective Service Occupations
35	Food Preparation and Serving Related Occupations
37	Building and Grounds Cleaning and Maintenance Occupations
39	Personal Care and Service Occupations
41	Sales and Related Occupations
43	Office and Administrative Support Occupations
45	Farming, Fishing, and Forestry Occupations
47	Construction and Extraction Occupations
49	Installation, Maintenance, and Repair Occupations
51	Production Occupations
53	Transportation and Material Moving Occupations

Histogram of $\hat{\rho}$ across subsamples of countries and industries

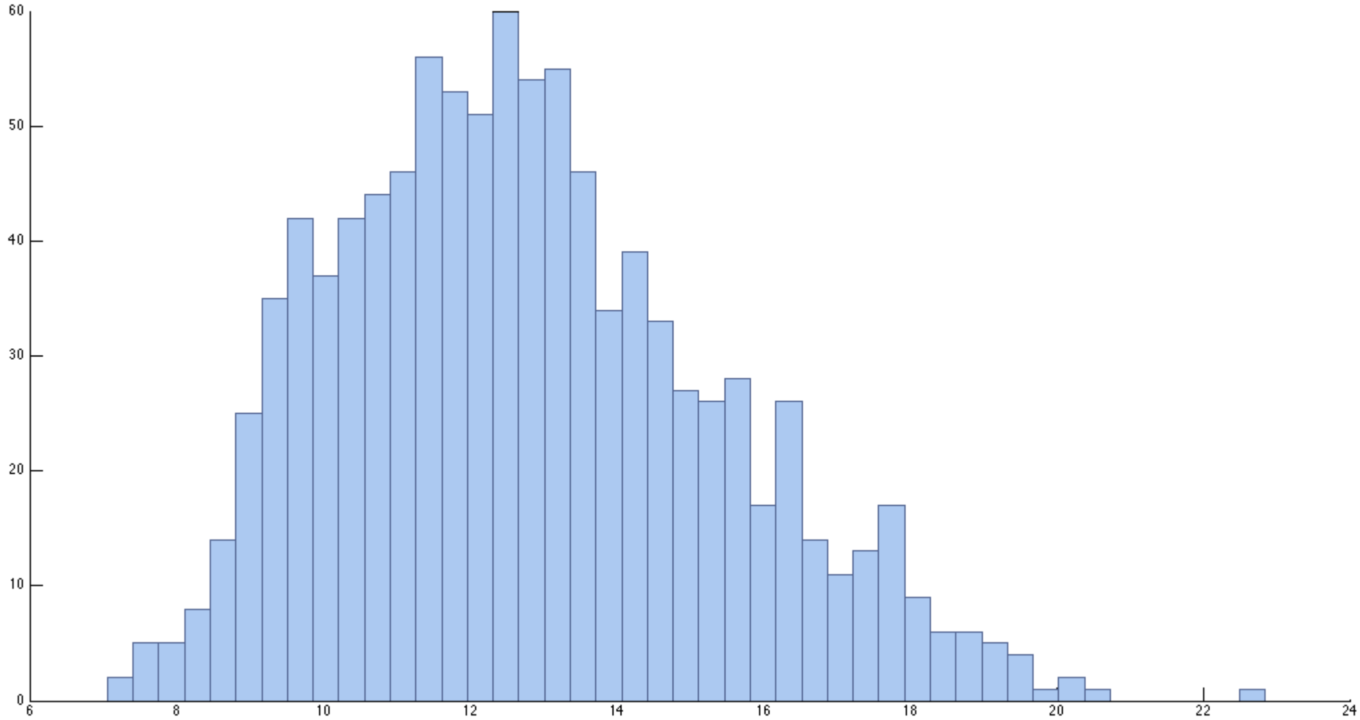

Figure 5: This graph plots the distribution of the $\hat{\rho}$ that solves minimization problem (46), across 1000 random subsamples of 22 countries and 10 industries, for the time period 1960's to 1970's.

Histogram of $\hat{\sigma}$ across subsamples of countries and industries

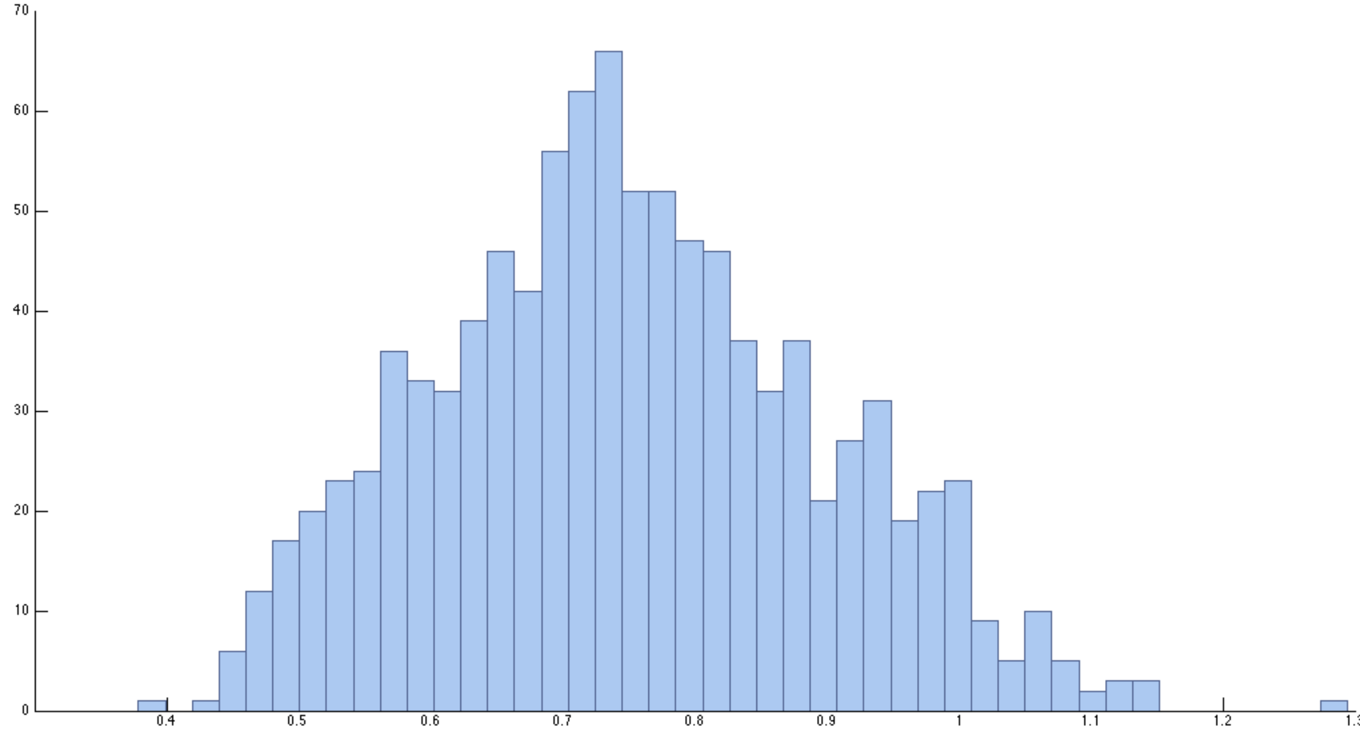

Figure 6: This graph plots the distribution of the $\hat{\sigma}$ that solves minimization problem (46), across 1000 random subsamples of 22 countries and 10 industries, for the time period 1960's to 1970's.

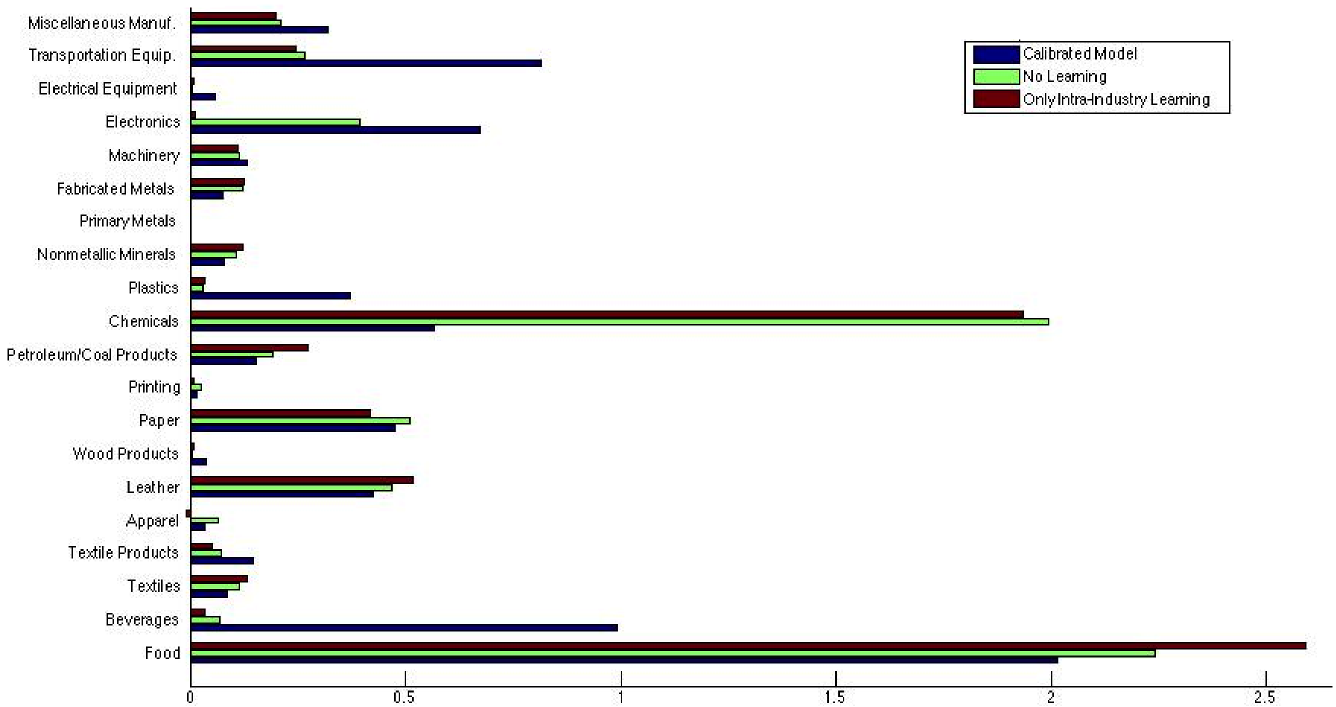

Figure 7: Percent increase in China's GDP, given 10% increase in industry-i productivity.

Table 9: Bonacich centrality of each industry in network of occupational learning diffusion

Industry	Bonacich Centrality
Electrical Equipment	2.03
Transportation Equipment	2.03
Primary Metals	2.12
Textiles	2.22
Misc. Manufacturing	2.02
Paper	2.07
Machinery	2.04
Plastics	2.16
Apparel	2.25
Electronics	1.63
Fabricated Metals	2.17
Printing	2.04
Minerals	1.86
Food	1.87
Chemicals	1.83
Petroleum/Coal Products	1.85
Wood Products	2.03