

On the Geography of Global Value Chains

Pol Antràs and Alonso de Gortari*
Harvard University and NBER

PRELIMINARY AND INCOMPLETE

April 28, 2017

Abstract

This paper develops a multi-stage general-equilibrium model of global value chains (GVCs) and studies the specialization of countries within GVCs in a world with barriers to international trade. With costly trade, the optimal location of production of a given stage in a GVC is not only a function of the marginal cost at which that stage can be produced in a given country, but is also shaped by the proximity of that location to the precedent and the subsequent desired locations of production. We show that, other things equal, it is optimal to locate relatively downstream stages of production in relatively central locations. We also develop a tractable, quantifiable version of our model that illustrates how changes in trade costs affect the extent to which various countries participate in domestic, regional or global value chains, and traces the real income consequences of these changes.

*We thank Arnaud Costinot and Iván Werning for useful conversations during the most upstream stages of this project. Roc Armenter, Rob Johnson, Myrto Kalouptsi, Sam Kortum, Marc Melitz, Eduardo Morales, Felix Tintelnot, and Kei-Mu Yi provided very useful feedback on preliminary versions of the paper. We are also grateful to seminar audiences at Princeton, the North American Econometric Society Summer Meeting in Philadelphia, the SED in Toulouse, the Federal Reserve Bank of Dallas, Geneva, Chicago Booth, MIT, Stanford, Harvard, Brown, Clark, UCLA, and Wisconsin for useful feedback. Le Kang, Daniel Ramos, BooKang Seol, Maria Voronina, and Diana Zhu provided valuable research assistance at various stages of this paper's production chain. Antràs acknowledges support from the NSF (proposal #1628852). All errors are our own.

1 Introduction

In recent decades, technological progress and falling trade barriers have allowed firms to slice up their value chains, retaining within their domestic economies only a subset of the stages in these value chains. The rise of global value chains (GVCs) has dramatically changed the landscape of the international organization of production, placing the specialization of countries *within* GVCs at the center stage. Where in GVCs are different countries specializing? Should countries use specific policies to place themselves in particularly appealing segments of GVCs? These are questions being posed in the policy arena for which the academic literature has yet to provide satisfactory answers.

This paper studies the specialization of countries within GVCs in a world with barriers to international trade. Although we are motivated by normative questions, the focus of this paper is on outlining the implications of the existence of *exogenously* given trade costs for the equilibrium shape of GVCs. The role of trade barriers on the geography of GVCs is interesting in its own right and has been relatively underexplored in the literature, perhaps due to the technical difficulties that such an analysis entails. More specifically, characterizing the allocation of production stages to countries is not straightforward because the optimal location of production of a given stage in a GVC is not only a function of the marginal cost at which that stage can be produced in a given country, but is also shaped by the proximity of that location to the precedent and the subsequent desired locations of production.

We start off our analysis in section 2 by illustrating these interdependencies in a simple partial equilibrium environment. We consider the problem of a *lead firm* choosing the location of its various production stages in an environment with costly trade. A key insight from our partial-equilibrium framework is that the relevance of geography (or trade costs) in shaping the location of the various stages of a GVC is more and more pronounced as one moves towards more and more downstream stages of a value chain. Intuitively, whenever trade costs are largely proportional to the gross value of the good being transported, these costs compound along the value chain, thus implying that trade costs erode more value added in downstream relative to upstream stages. In a parameterized example of our framework, this differential effect of trade costs takes the simple form of a stage-specific ‘trade elasticity’ that is increasing in the position of a stage in the value chain. The fact that trade costs are proportional to gross value follows from our iceberg formulation of these costs, a formulation that is not only theoretically appealing, but is also a reasonable approximation to reality.¹

Having characterized the key properties of the solution to the lead-firm problem, we next show how it can be ‘decentralized’. More specifically, we consider an environment in which there is no lead firm coordinating the chain, and instead stand-alone producers of the various stages in a GVC make cost-minimizing sourcing decisions by purchasing the good completed up to the prior

¹The fact that import duties and insurance costs are approximately proportional to the value of the goods being shipped should be largely uncontroversial. For shipping costs, weight and volume are naturally also relevant, but as shown by Brancaccio, Kalouptsi and Papageorgiou (2017), search frictions in the shipping industry allow shipping companies to extract rents from exporters by charging shipping fees that are typically increasing in the value of the goods in transit.

stage from their least-cost source. The partial equilibrium of this decentralized economy coincides with the solution to the lead-firm problem – and in fact can be recast as a dynamic programming formulation of the lead-firm problem – but it is dramatically simpler to compute. For a chain entailing N stages with each of these stages potentially being performed in one of J countries, characterizing the J optimal GVCs that service consumers in each country requires only $J \times N \times J$ computations, instead of the lead firm having to optimize over J^N potential paths for each of the J locations of consumption (for a total of $J \times J^N$ computations).

Although the results of our partial equilibrium model suggest that more central countries should have comparative advantage in relatively downstream stages within GVCs, formally demonstrating such a result requires developing a general-equilibrium model of GVCs in which production costs are endogenously determined and also shaped by trade barriers. With that goal in mind and also to explore the real income implications of changes in trade costs, in section 3 we develop a Ricardian model of trade in which the combination of labor productivity and trade costs differences across countries shapes the equilibrium position of countries in GVCs. More specifically, we adapt the Eaton and Kortum’s (2002) Ricardian model to a multi-stage production environment and derive sharp predictions for the *average* participation of countries in different segments of GVCs.

Previous attempts to extend the Ricardian model of trade to a multi-stage, multi-country environment (e.g., Yi, 2003, 2010, Johnson and Moxnes, 2016, Fally and Hillberry, 2016) have focused on the quantification of relatively low-dimensional models with two stages or two countries. Indeed, as we describe in section 3, it is not obvious how to exploit the extreme-value distribution results invoked by Eaton and Kortum (2002) in a multi-stage environment in which cost-minimizing location decisions are a function of the various cost ‘draws’ obtained by producers worldwide at various stages in the value chain. The reason for this is that neither the sum nor the product of Fréchet random variables are themselves distributed Fréchet, and thus previous approaches have been forced to resort to numerical analyses and simulated method of moments estimation.

We propose two alternative approaches to restore the tractability of Eaton and Kortum (2002) in a Ricardian model with multi-stage production. The first approach consists in simply treating the *overall* unit cost of production of a GVC flowing through a sequence of countries as a draw from a Fréchet random variable with a location parameter that is a function of the states of technology and wage levels of *all* countries involved in that GVC, as well as of the trade costs incurred in that chain. The second approach maintains the standard assumption that labor productivity is stage-specific and drawn from a Fréchet distribution, but instead considers a decentralized equilibrium in which, producers of a particular stage in a GVC have incomplete information about the productivity of certain suppliers upstream from them. More specifically, we assume that firms know their productivity and that of the suppliers immediately upstream from them (i.e., their tier-one suppliers) when they commit to sourcing from a particular supplier, but they do not know the precise productivity of their suppliers’ suppliers (i.e., tier-two suppliers, tier-three suppliers, and so on). Interestingly, we find that these two alternative approaches are isomorphic, in the sense that

they yield the exact same equilibrium equations.²

Under these two alternative assumptions, we show in section 4 that our model generates a closed-form expression for the probability of any potential path of production constituting the cost-minimizing path to service consumers in any country. These probabilities are analogous to the trade shares in Eaton and Kortum (2002), and indeed our model nests their framework in the absence of multi-stage production. Exploiting properties of the resulting distribution of final-good and input prices, we show that our model also delivers closed-form expressions for final-good and input trade flows across countries, which can easily be mapped to the various entries of a world Input-Output table, or WIOT for short. Various versions of these type of global Input-Output tables have become available in recent years, including the World Input Output Database, the OECD’s TiVA statistics, and the Eora MRIO database. Our Ricardian multi-stage framework also delivers a simple formula relating real income to the relative prevalence of purely domestic value chains, a formula that generalizes the ‘gains from trade’ formula in Arkolakis et al. (2012). Although the set of general-equilibrium equations is a bit more cumbersome than in Eaton and Kortum (2002), we show how the proof of existence and uniqueness in Alvarez and Lucas (2007) can be easily (though tediously) adapted to our setting. Finally, we formally establish the existence of a centrality-downstreamness nexus, by which the average downstreamness of a country in GVCs should be increasing in this country’s centrality (holding other determinants of comparative advantage constant). After introducing our main data sources, in section 5, we provide suggestive empirical evidence for this centrality-downstreamness nexus and for a key mechanism of the model – namely, the fact that the trade elasticity is larger for downstream stages than for upstream stages.

In section 6, we leverage the tractability of our framework to back out the model’s fundamental parameters from data on the various entries of a WIOT. Our empirical approach constitutes a blend of calibration and estimation. First, we show that when abstracting from variation in domestic costs across countries, our equilibrium conditions unveil a simple way to back out the matrix of bilateral trade costs across countries from data on final-good trade flows within and across countries. Our approach is akin to that in Head and Ries (2001), but it requires the use of only final-good trade flows. We also fix the a key parameter that governs the shape of the Fréchet distributions of productivity to (roughly) match the aggregate trade elasticity implied by our model. Conditional on a set of countries J and a number of stages N , we then estimate the remaining parameters of the model via a generalized method of moments (GMM), in which we target the diagonal entries of a WIOT.

We perform this exercise for two distinct and complementary samples. First, we use 2014 data from the World Input-Output Database, a source which is deemed to provide high-quality reliable data on intermediate input and final-good bilateral trade flows across countries for a sample of 43 countries and the rest of the world. The main downside of this database is that the bulk of the countries in the database are high- and medium-income countries in Europe, Asia and North

²The approach of building some form of incomplete information (or ex-ante uncertainty) into the Eaton and Kortum (2002) framework is similar in spirit to the one pursued by Tintelnot (forthcoming) and Antràs, Fort and Tintelnot (2017).

America. In order to study the geography of GVCs worldwide, we also present results using the broader sample of 190 countries in the Eora MRIO database. This data source is admittedly less reliable, but it allows us to estimate the model for 101 countries (or consolidated countries) in which all continents and income-levels are more properly represented. In both cases, we find that the model is able to match the targeted moments remarkably well, and it also provides a very good fit for the cells of the WIOT that were not directly targeted in the estimation.

Armed with estimates of the fundamental parameters of the model, we conclude the paper in section 7 by performing counterfactual exercises that illustrate how changes in trade barriers affect the extent to which various countries participate in domestic, regional or global value chains, and traces the real income consequences of these changes. We find that the gains from trade (i.e., the income losses from reverting to autarky) emanating from our model are modestly larger than those obtained from a version of our model without multi-stage production. This variant of our model is a generalization of the Eaton and Kortum (2002) model calibrated to match *exactly* the WIOT. When studying trade costs reductions relative to their calibrated levels, we find much higher income gains, both in absolute terms, but also relative the version of our model without multi-stage production. This larger gains partly reflect the increased participation of low-income countries in GVCs.

Our paper most closely relates to the burgeoning literature on GVCs. On the theoretical front, in recent years a few theoretical frameworks have been developed highlighting the role of the sequentiality of production for the global sourcing decisions of firms. Among others, this literature includes the work of Harms, Lorz, and Urban (2012), Baldwin and Venables (2013), Costinot *et al.* (2013), Antràs and Chor (2013), Kikuchi *et al.* (2014), Fally and Hillberry (2014), and Tyazhelnikov (2016).³ A key limitation of this body of theoretical work is that it either completely abstracts from modeling trade costs or it introduces such barriers in highly stylized ways (i.e., assuming common trade costs across all country-pairs). On the empirical front, a growing body of work, starting with the seminal work of Johnson and Noguera (2012), has been concerned with tracing the value-added content of trade flows and using those flows to better document the rise of GVCs and the participation of various countries in this phenomenon (see Koopman *et al.*, 2014, Johnson, 2014, Timmer *et al.*, 2014, de Gortari, 2017).⁴ A parallel empirical literature has developed indices of the relative positioning of industries and countries in GVCs (see Fally, 2012, Antràs *et al.*, 2012, Alfaro *et al.*, 2015). On the quantitative side, and as mentioned above, our work builds on and expands on previous work by Yi (2003, 2010), Johnson and Moxnes (2016) and Fally and Hillberry (2016). Other authors, and most notably Caliendo and Parro (2015), have developed quantitative frameworks with Input-Output linkages across countries, but in models with a roundabout production structure without an explicit sequentiality of production. Finally, some implications of the rise of offshoring and GVCs for trade policy have been studied by Antràs and

³This literature is in turn inspired by earlier contributions to modeling multi-stage production, such as Dixit and Grossman (1982), Sanyal and Jones (1982), Kremer (1993), Yi (2003) and Kohler (2004).

⁴An important precursor to this literature is Hummels *et al.* (2001), who combined international trade and Input-Output data to construct indices of vertical specialization.

Staiger (2012) and Bown et al. (2016), but in much more stylized frameworks than studied in this paper.

The rest of the paper is structured as follows. Section 2 develops our partial equilibrium model and highlights some of its key features. Section 3 describes the assumptions of the general equilibrium model, and section 4 characterizes its equilibrium. Section 6 covers the estimation of our model and section 7 explores several counterfactuals. All proofs and several details on data sources and the estimation are relegated to the Appendix A and Online Appendix B.

2 Partial Equilibrium: Interdependencies and Compounding

In this section, we develop a simple model of firm behavior that formalizes the problem faced by a firm choosing the location of its various production stages in an environment with costly trade. For the time being, we consider the problem of a firm (or, more precisely, of a competitive fringe of firms) producing a particular good. We defer a discussion of the general equilibrium aspects of the model to section 3.

2.1 Environment

There are J countries in which consumers derive utility from consuming a final good. The good is produced combining N stages that need to be performed sequentially. The last stage of production can be interpreted as assembly and is indexed by N . We will often denote the set of countries $\{1, \dots, J\}$ by \mathcal{J} and the set of production stages $\{1, \dots, N\}$ by \mathcal{N} . At each stage $n > 1$, production combines a local composite input (which encompasses primitive factors of production and a bundle of materials), with the good finished up to the previous stage $n - 1$. Production in the initial stage $n = 1$ only uses the composite input. The cost of the composite input varies across countries and is denoted by c_i in country i . Countries also differ in their geography, as captured by a $J \times J$ matrix of iceberg trade coefficients $\tau_{ij} \geq 1$, where τ_{ij} denotes the units of the finished or unfinished good that need to be shipped from i for one unit to reach j . Firms are perfectly competitive and the optimal location $\ell(n) \in \mathcal{J}$ of the different stages $n \in \mathcal{N}$ of the value chain is dictated by cost minimization. Because of marginal-cost pricing, we will somewhat abuse notation and denote by $p_{\ell(n)}^n$ the unit cost of production of a good completed up to stage n in country $\ell(n)$. That good is available in country $\ell(n+1)$ at a cost $p_{\ell(n)}^n \tau_{\ell(n)\ell(n+1)}$.

We summarize technology via the following sequential cost function associated with a path of production $\ell = \{\ell(1), \ell(2), \dots, \ell(N)\}$:

$$p_{\ell(n)}^n(\ell) = g_{\ell(n)}^n \left(c_{\ell(n)}, p_{\ell(n-1)}^{n-1}(\ell) \tau_{\ell(n-1)\ell(n)} \right), \text{ for all } n \in \mathcal{N}. \quad (1)$$

The stage- and country-specific cost functions $g_{\ell(n)}^n$ in equation (1) are assumed to feature constant-returns-to-scale and diminishing marginal products. As mentioned before, we let the cost of the first stage depend only on the local composite input bundle, so constant returns to scale implies

$p_{\ell(1)}^1(\boldsymbol{\ell}) = g_{\ell(1)}^1(c_{\ell(1)})$ for all paths $\boldsymbol{\ell}$, with the function $g_{\ell(1)}^1$ necessarily being linear in $c_{\ell(1)}$.

Note that equation (1) also applies to the assembly stage N , and a good assembled in $\ell(N)$ after following the path $\boldsymbol{\ell}$ is available in any country j at a cost $p_j^F(\boldsymbol{\ell}) = p_{\ell(N)}^N(\boldsymbol{\ell}) \tau_{\ell(N)j}$ (we use the superscript F to denote finished goods). For each country $j \in \mathcal{J}$, the goal is then choose the optimal path of production $\boldsymbol{\ell}^j = \{\ell^j(1), \ell^j(2), \dots, \ell^j(N)\} \in \mathcal{J}^N$ that minimizes the cost $p_j^F(\boldsymbol{\ell})$ of providing the good to consumers in that country j .

For the bulk of this paper, we will focus on the case in which cross-country differences in technology are associated with Ricardian differences in the efficiency with which the composite bundle of inputs is used in different stages, and in which the function $g_{\ell(n)}^n$ is a Cobb-Douglas aggregator of the input bundle and the product finished up to the previous stage. More specifically, we write

$$p_{\ell(n)}^n(\boldsymbol{\ell}) = \left(a_{\ell(n)}^n c_{\ell(n)} \right)^{\alpha_n} \left(p_{\ell(n-1)}^{n-1}(\boldsymbol{\ell}) \tau_{\ell(n-1)\ell(n)} \right)^{1-\alpha_n}, \text{ for all } n \in \mathcal{N}, \quad (2)$$

where α_n denotes the cost share of the input bundle at stage n and $a_{\ell(n)}^n$ is the unit input bundle requirement at stage n in country $\ell(n)$. Because the initial stage of production uses solely the local composite input, we have $\alpha_1 = 1$.

2.2 Lead-Firm Problem

We consider first the problem of a *lead firm* choosing the location of production of all stages $n \in \mathcal{N}$, in order to minimize the overall cost of serving consumers in a given country j . Using $p_j^F(\boldsymbol{\ell}) = p_{\ell(N)}^N(\boldsymbol{\ell}) \tau_{\ell(N)j}$ and iterating (2), this problem reduces to:

$$\boldsymbol{\ell}^j = \arg \min_{\boldsymbol{\ell} \in \mathcal{J}^N} p_j^F(\boldsymbol{\ell}) = \arg \min_{\boldsymbol{\ell} \in \mathcal{J}^N} \left\{ \prod_{n=1}^N \left(a_{\ell(n)}^n c_{\ell(n)} \right)^{\alpha_n \beta_n} \times \prod_{n=1}^{N-1} \left(\tau_{\ell(n)\ell(n+1)} \right)^{\beta_n} \times \tau_{\ell(N)j} \right\} \quad (3)$$

where

$$\beta_n \equiv \prod_{m=n+1}^N (1 - \alpha_m), \quad (4)$$

and where we use the convention $\prod_{m=N+1}^N (1 - \alpha_m) = 1$. Note that $\sum_{n=1}^N \alpha_n \beta_n = 1$.

We next highlight two important features of program (3). First, notice that when trade costs are identical for all country-pairs (i.e., $\tau_{ij} = \tau$ for all i and j), the first two terms reduce to a constant that is independent of the path of production. In such a case, we can break the cost-minimization problem in (3) into a sequence of N independent cost-minimization problems in which the optimal location of stage n is simply given by $\ell^j(n) = \arg \min_i \{a_i^n c_i\}$, and is thus independent of the country of consumption j . Notice, however that this result requires no differences between internal and external trade costs (i.e., $\tau_{ij} = \tau$ also for $i = j$), and thus this case is isomorphic, up to a productivity shifter, to an environment with costless trade. With a general geography of trade costs, a lead firm can no longer perform cost minimization independently stage-by-stage, and instead it needs to optimize over the whole path of production. Intuitively, the location $\ell(n)$ minimizing production costs $a_{\ell(n)}^n c_{\ell(n)}$ might not be part of a firm's optimal path if the optimal locations for

stages $n - 1$ and $n + 1$ are sufficiently far from $\ell(n)$. A direct implication of this result is that the presence of arbitrary trade costs turns a problem of dimensionality $N \times J$ into J much more complex problems of dimensionality J^N each. As we will see below, however, the dimensionality of program (3) can be dramatically reduced using dynamic programming.

A second noteworthy aspect of the minimand in equation (3) is that the trade-cost elasticity of the unit cost of serving consumers in country j increases along the value chain. More specifically, note from equation (4) that, as long as $\alpha_n > 0$ for all n , we have $\beta_1 < \beta_2 < \dots < \beta_N = 1$. For the particular case in which overall value added is a symmetric Cobb-Douglas aggregator of the value added of all stages (i.e., $\alpha_n \beta_n = 1/N$, for all n), the program in (3) reduces to

$$\ell^j = \arg \min_{\ell \in \mathcal{J}^N} p_j^F(\ell) = \arg \min_{\ell \in \mathcal{J}^N} \left\{ \prod_{n=1}^N \left(a_{\ell(n)}^n c_{\ell(n)} \right)^{1/N} \times \prod_{n=1}^{N-1} \left(\tau_{\ell(n)\ell(n+1)} \right)^{n/N} \times \tau_{\ell(N)j} \right\}, \quad (5)$$

and the trade-cost elasticity increases linearly with the downstreamness n of a stage.

The reason for this compounding effect of trade costs stems from the fact that the costs of transporting goods have been modelled (realistically, as we argued in the Introduction) to be proportional to the gross value of the good being transacted, rather than being assumed proportional to the value added at that stage. Thus, as the value of the good rises along the value chain, so does the amount of resources used to transport the goods across locations. An implication of this compounding effect is that, in choosing their optimal path of production, firms will be relatively more concerned about reducing trade costs in relatively downstream stages than in relatively upstream stages. As we will illustrate below and formally demonstrate when exploring the general equilibrium of our model, this feature of the cost function will generate a centrality-downstreamness nexus by which, *ceteris paribus*, relatively more central countries will tend to gain comparative advantage and specialize in relatively downstream stages.⁵

Although we have derived this compounding effect of trade costs for the case of Ricardian technological differences and Cobb-Douglas cost functions, we show in Appendix A.1 that the same result applies for arbitrary constant-returns-to-scale technologies of the type in equation (1). More specifically, denoting by β_n the elasticity of $p_j^F(\ell)$ with respect to $\tau_{\ell(n)\ell(n+1)}$, we show that β_n is again necessarily non-decreasing in n even when these elasticities are not pinned down by exogenous parameters. Thus, the result that firms will be particularly concerned about minimizing trade costs in downstream stages is quite general.⁶

⁵Building on the results in Costinot (2012), we can briefly anticipate this result with the following example. Suppose that trade costs can be decomposed as $\tau_{ij} = (\rho_i \rho_j)^{-1}$, where we take ρ_i is an index of the *centrality* of country i . In such a case, it is straightforward to show that leaving aside other determinants of comparative advantage, the unit cost of servicing consumers in country j is log-supermodular in a country's centrality ρ_i and a stage's downstreamness n (see more on this in section 4).

⁶For example, for the case of a symmetric Leontief technology and production costs equal to 1 in all countries and stages (i.e., $a_{\ell(n)}^n c_{\ell(n)} = 1$ for all n and $\ell(n)$), we obtain

$$p_j^F(\ell) = \tau_{\ell(N)j} + \tau_{\ell(N)j} \tau_{\ell(N-1)\ell(N)} + \tau_{\ell(N)j} \tau_{\ell(N-1)\ell(N)} \tau_{\ell(N-2)\ell(N-1)} + \tau_{\ell(N)j} \tau_{\ell(N-1)\ell(N)} \tau_{\ell(N-2)\ell(N-1)} \tau_{\ell(N-3)\ell(N-2)} + \dots,$$

which again illustrates the larger relative importance of downstream trade costs.

2.3 Decentralization and Dynamic Programming

We have so far characterized the problem of a *lead firm* with full information on the productivity of the various potential producers of each stage n in each country j . This characterization relies on strong informational assumptions, so we now consider an alternative environment in which no individual firm coordinates the whole value chain. Instead, we assume that a value chain consists of a series of stage-specific producers that simply minimize their cost of production taking into account their input bundle cost, their productivity, and the cost at which they can obtain the good finished up to the immediately preceding stage. Similarly, consumers in country j simply purchase the final good from whichever assembler (i.e., stage N producer) worldwide can provide the finished good at the lowest price.

From equation (1), a producer of stage n in country $\ell(n)$ would choose to procure the good finished up to stage $n - 1$ by simply solving $\min_{\ell(n-1) \in \mathcal{J}} \left\{ p_{\ell(n-1)}^{n-1} \tau_{\ell(n-1)\ell(n)} \right\}$, regardless of this producer's input bundle cost, productivity and the future path of the good after flowing through $\ell(n)$ at stage n . Furthermore, the resulting price at which this producer can sell the good finished up to stage n to producers of stage $n + 1$ is only a function of $a_{\ell(n)}^n c_{\ell(n)}$ and this minimum price $\min_{\ell(n-1) \in \mathcal{J}} \left\{ p_{\ell(n-1)}^{n-1} \tau_{\ell(n-1)\ell(n)} \right\}$. Producers of the initial stage $n = 1$ only use their local composite bundle, and thus $p_{\ell(1)}^1 = a_{\ell(1)}^1 c_{\ell(1)}$.

With constant returns to scale, the identity of the specific firms making these decisions is of course immaterial, so this formulation is entirely consistent with our previous *lead firm* using dynamic programming to solve for the optimal path of production leading to consumption in each country $j \in \mathcal{J}$. More specifically, instead of solving program (3) in a brute force manner, the lead firm can break the problem into a series of stage- and country-specific optimal sourcing problems (as in the decentralized formulation above), and then solve the problem via forward induction (starting in the most upstream stages). Invoking the principle of optimality, we can then establish (see Appendix A.2) that the resulting optimal path of production $\ell = \{\ell(1), \ell(2), \dots, \ell(N)\} \in \mathcal{J}^N$ that minimizes the cost $p_j^F(\ell)$ in this decentralized formulation of the problem will coincide with the one we obtained solving the *lead-firm* problem in (3) by exhaustive search. A key advantage of this dynamic programming approach is that it only requires $J \times N \times J$ computations to obtain the optimal production path for *all* destinations of final consumption, instead of having to optimize over J^N potential paths for each country j .⁷ For example, with 200 countries and 5 stages, this amounts to only 200,000 computations rather than 64 trillion computations.⁸ Although it might be clear from our discussion above, it is worth stressing that the isomorphism between the lead-firm problem and the decentralized problem holds true for any constant-returns-to-scale technology, and not only for the Cobb-Douglas one in (2).

⁷See Appendix A.2 for more details. This same point has been made in contemporaneous work by Tyazhelnikov (2016).

⁸Though the dimensionality of the lead firm's problem is huge, for the particular case with Cobb-Douglas technologies, in Appendix A.2 we show that the problem can also be written as a zero-one integer programming problem, for which many extremely quick and efficient algorithms are available (see, for instance, <http://www.gurobi.com>).




Figure 1: An Example with Four Countries

2.4 An Example

We close this section by illustrating some of the salient and distinctive features of this partial model of sequential production via a simple example. We consider a world with four countries ($J = 4$) and four stages ($N = 4$). Technology is given by the symmetric Cobb-Douglas specification in (5), with $\alpha_n \beta_n = 1/4$ for all n . The four countries are divided into two regions, the West (comprising countries A and B) and the East (comprising countries C and D). The ‘geography’ of this example is illustrated in Figure 1. Note that we impose a great deal of symmetry: intra-regional trade costs are common in both regions, and inter-regional costs between A and C are identical to those between B and D . On the other hand, trade costs between B and C are lower than between A and C . For simplicity, all domestic trade costs are set to 0, so $\tau_{ii} = 1$ for $i = A, B, C, D$. We are interested in solving for the optimal path of a four-stage production process leading to consumption in country D (in green in the figure). Note that shipping to D directly is least costly when shipping from D itself, followed by C (the other country in the East), then by A and finally by B , which is the most remote country relative to D .

We compute the optimal path leading to D for different levels of trade costs starting with a benchmark with $\tau_{AB} = \tau_{CD} = 1.3$, $\tau_{BC} = 1.5$, $\tau_{AD} = 1.75$, $\tau_{AC} = \tau_{BD} = 1.8$, and then scale these international trade costs up or down by a shifter s (so starting from τ_{ij} , we instead use $\tilde{\tau}_{ij}(s) = 1 + s \times (\tau_{ij} - 1)$).⁹ For each matrix of trade costs, we run one million simulations with production costs $a_j^n c_j$ being drawn independently for each stage n and each country j from a lognormal distribution with mean 0 and variance 1. By choosing a common distribution across countries and stages, we seek to isolate the role of trade costs in shaping the optimal path of sequential value chains.

The results of these simulations are depicted in Figure 2 for various levels of s ranging from 0 (free trade) to 50 (which results in close to prohibitive trade costs). The upper left panel shows the

⁹These parameters are chosen such that for all values of s considered, the triangle inequality holds for any three given countries.

average propensity of each country to appear in GVCs leading to consumption in D . The upper right panel depicts the average position (or downstreamness) of countries in these GVCs. Finally, the lower panel decomposes GVCs into purely domestic ones (with all production stages in D), purely regional ones (with some stages in C and D , but not in A or B) and global ones (involving at least one stage in A or B).

Several aspects of Figure 2 are worth highlighting. First, focusing on the upper left panel, notice that country B , which is farthest away from country D , appears slightly more often in value chains leading to D than its Western neighbor A does. The reason for this surprising fact is tightly related to the sequential nature of production. Even though, A is closer to D than B is, B is relatively close to D 's Eastern neighbor C , and this makes this 'remote' country B a particularly appealing location from which to set off value chains that will flow to D through C .¹⁰ A second noteworthy aspect, apparent from the upper right panel of Figure 2, is that remoteness appears to shape the average position of a country in GVCs, a fact we anticipated above. More specifically, country B , which is farthest away from D , is on average the most upstream of all countries, followed by its Western neighbor A , and then by C , with D being naturally the country positioned most downstream in value chains leading to consumption in D . Finally, the lower panel of Figure 2 illustrates how the progressive reduction of international trade costs first gives rise to GVCs that are largely regional in nature, and then later to truly global value chains involving inter-regional trade. It is also worth highlighting that even for fairly low trade costs, purely domestic GVCs remain quite prevalent, much more so than would be predicted by an analogous model without sequentiality (see the Online Appendix B.1). The reason for this is the compounding effect of trade costs, which other things equal makes it costly to offshore intermediate stages in chains in which D has comparative advantage in the most upstream *and* downstream stages.

3 General Equilibrium Model

We next embed the model of firm behavior developed in section 2 into a full-fledged general equilibrium model.

3.1 Environment

We continue to assume a world with J countries (indexed by i or j) where consumers now derive utility from consuming a continuum of final-good varieties (indexed by z). Preferences are CES and given by

$$u\left(\{y_i^N(z)\}_{z=0}^1\right) = \left(\int_0^1 (y_i^N(z))^{\frac{\sigma-1}{\sigma}} dz\right)^{\frac{\sigma}{\sigma-1}}, \quad \sigma > 1. \quad (6)$$

Production of each of the final-good varieties is as described in the previous section: production processes entails N sequential stages (indexed by n) and is characterized by the Ricardian, Cobb-

¹⁰As we show in Online Appendix B.1, in an analogous world without sequentiality, the above pattern would not hold and the relative prevalence of countries would be strictly monotonic in the level trade costs incurred when shipping to the assembly location.




Figure 2: Some Features of Optimal Production Paths

Douglas specification in (2).

We let countries differ in three key aspects: (i) their technological efficiency, as determined by the unit input bundle requirements $a_i^n(z)$, (ii) their geography, as captured by a $J \times J$ matrix of iceberg trade cost $\tau_{ij} \geq 1$, and (iii) their size, as reflected by the measure L_i of ‘equipped’ labor available for production in each country i (labor is inelastically supplied and commands a wage w_i).

The composite bundle of inputs used at each stage comprises labor and an aggregator of final-good varieties that corresponds exactly to the CES aggregator in (6). In other words, part of final-good production is not absorbed by consumers, but rather by firms that use those goods as a bundle of materials. This roundabout structure of production is standard in recent Ricardian models (see Eaton and Kortum, 2002, Alvarez and Lucas, 2007, or Caliendo and Parro, 2014), so we adopt it for comparability (see, in particular, section 7). We should stress, however, that our model features intermediate input flows across countries even in the absence of these production ‘loops’. We let the cost of an input bundle in country i be captured by a Cobb-Douglas aggregator $c_i = (w_i)^{\gamma_i} (P_i)^{1-\gamma_i}$, where P_i is the ideal price index associated with the CES aggregator in (6). Although allowing for variation in value added shares γ_i across countries is not important for our theoretical results, it will prove useful in allowing our model to better match world Input-Output tables.

This completes the discussion of the structure of our general-equilibrium model. In principle,

given values for the unit bundle requirements $a_i^n(z)$ and all other primitive parameters, the equilibrium of the model could be computed by (i) solving for the cost-minimizing path of production for each good z and each destination of consumption j given a vector of wages, and (ii) invoking labor-market clearing to reduce equilibrium wages to the solution of a fixed point problem. Such an approach, however, would not be particularly useful in order to formally characterize certain features of the equilibrium or to estimate the model in a particularly transparent manner. With that in mind, we next explore a particularly convenient parametrization of the unit bundle requirements $a_i^n(z)$.

3.2 Technology

Building on the seminal work of Eaton and Kortum (2002), we propose a probabilistic specification of the unit bundle requirements $a_i^n(z)$ that delivers a remarkably tractable multi-stage, multi-country Ricardian model. We are certainly not the first ones to explore such a multi-stage extension of the Eaton and Kortum (2002) framework. Yi (2010) and Johnson and Moxnes (2016), for instance, consider a ‘natural’ extension in which each productivity parameter $1/a_i^n(z)$ is assumed stochastic and drawn independently (across goods and stages) from a type II (or Fréchet) extreme-value probability distribution, as in Eaton and Kortum (2002).

A key limitation of their approach is that the minimum cost associated with a given GVC path is not characterized by a particularly tractable distribution. The reason for this is that, although the minimum of a series of Fréchet draws is itself distributed Fréchet, the product of Fréchet random variables is *not* distributed Fréchet.¹¹ As a result, these papers need to resort to either calibration (Yi, 2010) or a simulated method of moments (Johnson and Moxnes, 2016) to map the model to actual data, while restricting the analysis to two-stage chains. We instead develop two alternative approaches that will permit a sharp characterization of some of the features of the equilibrium for an arbitrary number of stages and that will be readily amenable to structural (maximum-likelihood) estimation using world input-output tables.

A. *Lead-Firm Approach*

We begin by revisiting the problem of a lead firm choosing the location of the various stages of production with full knowledge of the realized unit bundle requirements $a_i^n(z)$ for each stage in each country. The key innovation we propose, relative to Eaton and Kortum (2002), is to introduce randomness to the overall cost of production of a given value chain, rather than to the productivity of each stage independently. Intuitively, a given production path $\ell = \{\ell(1), \ell(2), \dots, \ell(N)\} \in \mathcal{J}^N$ will be associated with an *average* cost that is naturally a function of trade costs, input bundle costs and the state of technology of the various countries involved in the chain. Yet, compatibility problems, production delays, or simple mistakes can generate idiosyncratic noise around that average. More formally, and building on the cost function in (3), we assume that the

¹¹ Assuming a linear cost function (i.e., perfect complementarity) does not provide tractability either because the sum of Fréchet random variables is not distributed Fréchet either.

overall ‘productivity’ of a given chain ℓ is characterized by

$$\Pr \left(\prod_{n=1}^N \left(a_{\ell(n)}^n(z) \right)^{\alpha_n \beta_n} \geq a \right) = \exp \left\{ -a^\theta \prod_{n=1}^N (T_{\ell(n)})^{\alpha_n \beta_n} \right\}, \quad (7)$$

which amounts to assuming that $\prod_{n=1}^N \left(a_{\ell(n)}^n(z) \right)^{-\alpha_n \beta_n}$ is distributed Fréchet with a shape parameter given by θ , and a location parameter that is a function of the states of technology in all countries in the chain, as captured by $\prod_{n=1}^N (T_{\ell(n)})^{\alpha_n \beta_n}$. A direct implication of this assumption is that the unit cost associated with serving consumers in a given country j via a given chain ℓ is also distributed Fréchet. More precisely, denoting by $p_j^F(\ell)$ the price paid by consumers in j for a good produced following the path ℓ , we have

$$\Pr \left(p_j^F(\ell) \geq p \right) = \exp \left\{ -p^\theta \times \prod_{n=1}^N \left((c_{\ell(n)})^{-\theta} T_{\ell(n)} \right)^{\alpha_n \beta_n} \times \prod_{n=1}^{N-1} (\tau_{\ell(n)\ell(n+1)})^{-\theta \beta_n} \times (\tau_{\ell(N)j})^{-\theta} \right\}, \quad (8)$$

independently of the final good z under consideration. This result will be key for neatly characterizing the equilibrium, as we will show in the next section.

B. Decentralized Approach

We also develop an alternative approach closer in spirit to the stage-specific productivity randomness in Yi (2010) and Johnson and Moxnes (2016), which also achieves tractability and, in fact, delivers an identical set of equilibrium conditions to those we will derive under the specification in (7). On the technology side, we now assume that $1/a_i^n(z)$ is drawn independently (across goods and stages) from a Fréchet distribution satisfying

$$\Pr \left(a_j^n(z)^{\alpha_n \beta_n} \geq a \right) = \exp \left\{ -a^\theta (T_j)^{\alpha_n \beta_n} \right\}. \quad (9)$$

Note that this formulation imposes a common variance (as captured by the shape parameter θ) of the contribution of each stage n 's productivity – i.e., $a_j^n(z)^{-\alpha_n \beta_n}$ – to the overall productivity of a value chain. This assumption ensures that we do not mechanically introduce heterogeneity across stages in the trade-cost elasticity related to the importance of these stages in production.¹²

In order to make this alternative approach tractable, we relax the assumption that firms choose the optimal path of production with full knowledge of the productivity levels with which all stages of production in their chain could be produced in different countries. More specifically, we explore an environment akin to the decentralized equilibrium developed in section 2.3, in which stage-specific producers simply attempt to minimize the cost of production of their stage. Unlike in section 2.3, we assume, however, that these stage-specific producers do not observe realized upstream prices

¹²There is no economic reason to think that the trade-cost elasticity should vary with the contribution of a stage to value added, and such variation would obviously obfuscate our result showing that this elasticity rises along the value chain.

before making sourcing decisions, and can only forecast these prices based on information on the productivity levels of their potential direct (or tier-one) suppliers in various countries. These tier-one supplier productivity levels are not sufficient statistics for sourcing prices because upstream marginal costs also depend on the productivity of suppliers further upstream (i.e., tier-two suppliers, tier-three suppliers and so on). The idea behind this formulation is that firms need to pre-commit to purchase from particular suppliers based on information they gather from inspecting (e.g., through factory visits) all their potential immediate suppliers. Ex-post, a supplier's price might be higher or lower than expected because this supplier may face unexpectedly high or low sourcing costs itself.

Because this decentralized approach with incomplete information is a bit more cumbersome than the formulation in (7), we illustrate how it works for the simple case with only two stages, input production (stage 1) and assembly (stage 2). In Appendix A.3, we show how the approach naturally generalizes to the case $N > 2$. Input producers of a given final good z in a given country $\ell(1) \in \mathcal{J}$ observe their productivity $1/a_{\ell(1)}^1(z)$, and simply hire labor and buy materials to minimize $p_{\ell(1)}^1(z) = a_{\ell(1)}^1(z) c_{\ell(1)}$. Assemblers of good z in any country $\ell(2) \in \mathcal{J}$ observe their own productivity $1/a_{\ell(2)}^2(z)$, as well as that of all potential input producers worldwide, and solve

$$p_{\ell(2)}^2(z) = \min_{\ell(1) \in \mathcal{J}} \left\{ \left(a_{\ell(2)}^2(z) c_{\ell(2)} \right)^{\alpha_2} \left(a_{\ell(1)}^1(z) c_{\ell(1)} \tau_{\ell(1)\ell(2)} \right)^{1-\alpha_2} \right\}.$$

Independently of the values of $a_{\ell(2)}^2(z)$, $c_{\ell(2)}$, and α_2 , the solution of this problem simply entails procuring the input from the location $\ell^*(1)$ satisfying $\ell^*(1) = \arg \min \left\{ \left(a_{\ell(1)}^1(z) c_{\ell(1)} \tau_{\ell(1)\ell(2)} \right)^{1-\alpha_2} \right\}$. As is well-known, the Fréchet assumption in (9) will make characterizing this problem pretty straightforward. Consider finally the problem of retailers in each country j seeking to procure a final good z to local consumers at a minimum cost. These retailers observe the productivity $1/a_{\ell(2)}^2(z)$ of all assemblers worldwide, but *not* the productivity of input producers, and thus seek to solve

$$p_j^F(z) = \min_{\ell(2) \in \mathcal{J}} \left\{ \left(a_{\ell(2)}^2(z) c_{\ell(2)} \right)^{\alpha_2} \mathbb{E} \left[a_{\ell^*(1)}^1(z) c_{\ell^*(1)} \tau_{\ell^*(1)\ell(2)} \right]^{1-\alpha_2} \tau_{\ell(2)j} \right\}. \quad (10)$$

If retailers could observe the particular realizations of input producers, the expectation in (10) would be replaced by the realization of $a_{\ell(1)}^1(z) c_{\ell(1)} \tau_{\ell(1)\ell(2)}$ in all $\ell(1) \in \mathcal{J}$, and characterizing the optimal choice would be complicated because it would depend on the joint distribution of $a_{\ell(2)}^2(z)$ and $a_{\ell(1)}^1(z)$, which is not Fréchet under (9). As we will demonstrate in section 4, with our incomplete information assumption, the expectation in (10) does not depend on the particular realizations of upstream productivity draws, and this will allow us to apply the well-know properties of the univariate Fréchet distribution in (9) to characterize the problem of retailers.

4 Characterization of the Equilibrium

In this section, we characterize the general equilibrium of our model. We proceed in six steps. First, we leverage our extreme-value representation of GVC productivity to obtain closed-form expressions

for the relative prevalence (in value terms) of different GVCs in the world equilibrium. Second, we show how to manipulate these relative market shares of different GVCs to obtain expressions for bilateral intermediate input and final-good flows across countries, which can be mapped to observable data from world Input-Output tables. Third, we study the existence and uniqueness of the general equilibrium. Fourth, we obtain expressions for the gains from trade in our model and compare them to those in Eaton and Kortum (2002). Fifth, we formalize the link between downstreamness and centrality that we hinted at in section 2. In a final sixth subsection, we provide some suggestive evidence for this downstreamness-centrality nexus and also for a key mechanism in the model.

4.1 Relative Prevalence of Different GVCs and Equilibrium Prices

Let us begin with the lead-firm version of our model, in which the price paid by consumers in j for a good produced following the path $\ell \in \mathcal{J}^N$ is given by the Fréchet distribution in (8). In such a case, we can readily invoke a few of the results in Eaton and Kortum (2002) to characterize the equilibrium prices and the relative prevalence of different GVCs. First, it is straightforward to verify that the probability of a given GVC ℓ being the cost-minimizing production path for serving consumers in j is given by

$$\pi_{\ell j} = \frac{\prod_{n=1}^{N-1} \left((T_{\ell(n)})^{\alpha_n} ((c_{\ell(n)})^{\alpha_n} \tau_{\ell(n)\ell(n+1)})^{-\theta} \right)^{\beta_n} \times (T_{\ell(N)})^{\alpha_N} ((c_{\ell(N)})^{\alpha_N} \tau_{\ell(N)j})^{-\theta}}{\Theta_j}, \quad (11)$$

where

$$\Theta_j = \sum_{\ell \in \mathcal{J}^N} \prod_{n=1}^{N-1} \left((T_{\ell(n)})^{\alpha_n} ((c_{\ell(n)})^{\alpha_n} \tau_{\ell(n)\ell(n+1)})^{-\theta} \right)^{\beta_n} \times (T_{\ell(N)})^{\alpha_N} ((c_{\ell(N)})^{\alpha_N} \tau_{\ell(N)j})^{-\theta}, \quad (12)$$

and where remember that $c_i = (w_i)^{\gamma_i} (P_i)^{1-\gamma_i}$. With a unit measure of final goods, $\pi_{\ell j}$ also corresponds to the share of GVCs ending in j for which ℓ is the cost-minimizing production path.¹³

Second, and as in Eaton and Kortum (2002), the price $p_j^F(\ell)$ paid by consumers in j for a good z produced following the path ℓ satisfies

$$\Pr(p_j^F(\ell, z) \leq p) = 1 - \exp\{-\Theta_j p^\theta\} \quad (13)$$

regardless of the good z and the actual source (or, in our case, path of production) of the good. Because the distribution of final-good prices in j is independent of the path ℓ , it then follows that the probabilities in $\pi_{\ell j}$ also constitute the shares of country j 's income spent on final goods produced under all possible paths $\ell \in \mathcal{J}^N$.

As is clear from equation (11), GVCs that involve countries with higher states of technology T_i

¹³Note that when $N = 1$, we necessarily have $\alpha_N = 1$, and the formulas (11) and (12) collapse to the well-know trade share formulas in Eaton and Kortum (2002).

or lower input bundles costs c_i will tend to feature disproportionately in production paths leading to consumption in j . Furthermore, and consistently with our discussion in section 2, high trade costs penalize the participation of countries in GVCs, but such an effect is disproportionately large for downstream stages relative to upstream stages. This is captured by the fact that the ‘trade elasticity’ associated with stage n is given by $\theta\beta_n$, and β_n is increasing in n with $\beta_N = 1$.

Following the same steps as in Eaton and Kortum (2002), we can further solve for the exact ideal price index P_j in country j associated with (6)

$$P_j = \kappa (\Theta_j)^{-1/\theta}, \quad (14)$$

where $\kappa = [\Gamma(\frac{\theta+1-\sigma}{\theta})]^{1/(1-\sigma)}$ and Γ is the gamma function. For the price index to be well defined, we impose $\sigma - 1 < \theta$.

So far we have focused on the ‘randomness-in-the-chain’ formulation in (8). Consider now our alternative approach with stage-specific randomness captured by (9) and incomplete information. As in section 3, we will focus here on the case with two stages and leave the more general case to Appendix A.3. Take two countries $\ell(1)$ and $\ell(2)$ and consider the probability $\pi_{\ell j}$ of a GVC flowing through $\ell(1)$ and $\ell(2)$ before reaching consumers in j . This probability is simply the product of (i) the probability of $\ell(1)$ being the cost-minimizing location of input production conditional on assembly happening in $\ell(2)$, and (ii) the probability of $\ell(2)$ being the cost-minimizing location of assembly for GVC serving consumers in j . Denoting $\mathcal{E}_{\ell(2)} = \mathbb{E} \left[\tau_{\ell^*(1)\ell(2)} a_{\ell^*(1)}^1(z) c_{\ell^*(1)} \right]^{1-\alpha_2}$, and using the properties of the Fréchet distribution, it is easy to verify that we can write $\pi_{\ell j}$ as

$$\pi_{\ell j} = \underbrace{\frac{(T_{\ell(1)})^{1-\alpha_2} (c_{\ell(1)} \tau_{\ell(1)\ell(2)})^{-\theta(1-\alpha_2)}}{\sum_{k \in \mathcal{J}} (T_k)^{1-\alpha_2} (c_k \tau_{k\ell(2)})^{-\theta(1-\alpha_2)}}}_{\Pr(\ell(1)|\ell(2))} \times \underbrace{\frac{(T_{\ell(2)})^{\alpha_2} \left((c_{\ell(2)})^{-\alpha_2} \tau_{\ell(2)j} \right)^{-\theta} (\mathcal{E}_{\ell(2)})^{-\theta}}{\sum_{i \in \mathcal{J}} (T_i)^{\alpha_2} \left((c_i)^{\alpha_2} (\tau_{ij}) \right)^{-\theta} (\mathcal{E}_i)^{-\theta}}}_{\Pr(\ell(2))}. \quad (15)$$

A bit less trivially, but also exploiting well-known properties of the Fréchet distribution, it can be shown that

$$\mathcal{E}_{\ell(2)} = \mathbb{E} \left[\tau_{\ell^*(1)\ell(2)} a_{\ell^*(1)}^1(z) c_{\ell^*(1)} \right]^{1-\alpha_2} = \varsigma \left(\sum_{k \in \mathcal{J}} (T_k)^{1-\alpha_2} (c_k \tau_{k\ell(2)})^{-\theta(1-\alpha_2)} \right)^{-1/\theta},$$

for some constant $\varsigma > 0$, which allows us to reduce (15) to

$$\pi_{\ell j} = \frac{(T_{\ell(1)})^{1-\alpha_2} (c_{\ell(1)} \tau_{\ell(1)\ell(2)})^{-\theta(1-\alpha_2)} (T_{\ell(2)})^{\alpha_2} \left((c_{\ell(2)})^{-\alpha_2} \tau_{\ell(2)j} \right)^{-\theta}}{\sum_{k \in \mathcal{J}} \sum_{i \in \mathcal{J}} (T_k)^{1-\alpha_2} (c_k \tau_{ki})^{-\theta(1-\alpha_2)} (T_i)^{\alpha_2} \left((c_i)^{\alpha_2} (\tau_{ij}) \right)^{-\theta}}. \quad (16)$$

It should be clear that this expression is identical to (11) – plugging in (12) – for the special case $N = 2$. It is also straightforward to verify that the distribution of the price $p_j^F(z)$ paid by consumers

in j for any good z produced following the path $\ell = \{\ell(1), \ell(2)\}$ is independent of the actual path ℓ and is again characterized, as in equation (13), by $\Pr(p_j^F(\ell) \leq p) = 1 - \exp\{-\tilde{\Theta}_j p^\theta\}$, where $\tilde{\Theta}_j$ is the denominator in (16), and is the analog of Θ_j in (12) when $N = 2$.

In sum, this alternative specification of the stochastic nature of technology delivers the exact same distribution of GVCs and of consumer prices as the one in which the overall GVC unit cost is distributed Fréchet. As mentioned above and as demonstrated in Appendix A.3, this isomorphism carries over to the case $N > 2$.

4.2 Mapping the Model to Observables

So far, we have just described how to adapt the Eaton and Kortum (2002) probabilistic approach to apply to trade shares in terms of specific production paths (or GVCs) rather than in terms of trade volumes. Unfortunately, these ‘GVC trade shares’ are not observable in the data, so we next describe how to map the model to the type of information available in world Input-Output datasets. These sources of data provide information on (i) the share of final-good consumption in country j originating in assembly plants (producing stage N) in all other countries i , and (ii) the share of intermediate input purchases used by firms in j originating from producers in all other countries in j .

Consider first the implications of our model for final-good shares. Notice that final goods flow from a given source country i to a given destination country j only in GVCs serving consumers in country j in which country i is in position N . Define the set of GVCs flowing through i at position n by $\Lambda_i^n \in \mathcal{J}^{N-1}$. Formally, $\Lambda_i^n = \{\ell \in \mathcal{J}^N \mid \ell(n) = i\}$. The overall relative prevalence of all GVCs serving consumers in j in which country i is in assembly (position N) can then be expressed as

$$\pi_{ij}^F = \frac{\sum_{\ell \in \Lambda_i^N} \prod_{n=1}^{N-1} \left((T_{\ell(n)})^{\alpha_n} ((c_{\ell(n)})^{\alpha_n} \tau_{\ell(n)\ell(n+1)})^{-\theta} \right)^{\beta_n} \times (T_i)^{\alpha_N} ((c_i)^{\alpha_N} \tau_{ij})^{-\theta}}{\Theta_j}. \quad (17)$$

Because these flows occur at the same expected price for all goods regardless of the actual source country j , it follows that the shares π_{ij}^F also correspond to the final consumption shares reported in world Input-Output tables. Our model thus provides explicit formulas for these world Input-Output entries as a function of the parameters of our model and the endogenous input bundle cost c_i , which we can solve for in general equilibrium. Note also that final-good trade flows between any two countries i and j are then simply given by $\pi_{ij}^F \times w_j L_j$, since spending on final goods in country j must equal aggregate income, and labor is the only factor of production (when we estimate the model, we will incorporate trade imbalances).

Computing intermediate input flows between any two countries i and j is a bit more tedious, but equally straightforward. To begin, we need to distinguish between two types of intermediate input flows. First, at any stage of production, firms in country j purchase a bundle of intermediates at cost P_j from firms worldwide, and part of that spending originates in country i . Because the bundle of intermediate corresponds exactly to the consumption CES aggregator, the share of j ’s

input purchases originating in i is again given by π_{ij}^F in (17).¹⁴ Furthermore, note that any time the bundle of intermediates is used in production, spending on it in country j equals a multiple $(1 - \gamma_j) / \gamma_j$ of spending on labor. As a result, aggregate flows between i and j of this type of intermediates are given by $\pi_{ij}^F \times (1 - \gamma_j) / \gamma_j \times w_j L_j$.

In our multi-stage model, there is a second type of intermediate input flows across countries. In particular, firms from j also import a semi-finished product from i in sequential GVCs in which i immediately precedes j . To compute these flows, let us begin by denoting by $\Lambda_{k \rightarrow i}^n \in \mathcal{J}^{N-2}$ the set of GVCs that flow through k at position $n \leq N - 1$ and through i at position $n + 1$, or more formally, $\Lambda_{k \rightarrow i}^n = \{\ell \in \mathcal{J}^N \mid \ell(n) = k \text{ and } \ell(n + 1) = i\}$. The probability that this subset of GVCs emerges in equilibrium in GVCs serving consumers in j is given by

$$\Pr(\Lambda_{k \rightarrow i}^n, j) = \frac{\sum_{\ell \in \Lambda_{k \rightarrow i}^n} \prod_{n=1}^{N-1} \left((T_{\ell(n)})^{\alpha_n} ((c_{\ell(n)})^{\alpha_n} \tau_{\ell(n)\ell(n+1)})^{-\theta} \right)^{\beta_n} \times (T_{\ell(N)})^{\alpha_N} ((c_{\ell(N)})^{\alpha_N} \tau_{\ell(N)j})^{-\theta}}{\Theta_j}.$$

Note further that all final goods sold in j , command the same expected price regardless of the actual chain, and thus $\Pr(\Lambda_{k \rightarrow i}^n, j)$ corresponds to the share of total spending in country j associated with chains that flow through k at position $n \leq N - 1$ and through i at position $n + 1$ before reaching country j after assembly. Moreover, the value of the trade flow between countries k and i at positions n and $n + 1$ is a share β_n of the total spending on that chain in country j .¹⁵ The latter spending comprises final-good consumption ($w_j L_j$) and spending in the intermediate input bundle $((1 - \gamma_j) / \gamma_j \times w_j L_j)$. To find the overall spending of intermediate input purchases by firms in i importing from firms in k immediately upstream from them, we thus just need to aggregate across destinations markets j and neighboring stages n and $n + 1$ to obtain

$$\mathcal{X}_{ki} = \sum_{j \in \mathcal{J}} \sum_{n=1}^{N-1} \beta_n \Pr(\Lambda_{k \rightarrow i}^n, j) \frac{1}{\gamma_j} w_j L_j.$$

Together with the input flows associated with the more standard roundabout structure of production, we finally obtain that the share of input purchases by firms in i originating in country k is given by:

$$\pi_{ki}^X = \frac{\pi_{ki}^F \frac{1-\gamma_i}{\gamma_i} w_i L_i + \sum_{j \in \mathcal{J}} \sum_{n=1}^{N-1} \beta_n \Pr(\Lambda_{k \rightarrow i}^n, j) \frac{1}{\gamma_j} w_j L_j}{\sum_{k' \in \mathcal{J}} \pi_{k'i}^F \frac{1-\gamma_i}{\gamma_i} w_i L_i + \sum_{k' \in \mathcal{J}} \sum_{j \in \mathcal{J}} \sum_{n=1}^{N-1} \beta_n \Pr(\Lambda_{k' \rightarrow i}^n, j) \frac{1}{\gamma_j} w_j L_j}. \quad (18)$$

Although computing these intermediate input shares is somewhat cumbersome, notice that our model provides an explicit expression for these shares, which have an empirical counterpart in world Input-Output databases (see more on this in section 6).

¹⁴In the Eaton and Kortum (2002) model, these are the only type of intermediate input flows and thus there is a unique ‘trade share’ π_{ij} regardless of the nature of the goods flowing between country i and country j .

¹⁵This can be verified by iterating (2) and referring to the definition of β_n in (4).

4.3 General Equilibrium

So far, we have characterized trade flows as a function of the vectors of equilibrium wages $\mathbf{w} = (w_1, \dots, w_J)$ and of input bundle costs $\mathbf{P} = (P_1, \dots, P_J)$. We next describe how these vectors are pinned down in general equilibrium.

Notice first that invoking (14) and $c_i = (w_i)^{\gamma_j} (P_i)^{1-\gamma_j}$, we can solve for the vector \mathbf{P} as a function of the vector \mathbf{w} from the system of equations:

$$P_j = \kappa \left(\sum_{\ell \in \mathcal{J}^N} \prod_{n=1}^N \left((c_{\ell(n)})^{-\theta} T_{\ell(n)} \right)^{\alpha_n \beta_n} \times \prod_{n=1}^{N-1} (\tau_{\ell(n)\ell(n+1)})^{-\theta \beta_n} \times (\tau_{\ell(N)j})^{-\theta} \right)^{-1/\theta}. \quad (19)$$

To solve for equilibrium wages, notice that for all GVCs, stage n value added (or labor income) accounts for a share $\gamma_{\ell(n)} \alpha_n \beta_n$ of the value of the finished good emanating from that GVC. Furthermore, total spending in any country j is given by the sum of final-good spending ($w_j L_j$) and spending in the intermediate input bundle ($(1 - \gamma_j) / \gamma_j \times w_j L_j$). The share of that spending going to GVCs in which country i is in position n is given by $\Pr(\Lambda_i^n, j) = \sum_{\ell \in \Lambda_i^n} \pi_{\ell j}$, where remember that we have defined $\Lambda_i^n = \{\ell \in \mathcal{J}^N \mid \ell(n) = i\}$ and $\pi_{\ell j}$ is given in equation (11). It thus follows that the equilibrium wage vector is determined by the solution of the following system of equations

$$\frac{1}{\gamma_i} w_i L_i = \sum_{j \in \mathcal{J}} \sum_{n \in \mathcal{N}} \alpha_n \beta_n \times \Pr(\Lambda_i^n, j) \times \frac{1}{\gamma_j} w_j L_j. \quad (20)$$

The system of equations is nonlinear because $\Pr(\Lambda_i^n, j)$ is a nonlinear function of wages themselves, and of the vector \mathbf{P} , which is in turn a function of the vector of wages \mathbf{w} .

When $N = 1$, we have that $\alpha_N \beta_N = 1$ and $\Pr(\Lambda_i^1, j) = \pi_{ij} = (\tau_{ij} c_i)^{-\theta} T_i / \sum_k (\tau_{kj} c_k k)^{-\theta} T_k$. The equilibrium then boils down to a simple generalization of the general equilibrium in Eaton and Kortum (2002) and Alvarez and Lucas (2007), with cross-country variation in how the composite input aggregates value added and the bundle of intermediate inputs.

In Online Appendix B.2, we build on Alvarez and Lucas (2007) to show that, given a vector of wages \mathbf{w} , the system of equations in (19) delivers a unique vector of input bundle costs \mathbf{P} . In that Appendix, we also demonstrate the existence of a solution $\mathbf{w}^* \in \mathbb{R}_{++}^J$ to the system of equations in (20) – with (19) plugged in – and we derive a set of sufficient conditions that ensure that this solution is unique.

4.4 Gains from Trade

We next study the implications of our framework for how changes in trade barriers affect real income in all countries. Consider a ‘purely-domestic’ value chain that performs all stages in a given country i to serve consumers in the same country i . Let us denote this chain $\ell = (i, i, \dots, i)$ by i^N .

From equation (11), such a value chain would capture a share of country i 's spending equal to

$$\pi_{iN} = \Pr(i^N) = \frac{(\tau_{ii})^{-\theta} \left(1 + \sum_{n=1}^{N-1} \beta_n\right) \times (c_i)^{-\theta} T_i}{\Theta_i},$$

where we have used the fact that $\sum_{n=1}^N \alpha_n \beta_n = 1$. Combining this equation with (14) and $c_i = (w_i)^{\gamma_i} (P_i)^{1-\gamma_i}$, we can establish that real income in country i can be expressed as

$$\frac{w_i}{P_i} = \left(\kappa (\tau_{ii})^{1 + \sum_{n=1}^{N-1} \beta_n} \right)^{-1/\gamma_i} \left(\frac{T_i}{\pi_{iN}} \right)^{1/(\theta\gamma_i)}. \quad (21)$$

Because under autarky $\pi_{iN} = 1$, we can conclude that the (percentage) real income gains from trade, relative to autarky, are given by $(\pi_{iN})^{-1/(\theta\gamma_i)} - 1$. This formula is analogous to the one that applies in the Eaton and Kortum (2002) framework (and the wider class of models studied by Arkolakis et al., 2012). An important difference, however, is that π_{iN} is not the share of spending on domestic finished goods (π_{ii}^F in equation (17)), but rather the share of spending on goods that *only* embody domestic value added. The latter share π_{iN} is necessarily lower than π_{ii}^F (and increasingly so, the larger number of stages), and thus the gains from trade emanating from our model are expected to be larger on this account. This result is similar to the one derived by Melitz and Redding (2014) in an Armington framework with sequential production, and also bears some resemblance to Ossa's (2015) argument that the gains from trade can be significantly larger in a multi-sector models, with stages in our model playing the role of sectors in his framework. One can also show that our Cobb-Douglas assumption in technology is not essential for this result: the gains from trade would still be given $(\pi_{iN})^{-1/(\theta\gamma_i)} - 1$ for any CES multi-stage production technology with an elasticity of substitution lower than one between the value added at different stages.¹⁶

Another key distinctive feature of the formula in (21) is that, unlike π_{ii}^F , π_{iN} cannot be directly observed in the data, and thus the sufficient statistic approach advocated by Arkolakis et al. (2012) is not feasible in our setting. Instead, one needs a model to structurally back out π_{iN} from available data. For a similar reason, the hat algebra approach to counterfactual analysis proposed by Dekkle et al. (2008) is not feasible in our setting.

Although we have argued above that $\pi_{iN} < \pi_{ii}^F$ translates into large gains from trade than in models without sequential production, it should be noted that the values of γ_i and θ that are appropriate for our model might be different from those appropriate for a model without multi-stage production. First, remember that our model features an additional type of intermediate input flows relative to a model with roundabout production. In order to match the empirical ratio of value added to gross output in each country, our model will thus require setting relatively higher values of γ_i , which other things equal, will lead to lower gains from trade. As for the parameter θ governing the elasticity of trade flows to iceberg trade costs, we can no longer rely on simple gravity equation specifications to back out that parameter. Furthermore, our model suggests that

¹⁶We thank Arnaud Costinot for this observation. It should be clear, however, that the π_{iN} one would back out from available data would depend on the multi-stage production function one specifies.

the trade elasticity should on average be lower for intermediate inputs than for final goods, a prediction we will find some suggestive empirical support for below. This suggests that the proper way to calibrate our model entails setting a value of θ higher than the one that would be suitable to calibrate a Ricardian model without multi-stage production. As in the case of γ_i , the use of a larger value of θ would again generate a downward correction to the gains from trade. Overall, whether our model generates larger or smaller gains from trade than models without multi-stage production is an empirical question, one which we will explore in section 7.

4.5 The Centrality-Downstreamness Nexus

We next exploit the tractability of our framework to formally explore the role of a country's geography (and, in particular, its centrality) in shaping its average position in GVCs. In order to isolate the role of geography in shaping GVC positioning, we further focus on the 'symmetric' case $\alpha_n \beta_n = 1/N$ for all stages n , which amounts to assuming $\alpha_n = 1/n$ and $\beta_n = n/N$ for all $n \in \mathcal{N}$. Without this assumption, technology would not be symmetric in the value added originated at different stages, and thus the state of technology T_i of a country would affect different stages differentially, thereby generating *technological* comparative advantage.¹⁷

In order to formalize a centrality-downstreamness nexus, let us define the average upstreamness of production of a given country i in value chains that seek to serve consumers in country j , by

$$U(i; j) = \sum_{n=1}^N (N - n + 1) \times \frac{\Pr(\Lambda_i^n, j)}{\sum_{n'=1}^N \Pr(\Lambda_i^{n'}, j)}, \quad (22)$$

where remember that $\Pr(\Lambda_i^n, j)$ is the probability that country i features in position n in value chains leading to consumption in country j , and corresponds to $\Pr(\Lambda_i^n, j) = \sum_{\ell \in \Lambda_i^n} \pi_{\ell j}$, with $\pi_{\ell j}$ given in equation (11). The index $U(i; j)$ in (22) is thus a weighted average distance of country i from final consumers in value chains that service consumers in country j , and it is closely related to measure proposed by Antràs et al. (2012). Although $U(i; j)$ in equation (22) uses probabilities rather than expenditure shares as weights, it can be verified that given our symmetry assumption, these probabilities correspond to the share of country i 's value added at each stage in chains ending in consumption in country j .

We seek to establish a connection between the measure of upstreamness $U(i; j)$ and the centrality of country i . As in section 2, the structure of equation (11) already hints at a negative association between the two, since high values of trade costs (high τ_{ij}) in relatively downstream stages (high n) have a disproportionately negative effect on the likelihood of a given permutation of countries forming an equilibrium value chain. In order to develop a more precise formulation of this result, we assume that the easiness of trade between any two countries i and j can be decomposed as $(\tau_{ij})^{-\theta} = \rho_i \rho_j$, where we take ρ_i to be an index of the *centrality* of country i . Notice that if

¹⁷For instance, if downstream stages contributed more to overall value added than upstream stages, we would obtain a prediction analogous to that in Costinot et al. (2013), namely that countries with better technologies T_i have comparative advantage in downstream stages.

		Input use & value added			Final use			Total use
		Country 1	...	Country J	Country 1	...	Country J	
Intermediate inputs supplied	Country 1	X_{11}	...	X_{1J}	F_{11}	...	F_{1J}	Y_1

	Country J	X_{J1}	...	X_{JJ}	F_{J1}	...	F_{JJ}	Y_J
Value added		$w_1 L_1$...	$w_J L_J$				
Gross output		Y_1	...	Y_J				

Figure 3: A schematic world Input-Output table.

country i is more central than country j , then it is cheaper to ship from i to any other country in the world than it is to ship from country j . This is a rather strong notion of centrality but it has the virtue of providing the following stark result (which we prove in Appendix A.4):

Proposition 1 The more central a country i is (i.e., the higher is ρ_i), the lower is the average upstreamness $U(i; j)$ of this country in global value chains leading to consumers in any country $j \in J$.

In the next section, we will provide suggestive evidence consistent with this prediction.

5 Data Sources and Suggestive Evidence

The structural estimation of our model relies exclusively on aggregate world Input-Output data. To understand the nature of the data, it is useful to refer to the schematic representation of a WIOT in Figure 3. This matrix is split into two blocks with the block on the left containing data on bilateral intermediate input trade flows across countries (denoted by X_{ij}) and the block on the right containing the data for final good trade flows (denoted by F_{ij}). Each row (column) represents the sales (purchases) of each country to (from) every other country. More specifically, each row corresponds to the sales of a country for production in every other country (intermediate input sales) and for consumption in every other country (final good sales). Hence, the sum across a row equals a country's gross output (denoted by Y_i). Meanwhile, columns in the left block contain intermediate input purchases by each country so that the sum across a column equals gross output minus value added (the latter denoted by $w_i L_i$). Finally, summing down a column on the right block delivers aggregate final good consumption.¹⁸

Given the cells of a WIOT, it is straightforward to construct empirical analogs to our model's key equilibrium variables, namely the final-good π_{ij}^F and intermediate input shares π_{ij}^X in (17) and (18), as well as gross output and value added in each country. More specifically, denoting by a 'hat'

¹⁸Note that the difference between aggregate final consumption and value-added is the trade deficit/surplus. These deficits are nontrivial for certain countries and are taken into account in both our estimation and counterfactual exercises, as discussed below.




Figure 4: Some Key Features of the World Input Output Database

these empirical moments, we have

$$\hat{\pi}_{ij}^F = \frac{F_{ij}}{\sum_{i' \in \mathcal{J}} F_{i'j}}, \quad \hat{\pi}_{ij}^X = \frac{X_{ij}}{\sum_{i' \in \mathcal{J}} X_{i'j}}, \quad \hat{Y}_j = \sum_{i \in \mathcal{J}} X_{ji} + \sum_{i \in \mathcal{J}} F_{ji}, \quad \hat{w}_j L_j = \hat{Y}_j - \sum_{i \in \mathcal{J}} X_{ji}. \quad (23)$$

Building a WIOT of the type in Figure 3 is a formidable endeavor because it requires collecting trade and production data from many different sources, including national and supra-national statistical offices, but also because it necessarily requires assumptions and data analysis in order to make the data comparable. In this paper we work, for the most part, with the World Input Output Database (or WIOD for short), the outcome of a project was carried out by a consortium of 12 research institutes headed by the University of Groningen in the Netherlands (see Timmer et al., 2015). We choose this dataset for our benchmark estimation because we believe that the assumptions put into its construction are less heroic than those contained in other sources. The main limitation of the WIOD is that it only covers 43 relatively developed countries, and includes no African country and only one country in Latin America (Brazil).¹⁹ With that in mind, we will also estimate our model using the more comprehensive Eora MRIO database, which provides yearly world input-output tables covering 190 countries and the rest of the world for the period 1990-2013. Although we find this dataset less reliable than the WIOD, its broader sample will enrich the set of counterfactuals studied in section 7.

Figure 4 depicts some salient characteristics of the data we employ in our structural estimation. The left panel plots the diagonal elements $\hat{\pi}_{ii}^F$ of the final-good matrix against the diagonal elements $\hat{\pi}_{ii}^X$ of the input matrix in the WIOD for the year 2013, with the size of each observation being proportional to each country's GDP. As is clear from the graph, most observations lie above the 45 degree line, indicating that the input $\hat{\pi}_{ij}^X$ and output $\hat{\pi}_{ij}^F$ matrices are asymmetric, and that

¹⁹Two releases of the WIOD are available. The 2016 release contains a WIOT covering 43 countries and the rest of the world for the period 2000-2014. A previous release (2013) contained information for 40 countries and the rest of the world, for the period 1995-2011. See <http://www.wiod.org>.

countries tend to rely on foreign sources more prevalently for inputs than for final goods. This pattern, which is also observed in the Eora dataset (see Figure A.2 in the Appendix), is consistent with the notion that trade costs are more detrimental for downstream versus upstream stages. The right panel of Figure 4 plots each country’s gross output to final-good sales (GO/F) against its gross output to value added ratio (GO/VA). In a closed-economy, these two ratios would naturally coincide and all observations would lie on the 45 degree line. In a globalized world, differences in these ratios provide a rough measure of the positioning of countries in GVCs. More specifically, for a given GO/VA , a high share of final output in overall sales (i.e., a low GO/F in the vertical axis) would indicate that a country is relatively downstream in GVCs. Similarly, for a given GO/F , a high ratio GO/VA would indicate that a country uses a relatively large amount of foreign inputs in production, which again would suggest a relatively downstream position of this country in GVCs. With this background, the figure indicates that although there are some deviations from the 45 degree line, cross-country variation in these ratios is much larger than within-country differences. These considerations motivate the introduction of cross-country variation in the value-added shares γ_i in our model. Without such variation, it would be impossible for us to match the ratios in the right-panel of 4.

Suggestive Evidence

Having described our main data sources, we complete this section by exploring the empirical relevance of a key mechanism of the model – namely, the fact that the trade elasticity is larger for downstream stages than for upstream stages – and of the centrality-downstreamness nexus result in Proposition 1. The empirical analysis in this section is reduced-form in nature and not structurally related to our model, but we deem them to be informative nonetheless. We also note that the tests below use additional sources of data that are not employed in the structural estimation, so we will only discuss them succinctly as we introduce them.

We begin by studying empirically the compounding effect of trade costs, a key feature of our model. A crude way to assess the differential sensitivity of trade flows to trade costs at different stages of the value chain is to compare the elasticity of intermediate-input and final-good flows to various proxies for trade costs τ_{ij} . In particular, and building on the gravity equation literature, consider projecting the bilateral trade cost parameters τ_{ij} on a vector of pair-specific variables including distance, contiguity and a common language indicator. More specifically, let

$$\ln \tau_{ij} = \ln \kappa + \delta_{dist} \ln Distance_{ij} + \delta_{con} Contiguity_{ij} + \delta_{lang} SameLanguage_{ij}.$$

As long as the coefficients δ_{dist} , δ_{con} and δ_{lang} are common for intermediate inputs and final goods, then any difference in the sensitivity of final-good versus intermediate-input trade flows to these bilateral gravity variables will be indication of a differential sensitivity of ‘upstream’ versus ‘downstream’ trade flows to trade costs. To assess the plausibility of this approach, consider the case of the distance elasticity δ_{dist} . Our key identification assumption in that case is that trade

costs, as a percentage of the value of the good being shipped, are identical regardless of whether the good is an input or a final-good. If we observe final-good trade being more sensitive to distance than input trade, we will then conclude that final-good trade is more sensitive to trade costs than input trade is.

We implement this test in Table 1. In columns (1) and (2) of that table we report the results of a standard gravity specification in which the log of *aggregate* shipments from country i to country j are run on exporter and importer fixed effects, as well as the log of distance between i and j and dummy variables for whether i and j share a contiguous or share a common border. Our shipments data are from 2011 and correspond to the Eora MRIO database (described above). The data cover 190 countries and include information on domestic shipments (i.e., sales from i to i). The gravity variables are from the CEPII dataset for the year 2006 (the most recent one available), and the merge between these two datasources leaves us with information on 180 countries.

Table 1. Trade Cost Elasticities for Final Goods and Intermediate Inputs

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Distance	-1.111*** (0.019)	-0.823*** (0.014)	-1.144*** (0.019)	-0.851*** (0.014)	-1.210*** (0.021)	-0.903*** (0.015)	-0.794*** (0.015)
Distance \times Input					0.133*** (0.006)	0.106*** (0.006)	0.098*** (0.006)
Contiguity		2.187*** (0.111)		2.198*** (0.112)		2.287*** (0.120)	1.184*** (0.099)
Contiguity \times Input						-0.177*** (0.037)	-0.054 (0.040)
Language		0.480*** (0.026)		0.507*** (0.027)		0.596*** (0.029)	0.513*** (0.027)
Language \times Input						-0.179*** (0.013)	-0.169*** (0.013)
Domestic							5.635*** (0.187)
Domestic \times Input							-0.599*** (0.067)
Observations	32,400	32,400	64,800	64,800	64,800	64,800	64,800
R^2	0.98	0.982	0.972	0.974	0.972	0.974	0.976

Notes: Standard errors clustered at the country-pair level reported. ***, **, and * denote 1, 5 and 10 percent significance levels. All regressions include exporter and importer fixed effects. Regressions in columns (3)-(7) also include a dummy variable for inputs flows. See Appendix ?? for details on data sources.

Our results in columns (1) and (2) are fairly standard. Distance reduces trade flows with an elasticity of around -1 , while contiguity and common language have a sizeable positive effect on

bilateral flows. Starting in column (3), we exploit a key advantage of the Eora MRIO database, namely the fact that it reports separately bilateral shipments of intermediate inputs (X_{ij}) and of finished goods (F_{ij}). In columns (3) and (4), we pool these observations and re-run the specifications in columns (1) and (2), while clustering at the country-pair level. It is clear that the results are almost identical to those in columns (1) and (2). More interestingly, in column (5) we document that the elasticity of trade flows to distance is significantly larger for final-good trade (-1.210) than for intermediate-input trade (-1.077). The difference is sizeable and highly statistically significant. In column (6), we document a similar phenomenon: the positive effect of contiguity and common language on trade flows is significantly attenuated when focusing on the intermediate-input component of trade. Finally, in column (7) we introduce a dummy variable for intranational shipments as well as its interaction with input trade. As is well-known from the border-effect literature, the domestic trade dummy is very large, but we again observe that it is significantly lower for input trade, a result consistent with the scatter plot in the left panel of Figure 4 discussed above.

Taken together, the results in Table 1 are highly suggestive of trade barriers impeding trade more severely in downstream stages than in upstream stages. In Online Appendix B.4, we further show that our results are not materially affected when pooling data from all years (1995-2013) for which the Eora dataset is available (instead of just using 2011 data).²⁰ We also repeat our tests using data from the two releases of the WIOD database, which cover a smaller and more homogenous set of countries. The results with the 2013 release of the WIOD continue to indicate a significantly lower distance elasticity and lower ‘home bias’ in intermediate-input relative to final-good trade. Nevertheless, with the 2016 release of the same dataset, we only find support for the second differential effect (see Online Appendix B.4 for details). This last result makes us interpret our results with caution. Another important caveat with the evidence above is that it is based on gravity-style specifications that are inconsistent with our theoretical framework. As equations (17) and (18) indicate, bilateral trade flows of final goods and intermediate inputs will typically be affected by trade costs associated with third countries (see Morales et al., 2014, and Adao et al., 2017, for recent evidence of these third-market effects).

We next turn to examining the empirical relevance of the downstreamness-centrality nexus formalized in Proposition 1. For that purpose, we build on Antràs et al. (2012) who propose a measure of the positioning of countries in GVCs and study how this measure correlates with various country-level variables. More specifically, Antràs et al. (2012) propose a measure of industry “upstreamness” (or average distance of an industry’s output from final use) and then compute the average upstreamness of a country’s export vector using trade flow data from the BACI dataset for the year 2002. Column (1) of Table 2 reproduces exactly their baseline specification, which includes 120 countries, and correlates a country’s upstreamness with its GDP per capita, rule of law, financial development, capital-labor ratio and human capital (schooling).²¹ Only financial development and schooling have a statistically significant partial correlation with upstreamness.

²⁰In fact, we have obtained extremely stable results when running these same regressions year by year for this same period 1995-2013.

²¹The source for each of these variables is discussed in Antràs et al. (2012).

In order to assess the relationship between upstreamness and centrality, we simply add a measure of centrality to these core specification in column (1). In particular, for each country i we compute $C_i^{GDP} = \sum_j (GDP_j / Distance_{ij})$ and $C_i^{pop} = \sum_j (Population_j / Distance_{ij})$, which capture a country's proximity to other countries with either large GDP or large population (or both). We are able to compute these measures for only 118 of the original 120 countries in Antràs et al. (2012), so for completeness, column (2) reproduces the results of running the same specification as in column (1) with only those 118 countries. Clearly, the results are not materially affected. More interestingly, in columns (3) and (4) we document a highly statistically significant negative relationship between upstreamness and each of the two measures of centrality. This partial correlation is not driven by the presence of the other covariates: column (5) shows that it persists when only controlling for GDP per capita, and column (6) demonstrates that it holds even unconditionally. In Online Appendix B.4, we plot this relationship and show that it is not driven by any outliers. Though these correlations cannot be interpreted causally, they are again suggestive of the empirical relevance of the nexus between centrality and downstreamness highlighted in Proposition 1.

Table 2. Export Upstreamness and Centrality

	(1)	(2)	(3)	4)	(5)	(6)
Centrality (GDP weighted)			-0.173 ^{***}		-0.233 ^{***}	-0.155 ^{***}
			(0.065)		(0.061)	(0.044)
Centrality (population weighted)				-0.228 ^{***}		
				(0.084)		
Log(Y/L)	0.083	0.082	0.102	0.046	0.083*	
	(0.142)	(0.142)	(0.138)	(0.148)	(0.046)	
Rule of Law	-0.029	-0.026	0.010	0.010		
	(0.103)	(0.104)	(0.105)	(0.105)		
Credit/Y	-0.437 ^{***}	-0.440 ^{***}	-0.375 ^{***}	-0.407 ^{***}		
	(0.136)	(0.137)	(0.130)	(0.135)		
Log(K/L)	0.156	0.159	0.163	0.188		
	(0.131)	(0.132)	(0.129)	(0.132)		
Schooling	-0.085 ^{***}	-0.085 ^{***}	-0.083 ^{***}	-0.094 ^{***}		
	(0.031)	(0.031)	(0.030)	(0.029)		
Observations	120	118	118	118	118	118
R^2	0.154	0.153	0.194	0.199	0.083	0.056

Notes: Robust standard errors reported. ^{***}, ^{**}, and ^{*} denote 1, 5 and 10 percent significance levels. See Appendix ?? for details on data sources.

6 Estimation

We next turn to a more structural empirical exploration of the model. In our baseline, we will employ data from the WIOD for the year 2013, but we will subsequently replicate our estimation using the broader Eora dataset for the year 2012.

It is useful to begin by outlining the parameters we need to estimate or calibrate for a given number J of countries and N of stages. Geography is pinned down by the $J \times J$ matrix of iceberg trade costs τ_{ij} . Production depends on the Cobb-Douglas input expenditure shares α_n , which are stage-specific but common across countries, while the labor value-added shares γ_j are country-specific but common across stages. There are thus N α_n 's to estimate and J γ_j 's. Lastly, labor productivity depends on a vector $J \times 1$ of country-specific state-of-technology levels T_j , while comparative advantage is governed by a single parameter θ . Although countries are also allowed to vary in terms of their supply of equipped labor, the particular values of L_j only affect the estimates of T_j and of equilibrium wages, but have no bearing on the counterfactuals discussed below. In order to keep the estimates of T_i in economically meaningful levels, we simply normalize equipped labor as $L_i = (\text{capital}_i)^{\frac{1}{3}}(\text{population}_i)^{\frac{2}{3}}$, where both capital and population are drawn from the Penn World Tables for the years 2013 (in the Eora estimation) and 2014 (in the WIOD estimation).

To pin down trade costs, we follow the method proposed by Head and Ries (2001) and make the simplifying assumption that domestic trade costs are common across countries and normalized to 0, i.e. $\tau_{jj} = 1$ for all $j \in \mathcal{J}$. International trade costs, up to a power $-\theta$, can then be immediately read off the data through the use of equation (17) and our empirical analogs in (23):

$$\tau_{ij}^{-\theta} = \sqrt{\frac{\hat{\pi}_{ij}^F \hat{\pi}_{ji}^F}{\hat{\pi}_{ii}^F \hat{\pi}_{jj}^F}}. \quad (24)$$

Trade costs are symmetric by construction, i.e. $\tau_{ij} = \tau_{ji}$, and in practice the triangle inequality (i.e., $\tau_{ij} \leq \tau_{ik}\tau_{kj}$) holds across more than 99.9% of triples.

A consequence of our approach to backing out trade costs is that the calibrated values for $\tau_{ij}^{-\theta}$ we obtain are unaffected by the particular value of θ chosen. Although the value of θ affects the equilibrium of our model beyond its effect on $\tau_{ij}^{-\theta}$, it turns out that the moments we employ for our structural estimation (see below) do *not* identify θ . More precisely, for every possible value of θ , there exists a re-normalization of T_j that yields the same equilibrium (conditional on the same set of parameters γ_j and α_n). With that in mind, we simply set $\theta = 5$ in our estimation. This value is slightly higher than is typically assumed in the literature, but our model predicts that the trade elasticity for final goods (i.e., θ) should be larger than the elasticity one would estimate with overall trade flows (which is a weighted average of $\theta\beta_1, \theta\beta_2, \dots, \theta$).²² We will return to this point below, after having estimated the β_n 's (i.e., the α_n 's).

Another tricky parameter to calibrate is the number of stages N . We initially estimate the

²²Simonovska and Waugh (2014), in a widely cited study, find a range for the elasticity of trade between 2.47 and 5.51. Using U.S. import data, Antràs et al. (2017) estimate an elasticity of trade of 4.54.

model for the case $N = 2$. This is the simplest multi-stage model one could estimate, and as we will discuss extensively below, the data actually appears to favor this value of N over larger ones.

Having pinned down the matrix of trade costs $\tau_{ij}^{-\theta}$, θ , and N , we estimate the remaining parameters of the model by targeting specific moments of a WIOT via the generalized method of moments (GMM). More specifically, we target four sets of moments, each constituting a $J \times 1$ vector: i) the diagonal elements $\hat{\pi}_{ii}^F$ of the final-good matrix; ii) the diagonal elements $\hat{\pi}_{ii}^X$ of the input matrix, iii) the gross output to value added ratio GO_i/VA_i in each country, and iv) the GDP share $\hat{w}_i L_i / \sum_j \hat{w}_j L_j$ of each country.²³

The choice of these moments is motivated by the following considerations. First, note that the vector of state of technology parameters T_i naturally shapes wages and thus (for given L_i) variation in GDP shares across countries. These technology parameters also affect the extent to which countries rely on local versus foreign sources of inputs and final goods (and thus the shares $\hat{\pi}_{ii}^X$ and $\hat{\pi}_{ii}^F$). Meanwhile, the input expenditure shares α_n determine how fast the trade elasticity increases along GVCs, and are thus crucial in shaping the observed differences between the input own share $\hat{\pi}_{ii}^X$ and the output own share $\hat{\pi}_{ii}^F$. With $N = 2$, we need only estimate α_2 (since $\alpha_1 = 1$ by assumption). If α_2 were to be close to 1, the sequentiality of production would become immaterial, and all input trade would reflect a roundabout structure of production, as in Eaton and Kortum (2002). In such a case, we would not expect large asymmetries between $\hat{\pi}_{ii}^X$ and $\hat{\pi}_{ii}^F$. Our previously discussed evidence in the left panel of Figure 4 already indicated the empirical importance of those asymmetries, so we would not expect α_2 to be too close to 1.

Finally, the vector of gross output to value added ratios GO_i/VA_i is a natural target for the vector of country-specific value-added share parameters γ_j . To see this, note that with $N = 1$, the gross output to value-added ratio is given by

$$\frac{GO_i}{VA_i} = \frac{Y_j}{w_j L_j} = \frac{w_j L_j + \frac{1-\gamma_j}{\gamma_j} w_j L_j}{w_j L_j} = \frac{1}{\gamma_j}.$$

When $N > 1$, the expression for the gross output to value-added ratio is more complicated and the other parameters of the model – and most notably the input shares α_n – have an influence over GO_i/VA_i . To see this, consider our estimation with $N = 2$. For a given γ_j , the gross-output to value-added ratio will be close to $1/\gamma_j$ when the upstream stage of production is irrelevant for production (i.e., when $\alpha_2 \rightarrow 1$), since this corresponds to reducing N from 2 to 1. Conversely, when $\alpha_2 \rightarrow 0$, the downstream stage of production adds very little value, and the gross output to value added ratio is close to $2/\gamma_j$, since the same output is shipped twice but value is added essentially only once. In practice, for a general N , the gross output to value added ratio features variation (see the right panel of Figure 4 above) both because countries have different labor value-added shares but also because they find themselves at different degrees of upstreamness along the GVC;

²³In order to guarantee that our model provides a proper quantitative evaluation of the general-equilibrium workings of the world economy, we place a higher weight on matching the empirical moments of larger economies. More precisely, our weighting matrix is a diagonal matrix with GDP shares in the diagonal.

the interaction of both forces determine this statistic.

Before turning to a discussion of our estimation results, we briefly comment on our treatment of trade imbalances. As mentioned above (see footnote 18), these imbalances are empirically nontrivial and correspond to the difference between aggregate final consumption and value added. Following a common approach in the trade literature (see, in particular, Costinot and Rodríguez-Clare, 2015), we treat these deficits as exogenous parameters, and we adjust our general-equilibrium equations to account for the difference between income and spending (see Online Appendix B.3).

Estimation Results

We now turn to discussing our estimation results and overall fit of the model. We mostly focus our discussion on the results we obtain using the WIOD, but at the end of this section, we also briefly describe the our results with the broader Eora database.

As anticipated above, the asymmetries between the input and final-output diagonal elements of the WIOD lead to an estimate of α_2 far removed from one. In particular, we obtain $\alpha_2 = 0.16$. The estimated values for the vectors of γ_j and T_j are reported in Appendix A.5. Figure 5 presents a comparison between the data and the targeted moments, with the size of each observation proportional to GDP. The values for the diagonal elements π_{ii}^X , the gross output to GDP ratios, and GDP shares are all estimated very accurately, with correlations equal to 0.99, 0.96, and 0.99 with their empirical counterparts, respectively. The fit of the final-output diagonal elements π_{ii}^F is also very good (the correlation with data is 0.90), but it also presents some slight discrepancies, especially for some small countries (remember that our estimation algorithm weighs observations by country size).

Figure 6 performs a similar comparison between model and data but for moments that were not directly targeted in the estimation. The upper two charts present the non-diagonal elements of π^X and π^F , and those entries are also matched relatively accurately in both cases (with correlations equal to 0.83 and 0.91, respectively). The lower two charts explore how well our model matches the backward and forward GVC participation of various countries, two measures of the positioning of countries in GVCs proposed by Wang et al. (2017). Because we will later explore counterfactuals exercises that illustrate changes in the participation of countries in GVCs, it is desirable that our calibration matches these type of moments properly. The backward GVC participation index measures the share of value added, both domestic and foreign, in final good consumption that is imported through intermediate goods and which was produced upstream. The forward GVC participation index measures the share of domestic value-added that is exported through intermediate goods to be consumed downstream. Our benchmark calibration fits both moments very well with correlations of 0.99 and 0.95, respectively.

We next repeat our estimation with the use of the EORA data for 2013. Though the full EORA database contains 190 countries we consolidate it into a set of 101 country/regions in order to alleviate the burden of calibrating so many parameters.²⁴ Remarkably, we estimate an upstream

²⁴Specifically, we keep all countries with a population of 10 million or more and aggregate the rest into a set of 9

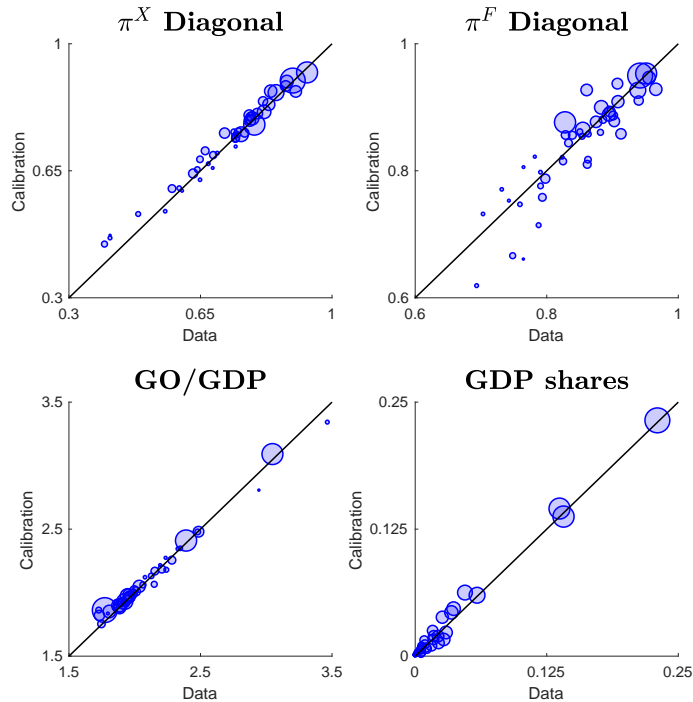


Figure 5: Targeted Moments

input share of $\alpha_2 = 0.19$, which is very similar to the value of $\alpha_2 = 0.16$ found for the WIOD. In Appendix A.5, we provide estimates for the remaining parameters, and also illustrate the fit of the estimation via figures analogous to those in Figure 5 and 6. For both targeted and non-targeted moments, the fit continues to be extremely good.²⁵

It is useful to compare our estimates of α_2 with those implied by our reduced-form results in Table 1, which also used the Eora dataset. Although the gravity-style specification in Table 1 cannot be mapped structurally to our model, the differential effect of distance on input and final-good trade is informative on the relative size of α_2 . More specifically, the ratio of the elasticity of stage 1 output trade to stage 2 output trade is given by $1 - \alpha_2$ in our model. Given the distance elasticities estimated in column (7) of Table 1, and assuming that all input trade is stage 1 output, we would then infer $\alpha_2 = 1 - 0.696/0.794 = 0.12$. Now, of course, in our model not all input trade is stage 1 output, since value added is combined with a bundle of materials at each stage, and the trade elasticity of that “roundabout” input trade is equal to that of final good trade. Using the structural estimates of our model we find that around 18 percent of input trade takes this “roundabout” form. The actual elasticity of stage 1 input is thus lower than is implied by the

regions: Latin America and Caribbean, Central Europe, Eastern Europe, Western Europe, Scandinavia, Middle East and North Africa, Sub-Saharan Africa, Central Asia, and East Asia and Pacific

²⁵The correlations between data and model for the four targeted moments are 0.96, 0.93, 0.95, and 0.99, respectively. The correlations for the off diagonal elements of π^X and π^F are 0.83 and 0.9, while the correlations for the backward and forward participation index stand at 0.83 and 0.68.

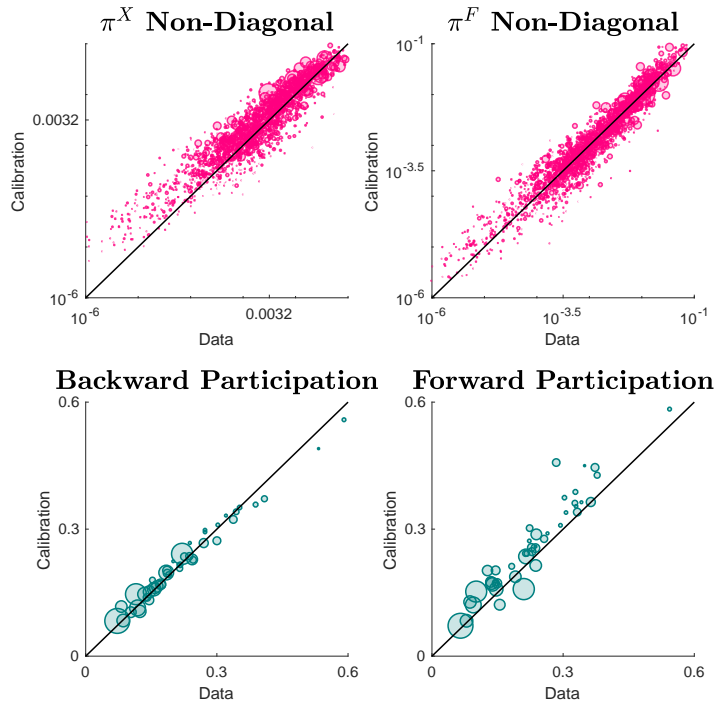


Figure 6: Untargeted Moments.

results in Table 1 (0.675 rather than 0.696), and the implied α_2 is slightly larger ($\alpha_2 = 0.15$), and very close to the one we have estimated structurally.

Revisiting the Calibration of N

Up to now, we have fixed the number of stages to $N = 2$. Estimating our model for $N > 2$ is computationally more demanding but straightforward to carry out. In terms of the parameters to estimate, notice that this only amounts to estimating a longer vector of input shares α_n . Perhaps surprisingly, we have found that the structural estimation shuts down production stages that are more than one stage removed from final consumption, and delivers estimates for the other parameters that are identical to those in the benchmark with $N = 2$. To give a precise example, when we estimate the model with $N = 3$, our calibration delivers $\alpha_3 = 0.16$ and $\alpha_2 = 1$. The most upstream stage of production, $n = 1$, is thus effectively shut down (i.e., its output is negligible). The recovered parameters for γ_j are exactly the same as in our benchmark calibration while those for T_j are exactly those consistent with our benchmark calibration as well.²⁶

Why does our model reject $N > 3$? A first important point to make is that we are calibrating an *average* N for the *whole* world economy, including sectors in which chains might be large (e.g., in some manufacturing sectors) but sectors in which they might be very short (e.g., certain types

²⁶Note that the model with $N = 3$ involves an additional summatory for $n = 1$ in the country level index Θ_j even when this stage is shut down. Hence, the calibrated T_j 's for $N > 2$ should equal $1/J^{N-2}$ times the T_j 's for the benchmark calibration with $N = 2$.

of sectors). Relatedly, the worldwide ratio of gross output to value added in manufacturing is 3.82 (in the 2014 WIOD), while it is 1.78 for non-manufacturing sectors. The fact that the aggregate value of N appears to be tightly related to the aggregate gross output to value added ratio (which is 2.13 in the WIOD) resonates with the theoretical results in the ‘roundabout’ (Input-Output) model of Fally (2012). Yet, we should stress that Fally’s result does not apply in our setting: by appropriate choices of α_n , a variant of our model with a large number of stages could be made consistent with gross output to value added ratios in the neighborhood of 2. Doing so, however, would demand setting relatively larger values of α_2 , which would then generate excessively high asymmetries between the diagonal elements of the input and final output matrices. Because, our GMM estimation penalizes those excessive deviations, we estimate a relatively low value of α_2 , which then constraints us to set $N = 2$ to match the observed aggregate GO/VA ratios.

Some readers might still object that recovering the same estimated values for $N > 2$ is not synonymous with correctly identifying $N = 2$. For example, the moments that we target may be misspecified or not contain sufficient information for backing out the correct N . We next show through simulations that there is a precise sense in which recovering the same calibration for $N \geq 2$ implies that the true N is indeed equal to 2 and that $N > 2$ can be rejected.

Let us work with several values for the chain length ranging from $N = 1$, which is the model of Eaton and Kortum (2002), all the way to $N = 4$, and for a set of $J = 5$ imaginary countries. For each N , we simulate a set of primitives of the model and compute the general equilibrium. We then take the resulting simulated WIOT entries and apply our GMM estimation method with the exact same four sets of moments as above. Furthermore, for each true value of N we run our calibration for various possible values for the number of stages, i.e., $\hat{N} = 1, 2, 3, 4$. The spirit of the exercise is thus to examine whether our estimation method can successfully recover the true value of N . We perform this exercise 100 times for each N , so this amounts to 1600 calibrations in total (100 simulations for each N and four estimations per simulation, for each value $\hat{N} = 1, 2, 3, 4$).

Figure 7 plots our simulation results split into four panels, one for each true value of N . In each panel, the x -axis plots the value of the objective function that the calibration minimizes, i.e. the difference between the observed and estimated moments, while the y -axis plots the sum of squared percentage differences between the true values of the parameters underlying the simulated data and those estimated in the calibration (note the log-scale on both axis). In a nutshell, a lower value in the x -axis implies that our calibration is fitting the data more accurately while a lower value in the y -axis implies that our calibrated parameters are closer to the true parameters. Obviously, estimations that place us in the bottom left corner of each plot are particularly accurate.

Focus first on the top left scatterplot for $N = 1$, and note that the calibrations for all \hat{N} do very well. This should not be surprising since the model with $\hat{N} > N$ is more flexible and thus nests a model with a lower \hat{N} . A crucial observation, however, is that all the points in the scatterplot lie in the bottom left of the graph, implying that a good fit of the moments occurs only if the true parameters are recovered. Turn next to the bottom right scatterplot for $N = 4$. In this case only the estimates for the empirical model with $\hat{N} = 4$ fit the data well, and notice that the true

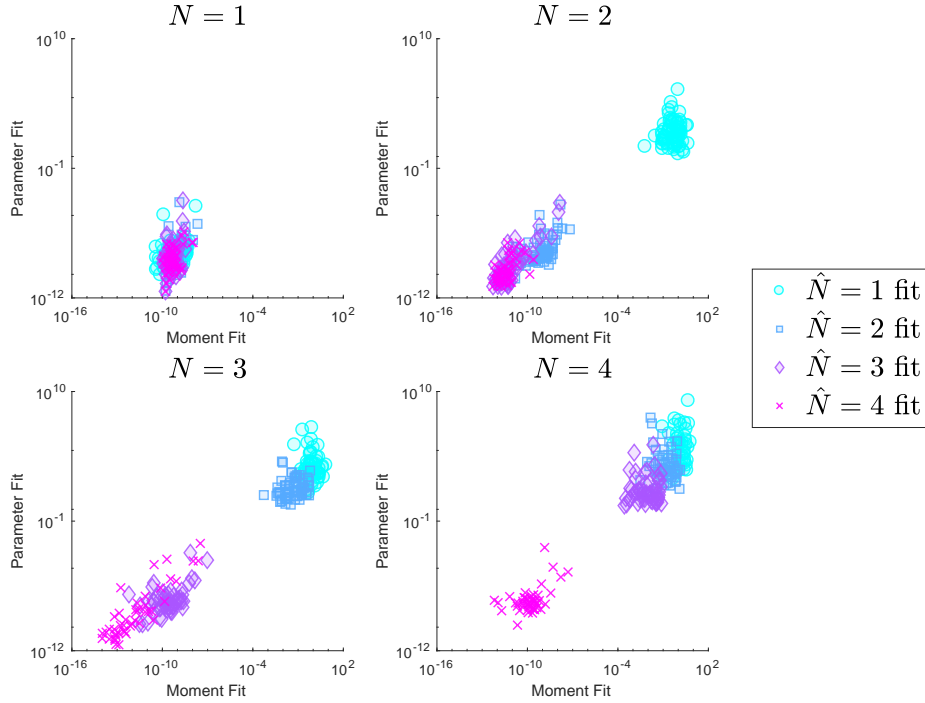


Figure 7: Calibration of N through simulations.

parameters are again recovered. In that same scatter plot, it is clear that the empirical model with $\hat{N} = 3$ does better than $\hat{N} = 2$, and both do better than $\hat{N} = 1$. This is also obvious since lower \hat{N} implies less degrees of freedom. The key takeaway from Figure 7 is that in order to recover the correct N one need only have $\hat{N} \geq N$ and the calibration will be recover the correct parameters regardless of the particular value of \hat{N} . This appears to be analogous to what occurs in our datasets with $N = 2$ and thus, to the extent that the data generating process behind the observable data is consistent with our model, we are able to reject $N > 2$.

It is important to stress, however, that our identification of N relies heavily on our assumption that the matrix of trade costs τ_{ij} and the vector of technology levels T_j is common for inputs and final goods. For example, one can show that an extension of the Eaton and Kortum (2002) framework without multi-stage production (i.e., $N = 1$) could be calibrated to *exactly* match a WIOT, provided that one allows for (i) cross-country variation in value added shares γ_j , and (ii) arbitrary and potentially asymmetric trade costs for inputs and final goods. Intuitively, one could choose an appropriate matrix of input trade costs τ_{ij}^X to reproduce the observed asymmetries in the input and output matrices π_{ij}^X and π_{ij}^F . The vector γ_j could then be set ensure that the GO/VA ratios across countries are exactly nailed, while the technology parameters T_j could be chosen to match the observed cross-country variation in GDP levels. In sum, the data we use cannot reject $N = 1$ if one allows enough flexibility in the modeling of input and output trade costs.

7 Counterfactuals

Having estimated the fundamental parameters of the model and having shown that our framework provides an excellent fit of both the WIOD and Eora datasets, we next explore how counterfactual changes in trade costs, holding other parameters constant, alter the entries of these WIOTs, thereby affecting the real income and positioning of countries in GVCs.

Autarky and Zero Gravity

We begin by revisiting two focal counterfactual exercises in quantitative international trade, namely an increase in trade costs large enough to bring back autarky, and a complete elimination of trade barriers. Both of these counterfactuals are extreme in nature, but they are useful in understanding some distinctive features of our framework.

The real income gains of trade relative to autarky can be computed with the formula $(\pi_{iN})^{-1/(\theta\gamma_i)} - 1$, as indicated by equation (21), although π_{iN} is not directly observable in the data and needs to be inferred from our model. For the sample of countries in the WIOD, the gains from trade range from a value of 3.3 percent for the United States to 75.9 percent for Luxembourg. The left-panel of Figure 8 plots these real income gains for the largest 25 economies in the WIOD sample.²⁷




Figure 8: Gains from trade relative to autarky in GVC model ($N = 2$) versus EK model ($N = 1$), WIOD sample.

The figure also compares these gains to those obtained in a comparison model without multi-stage production calibrated to match the WIOD for the year 2014. This comparison model is a modified Eaton and Kortum (2002) framework, with input trade reflecting roundabout production, but with cross-country variation in value added shares γ_i , and differential (and potentially asymmetric) trade costs for inputs and final goods. As mentioned at the end of section 6, by an appropriate choice of parameters, such a model can always exactly match a WIOT. Furthermore, similarly to Arkolakis et al. (2012), the real income losses from going to autarky can be computed

²⁷This formula still measures the real income gains from trade in the presence of trade imbalances. If autarky entailed a closing of these imbalances, the implications for real spending could be quite different than for real income.

using the formula $\left(\hat{\pi}_{ii}^F (\hat{\pi}_{ii}^X)^{1/\gamma_i^{EK}-1}\right)^{-1/\theta_{EK}} - 1$, where $\gamma_i^{EK} = \hat{w}_i L_i / \hat{Y}_i$ is the value added to gross output ratio in country i , and where the variables with hats can be read off the data as in (23). As explained in section 4.4, the value of θ relevant for this Eaton and Kortum (2002) model (i.e., θ_{EK} in the formula) is naturally smaller than the one relevant for our framework (i.e., $\theta = 5$), since θ_{EK} here corresponds to the elasticity of overall trade, while θ in our GVC model corresponds to the trade elasticity for only final good trade. Using our estimate of α_2 and the relative prevalence of final-good trade, sequential input trade, and ‘roundabout’ input trade in our structural estimation leads us to calibrate $\theta_{EK} = 4.635$, which is very much consistent with available estimates of the overall trade elasticity (see footnote 22).

With these background in mind, the left panel of Figure 8 shows that our model with GVCs generates gains from trade that are generally higher than those emanating from a comparable model without multi-stage production. The differences are, however, modest. Averaging across all 44 countries in the WIOD, the ratio of the (net) gains from trade in our GVC model versus those in a modified Eaton and Kortum model equals 1.075. This is despite the fact that the share π_{iN} of purely domestic GVCs is on average 29% lower than the final-good trade share π_{ij}^F (0.60 versus 0.85). As anticipated in section 4.4, the lower γ_i^{EK} and θ_{EK} (relative to γ_i and θ) attenuate the difference in the real income gains from trade. When computing the worldwide income gains from trade, the differences are even smaller: relative to autarky, world income is 7.0% higher with GVCs and 6.6% higher without them (with the resulting ratio 1.059 depicted as a dashed line in the right panel of Figure 8).

The right-panel of Figure 8 shows, however, that there is quite a lot of variation in the understatement of the gains from trade. China and Mexico, two of USA’s largest trading partners, are the countries for which the EK model underestimates these gains the most (by a factor 1.29 and 1.22, respectively). On the other hand, in a world with sequential production, the gains from trade are lower for certain countries, such as Australia and Russia.

So far, we have discussed our benchmark results with the WIOD. When performing counterfactuals with the broader sample of countries 101 countries and regions in the Eora database, we find similar results. The gains from trade in a world with GVCs are on average a factor XX larger than in a comparable model without multi-stage production, but again there is substantial variation across countries (see Figure A.5 in Appendix A.5). Because the bias is smaller for larger economies the worldwide income gains from trade are only a factor XX larger (XX% versus XX%) in our model relative to a comparable Eaton and Kortum (2002) model.

We next explore the implications of a (hypothetical) complete elimination of trade barriers. The real income consequences of a move to a world with zero gravity are much pronounced. Focusing on the 25 largest economies in the WIOD, Figure 9 shows that these gains range from 163% for the United States to a staggering 913% for Taiwan. Furthermore, these (net) percentage gains are on average a factor 1.359 higher than in a model without multi-stage production. For some countries, such as Norway or Taiwan, the modified EK model underestimates the real income gains by a very large factor (2.41 and 1.77, respectively). Yet, for other countries, the differences are

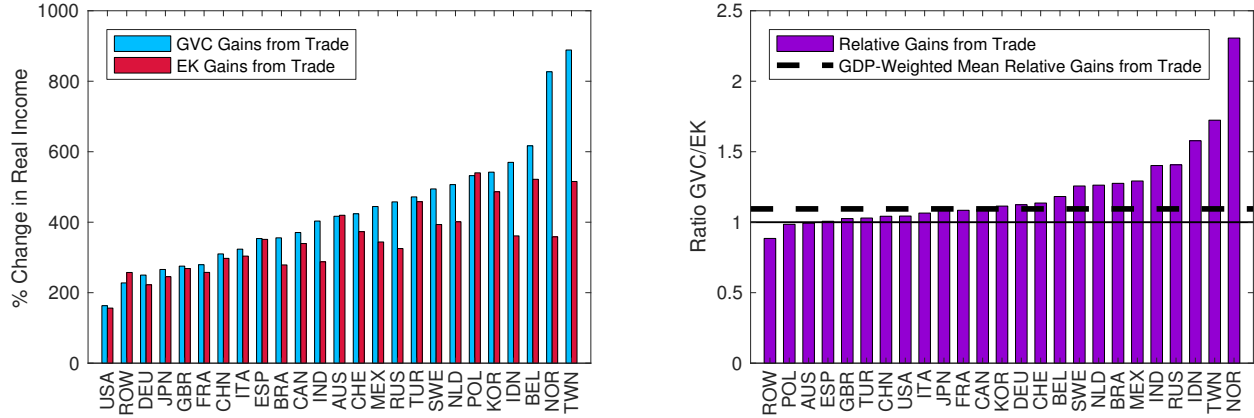


Figure 9: Gains from moving to zero gravity in GVC model ($N = 2$) versus EK model ($N = 1$), WIOD sample.

more modest. Overall, world income grows by 318% in our GVC model versus 283% in a model without multi-stage production (with the associated ratio 1.12 appearing as a dashed line in the right panel of Figure 9). When repeating this exercise for the sample of countries in Eora, we find that worldwide income grows by XX, a factor of XX more than in the comparison E-K model (see Appendix A.5 for an illustration analogous to Figure 9). Figure 10 breaks these worldwide income gains by continent.

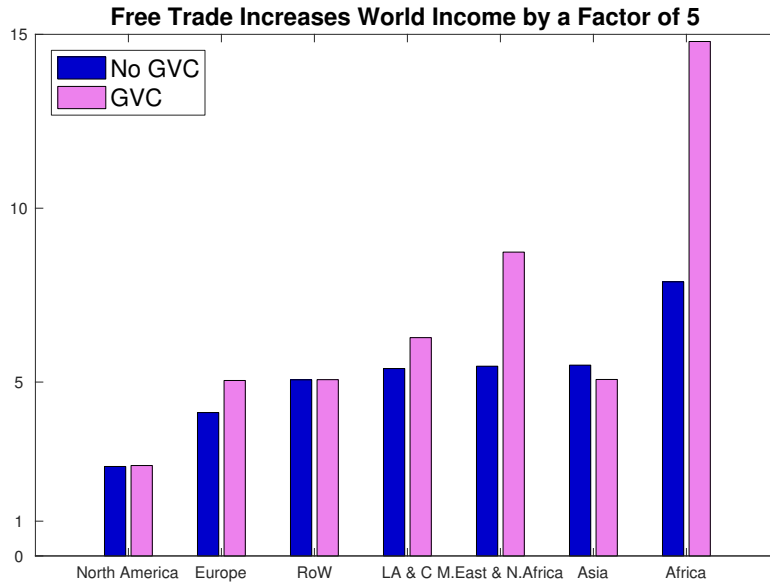


Figure 10: Gains from moving to zero gravity in GVC model ($N=2$) versus EK model ($N=1$), Eora sample

A Fifty Percent Reduction in Trade Barriers

[TO BE COMPLETED]

8 Conclusion

[TO BE WRITTEN]

References

- Adao, Rodrigo, Arnaud Costinot and Dave Donaldson (2017) “Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade,” *American Economic Review*, 107(3): 633-89.
- Alfaro, Laura, Pol Antràs, Paola Conconi, and Davin Chor (2015), “Internalizing Global Value Chains: A Firm-Level Analysis,” NBER Working Paper 21582, September.
- Anderson, James E., and Eric van Wincoop (2003), “Gravity with Gravitas: A Solution to the Border Puzzle,” *American Economic Review* 93(1): 170-192.
- Antràs, Pol, and Davin Chor (2013), “Organizing the Global Value Chain,” *Econometrica* 81(6): 2127-2204.
- Antràs, Pol, Davin Chor, Thibault Fally, and Russell Hillberry, (2012), “Measuring the Upstreamness of Production and Trade Flows,” *American Economic Review Papers & Proceedings* 102(3): 412-416.
- Antràs, Pol, Teresa Fort, and Felix Tintelnot (2017), “The Margin of Global Sourcing: Theory and Evidence from U.S. Firms,” forthcoming *American Economic Review*.
- Antràs, Pol, and Robert W. Staiger (2012), “Offshoring and the Role of Trade Agreements,” *American Economic Review* 102, no. 7: 3140-3183.
- Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare (2012), “New trade models, same old gains?” *The American Economic Review* 102.1: 94-130.
- Baldwin, Richard, and Anthony Venables (2013), “Spiders and Snakes: Offshoring and Agglomeration in the Global Economy,” *Journal of International Economics* 90(2): 245-254.
- Blanchard, Emily, Chad Bown, and Robert Johnson (2016), “Global Supply Chains and Trade Policy”, mimeo Dartmouth College.
- Brancaccio, Giulia, Myrto Kalouptsi, and Theodore Papageorgiou (2017), “Geography, Search Frictions and Trade Costs,” mimeo Harvard University.
- Caliendo, Lorenzo, and Fernando Parro (2015), “Estimates of the Trade and Welfare Effects of NAFTA,” *Review of Economic Studies* 82 (1): 1-44.
- Costinot, Arnaud (2009), “An Elementary Theory of Comparative Advantage,” *Econometrica* 77:4, pp. 1165-1192.
- Costinot, Arnaud, and Andrés Rodríguez-Clare (2015), “Trade Theory with Numbers: Quantifying the Consequences of Globalization” in *Handbook of International Economics*, vol. 4: pp. 197-261.

- Costinot, Arnaud, Jonathan Vogel, and Su Wang, (2013), “An Elementary Theory of Global Supply Chains,” *Review of Economic Studies* 80(1): 109-144.
- de Gortari, Alonso (2017), “Disentangling Global Value Chains,” mimeo Harvard University.
- Dekle, Robert, Jonathan Eaton, and Samuel Kortum (2008), “Global rebalancing with gravity: measuring the burden of adjustment.” *IMF Economic Review* 55.3: 511-540.
- Dixit, Avinash, and Gene Grossman, (1982), “Trade and Protection with Multistage Production,” *Review of Economic Studies* 49(4): 583-594.
- Ethier, Wilfred J. (1984), “Higher Dimensional Issues in Trade Theory,” Ch. 3 in *Handbook of International Economics*, Grossman, Gene M., and Kenneth Rogoff (Eds), Vol. 1.
- Eaton, Jonathan and Samuel Kortum, (2002), “Technology, Geography, and Trade,” *Econometrica*, 70:5, 1741-1779.
- Fally, Thibault (2012), “Production Staging: Measurement and Facts” mimeo University of Colorado Boulder.
- Fally, Thibault, and Russell Hillbery (2015), “A Coasian Model of International Production Chains,” mimeo UC Berkeley.
- Harms, Philipp, Oliver Lorz, and Dieter Urban, (2012), “Offshoring along the Production Chain,” *Canadian Journal of Economics* 45(1): 93-106.
- Head, Keith, and Philippe Mayer (2014), “Gravity Equations: Workhorse, Toolkit, and Cookbook,” Ch. 3 in *Handbook of International Economics*, Gopinath, Gita, Elhanan Helpman and Kenneth Rogoff (Eds), Vol. 4.
- Head, Keith, and John Ries (2001), “Increasing Returns Versus National Product Differentiation as an Explanation for the Pattern of U.S.-Canada Trade,” *American Economic Review*, 91(4): 858-876.
- Hummels, David, Jun Ishii, and Kei-Mu Yi (2001), “The Nature and Growth of Vertical Specialization in World Trade,” *Journal of International Economics* 54 (1): 75–96.
- Johnson, Robert, and Andreas Moxnes (2013), “Technology, Trade Costs, and the Pattern of Trade with Multi-Stage Production,” mimeo.
- Johnson, Robert C. and Guillermo Noguera (2012), “Accounting for Intermediates: Production Sharing and Trade in Value Added,” *Journal of International Economics* 86(2): 224-236.
- Johnson, Robert (2014), “Five Facts about Value-Added Exports and Implications for Macroeconomics and Trade Research,” *Journal of Economic Perspectives*, 28(2), pp. 119-142.

- Lenzen, M., Kanemoto, K., Moran, D., Geschke, A. (2012), “Mapping the Structure of the World Economy,” *Environmental Science and Technology*, 46(15) pp 8374-8381.
- Lenzen, M., Moran, D., Kanemoto, K., Geschke, A. (2013), “Building Eora: A Global Multi-regional Input-Output Database at High Country and Sector Resolution,” *Economic Systems Research*, 25:1, 20-49.
- Kikuchi, Tomoo, Kazuo Nishimura, and John Stachurski (2014), “Transaction Costs, Span of Control and Competitive Equilibrium,” mimeo.
- Kohler, Wilhelm, (2004), “International Outsourcing and Factor Prices with Multistage Production,” *Economic Journal* 114(494): C166-C185.
- Koopman, Robert, William Powers, Zhi Wang, and Shang-Jin Wei (2010), “Give Credit Where Credit Is Due,” NBER Working Paper 16426, September.
- Koopman, Robert, Zhi Wang, and Shang-Jin Wei (2014), “Tracing Value-Added and Double Counting in Gross Exports”, *American Economic Review* 104 (2): 459-94.
- Kremer, Michael, (1993), “The O-Ring Theory of Economic Development,” *Quarterly Journal of Economics* 108(3): 551-575.
- Melitz, Marc J, and Stephen J Redding (2014), “Missing Gains from Trade?” *American Economic Review* 104 (5): pp. 317-21.
- Morales, Eduardo, Gloria Sheu, and Andrés Zahler (2014), “Gravity and Extended Gravity: Using Moment Inequalities to Estimate a Model of Export Entry,” NBER Working Paper No. 19916.
- Ossa, Ralph (2015), “Why Trade Matters After All,” *Journal of International Economics* 97(2): pp. 266-277.
- Sanyal, Kalyan K., and Ronald W. Jones, (1982), “The Theory of Trade in Middle Products,” *American Economic Review* 72(1): 16-31.
- Timmer, Marcel P., Abdul Azeez Erumban, Bart Los, Robert Stehrer, and Gaaitzen J. de Vries (2014), “Slicing Up Global Value Chains.” *Journal of Economic Perspectives* 28, no. 2: 99-118.
- Timmer, Marcel P., Eric Dietzenbacher, Bart Los, Robert Stehrer, and Gaaitzen J. de Vries (2015), “An Illustrated User Guide to the World Input–Output Database: the Case of Global Automotive Production,” *Review of International Economics*, 23: 575–605.
- Tintelnot, Felix (forthcoming), “Global Production with Export Platforms,” *Quarterly Journal of Economics*.
- Tyazhelnikov, Vladimir (2016), “Production Clustering and Offshoring,” mimeo UC Davis.

Wang, Zhi, Shang-Jin Wei, Xinding Yu, Kunfu Zhu (2017), “Measures of Participation in Global Value Chains and Global Business Cycles,” NBER Working Paper 23222, March.

Yi, Kei-Mu, (2003), “Can Vertical Specialization Explain the Growth of World Trade?” *Journal of Political Economy* 111(1): 52-102.

A Appendix

A.1 Increasing Trade-Cost Elasticity

Define $\tilde{p}_{\ell(n)}^{n-1}(\boldsymbol{\ell}) = p_{\ell(n-1)}^{n-1}(\boldsymbol{\ell}) \tau_{\ell(n-1)\ell(n)}$ as the price paid in $\ell(n)$ for the good finished up to stage $n-1$ in country $\ell(n-1)$, so that we can express the sequential unit cost function as

$$p_{\ell(n)}^n(\boldsymbol{\ell}) = g_{\ell(n)}^n \left(c_{\ell(n)}, \tilde{p}_{\ell(n)}^{n-1}(\boldsymbol{\ell}) \right).$$

Define the elasticity of $p_j^F(\boldsymbol{\ell})$ with respect to the trade costs that stage n 's production faces as

$$\beta_n^j = \frac{\partial \ln p_j^F(\boldsymbol{\ell})}{\partial \ln \tau_{\ell(n)\ell(n+1)}},$$

with the convention that $\ell(N+1) = j$ so that β_N^j is the elasticity of $p_j^F(\boldsymbol{\ell})$ with respect to the trade costs faced when shipping assembled goods to final consumers in j . Because $\tau_{\ell(n)\ell(n+1)}$ increases $\tilde{p}_{\ell(n+1)}^n(\boldsymbol{\ell})$ with a unit elasticity, the following recursion holds for all $n' > n$

$$\frac{\partial \ln p_{\ell(n'+1)}^{n'+1}(\boldsymbol{\ell})}{\partial \ln \tau_{\ell(n)\ell(n+1)}} = \frac{\partial \ln p_{\ell(n'+1)}^{n'+1}(\boldsymbol{\ell})}{\partial \ln \tilde{p}_{\ell(n'+1)}^{n'}(\boldsymbol{\ell})} \frac{\partial \ln p_{\ell(n')}^{n'}(\boldsymbol{\ell})}{\partial \ln \tau_{\ell(n)\ell(n+1)}}.$$

At the same time, the unit cost elasticity at stage $n+1$ satisfies

$$\frac{\partial \ln p_{\ell(n+1)}^{n+1}(\boldsymbol{\ell})}{\partial \ln \tau_{\ell(n)\ell(n+1)}} = \frac{\partial \ln p_{\ell(n+1)}^{n+1}(\boldsymbol{\ell})}{\partial \ln \tilde{p}_{\ell(n+1)}^n(\boldsymbol{\ell})}.$$

Hence, the elasticity of finished good prices can be decomposed as

$$\beta_n^j = \prod_{n'=n+1}^N \frac{\partial \ln p_{\ell(n')}^{n'}(\boldsymbol{\ell})}{\partial \ln \tilde{p}_{\ell(n')}^{n'-1}(\boldsymbol{\ell})}, \quad (\text{A.1})$$

invoking the convention $\prod_{n'=N+1}^N f(n') = 1$ for any function $f(\cdot)$. Constant returns to scale in production implies that the function $g_{\ell(n)}^n$ is homogeneous of degree one. As a result, the elasticity of unit costs with respect to input prices is always less or equal than one, so for all $n > 1$ we have

$$\frac{\partial \ln p_{\ell(n)}^n(\boldsymbol{\ell})}{\partial \ln \tilde{p}_{\ell(n)}^{n-1}(\boldsymbol{\ell})} \leq 1,$$

with strict inequality whenever a stage adds value to the product. From equation (A.1), it is then clear that

$$\beta_j^1 \leq \beta_j^2 \leq \dots \leq \beta_j^N = 1,$$

with strict inequality when value added is positive at all stages.

A.2 Fighting the Curse of Dimensionality: Dynamic and Linear Programming

When discussing the lead-firm problem in section 2.2, we mentioned that there are J^N sequences that deliver distinct finished good prices $p_j^F(\ell)$ in country j . Hence, solving for the optimal sequences ℓ^j for all j by brute force requires J^{N+1} computations and is infeasible to do when J and N are sufficiently large. However, we show below that use of dynamic programming surmounts this problem by reducing the computation of all sequences to only $J \times N \times J$ computations. Furthermore, in the special case in which production is Cobb-Douglas the minimization problem can be modeled with zero-one linear programming, for which very efficient algorithms exist.

Dynamic Programming

Define $\ell_n^j \in \mathcal{J}^n$ as the optimal sequence for delivering the good completed up to stage n to producers in country j . This term can be found recursively for all $n = 1, \dots, N$ by simply solving

$$\ell_n^j = \arg \min_{k \in \mathcal{J}} p_k^n \left(\ell_{n-1}^k \right) \tau_{kj}, \quad (\text{A.2})$$

since the optimal source of the good completed up to stage n is independent of the local cost c_j at stage j , the specifics of the cost function g_j^n , or the future path of the good. For this same reason, we have written the pricing function p_k^n in terms of the $n-1$ stage sequence ℓ_{n-1}^k since it does not depend on future stages of production (though it should be clear that p_k^n will also be a function of the production costs and technology available for producers at that chosen location k). The convention at $n=1$ is that there is no input sequence so that $\ell_0^k = \emptyset$ for all $k \in \mathcal{J}$ and the price depends only the input bundle: $p_k^1(\emptyset) = g_k^1(c_k)$.

The formulation in (A.2) makes it clear that the optimal path to deliver the assembled good to consumers in each country j , i.e., $\ell^j = \ell_N^j$, can be solved recursively by comparing J numbers for each location $j \in \mathcal{J}$ at each stage $n \in \mathcal{N}$, for a total of only $J \times N \times J$ computations.

To further understand this dynamic programming approach, Figure A.1 illustrates a case with 3 stages and 4 countries. Instead of computing $J^N = 64$ paths for each of the four locations of consumption, it suffices to determine the optimal source of (immediately) upstream inputs (which entails $J \times J = 16$ computations at stages $n=2$ and $n=3$, and for consumption). In the example, the optimal production path to serve consumers in A , B , and C is $A \rightarrow B \rightarrow B$, while the optimal path to serve consumers in D is $C \rightarrow D \rightarrow D$.

Linear Programming

In the special case in which production is Cobb-Douglas, the optimal sourcing sequence can be written as a log-linear minimization problem

$$\ell^j = \arg \min_{\ell \in \mathcal{J}^N} \sum_{n=1}^{N-1} \beta_n \ln \tau_{\ell(n)\ell(n+1)} + \ln \tau_{\ell(N)j} + \sum_{n=1}^N \alpha_n \beta_n \ln \left(a_{\ell(n)}^n c_{\ell(n)} \right).$$

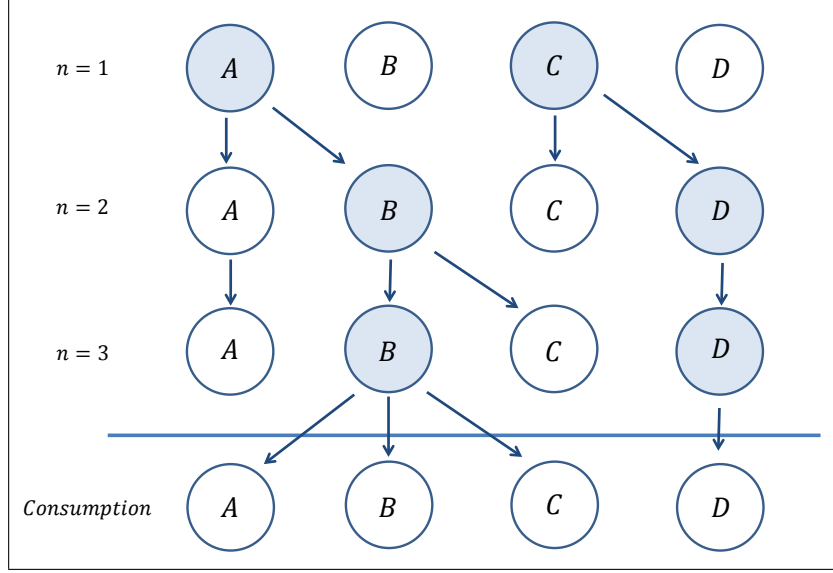


Figure A.1: An Example with Four Countries and Three Stages

This can in turn be reformulated as the following zero-one integer programming problem

$$\begin{aligned}
\ell^j &= \arg \min \sum_{n=1}^{N-1} \beta_n \sum_{k \in \mathcal{J}} \sum_{k' \in \mathcal{J}} \zeta_{kk'}^n (\ln \tau_{kk'} + \alpha_n a_k^n c_k) + \sum_{k \in \mathcal{J}} \zeta_k^N (\ln \tau_{kj} + \alpha_N a_k^N c_k) \\
\text{s.t. } &\sum_{k' \in \mathcal{J}} \zeta_{k'k}^n = \sum_{k' \in \mathcal{J}} \zeta_{kk'}^{n+1}, \forall k \in \mathcal{J}, n = 1, \dots, N-2 \\
&\sum_{k' \in \mathcal{J}} \zeta_{k'k}^{N-1} = \zeta_k^N, \forall k \in \mathcal{J} \\
&\sum_{k \in \mathcal{J}} \zeta_k^N = 1 \\
&\zeta_{kk'}^n, \zeta_k^N \in \{0, 1\}.
\end{aligned}$$

A.3 Decentralized Approach with $N \in \mathbb{N}^+$

This Appendix demonstrates how to generalize our approach with stage-specific randomness and incomplete information to an environment with more than two stages. It should be clear that the input sourcing decisions for the two most upstream stages work in the same way as outlined in section 3.2.B for a general number of stages $N > 2$. Let us quickly recap those decisions. Input producers of good z at the first stage set prices equal to the cost of labor and materials needed to produce a unit of the first-stage good: $p_{\ell(1)}^1(z) = a_{\ell(1)}^1(z) c_{\ell(1)}$. Meanwhile, a producer of z at stage $n = 2$ in country j observes the productivity draws of its tier-one input suppliers and thus sources inputs from $\ell_z^j(1) = \arg \min_{\ell(1) \in \mathcal{J}} \left\{ (a_{\ell(1)}(z) c_{\ell(1)} \tau_{(1)j})^{1-\alpha_2} \right\}$. However, producers at stage $n > 2$ only observe the productivity draws of their tier-one suppliers (i.e., those at $n-1$), and are forced to use their expectations over the productivity draws of upper tier input suppliers in order

to form expectations over the prices at which they will ultimately buy from their tier-one suppliers (at $n - 1$). This is because we have assumed that sourcing decisions are made before observing the prices at which tier-one suppliers will ultimately be able to sell at. In other words, when deciding on their optimal input sources, firms producing at stage $n + 1$ can only form expectations over the input prices from stage $n - 1$ that each of its own possible suppliers producing at stage n faces (or will face).

Let $\ell_z^j(n)$ be the tier-one sourcing decision of a firm producing good z at stage $n + 1$ in j . Generalizing the approach in the main text, define the expectation

$$\mathcal{E}_j^n [s] = \mathbb{E}_n \left[\left(p_{\ell_z^j(n)}^n(z) \tau_{\ell_z^j(n)j} \right)^s \right],$$

for any $s > 0$ and where we have written the expectation with an n subscript indicating that the expectation takes that unit costs (and prices) from stages $1, \dots, n$ as unobserved. To be fully clear, a firm at $n+2$ observes the productivity draws from stage $n+1$ but does not know previous sourcing decisions. Hence it must form an expectation over the location from which its stage n suppliers source, $\ell_z^j(n)$, and use this to calculate the expected input prices $\mathcal{E}_j^n [s]$. As will become clear in the next paragraph, denoting the expectations for a general $s > 0$ is useful since downstream firms between $n + 2, \dots, N$ and final consumers will all use the information on expected input prices at n but in different ways depending on the objective function they seek to minimize.

Substituting in the Cobb-Douglas production process in (2), we can write

$$\mathcal{E}_j^n [s] = \mathbb{E}_n \left[\left(a_{\ell_z^j(n)}^n(z) c_{\ell_z^j(n)} \right)^{\alpha_n s} \times \mathcal{E}_{\ell_z^j(n)}^{n-1} [(1 - \alpha_n) s] \times \left(\tau_{\ell_z^j(n)j} \right)^s \right].$$

The crucial observation is that to determine expected input prices from stage n a firm must also incorporate expected input prices from stage $n - 1$, and so on until input prices from all upstream stages have been incorporated. Note that productivity draws across stages of production are independent, but even more importantly, sourcing decisions across stages of production are also independent. Hence, one can use the law of iterated expectations to compute expected input prices from $n - 1$, $\mathcal{E}_{\ell_z^j(n)}^{n-1} [\cdot]$, in the computation of expected prices at n in $\mathcal{E}_j^n [\cdot]$. The latter expectation is over $\ell_z^j(n)$ but once we condition on a specific value for $\ell_z^j(n)$, the expectation $\mathcal{E}_{\ell_z^j(n)}^{n-1} [\cdot]$ is a constant. Finally, note also that this recursion starts at $n = 1$ with $\mathcal{E}_j^0 [s] = 1$ since only labor and materials are used in that initial stage.

Let us next illustrate why these definitions are useful. Consider the optimal sourcing strategies related to procuring the good finished up to stage $n < N$. Given the sequential cost function in (2), the problem faced by a stage $n + 1$ producer in j can be written as

$$\ell_z^j(n) = \arg \min_{\ell(n) \in \mathcal{J}} \left\{ \left(a_{\ell(n)}^n(z) c_{\ell(n)} \right)^{\alpha_n(1-\alpha_{n+1})} \times \mathcal{E}_{\ell(n)}^{n-1} [(1 - \alpha_n)(1 - \alpha_{n+1})] \times \left(\tau_{\ell(n)j} \right)^{1-\alpha_{n+1}} \right\}.$$

where the $1 - \alpha_{n+1}$ superscript comes from the stage $n + 1$ producer wanting to minimize its own expected input price and in which the stage n input price enters its own unit cost to this power.

Meanwhile, final consumers (or local retailers on their behalf) source their goods by solving

$$\ell_z^j(N) = \arg \min_{\ell(N) \in \mathcal{J}} \left\{ \left(a_{\ell(N)}^N(z) c_{\ell(N)} \right)^{\alpha_N} \times \mathcal{E}_{\ell(N)}^{N-1} [1 - \alpha_N] \times \tau_{\ell(N)j} \right\}.$$

The probability of sourcing inputs from a specific location i at any stage n can be determined by invoking the properties of the Fréchet distribution, given that $1/a_i^n(z)$ is drawn independently (across goods and stages) from a Fréchet distribution satisfying

$$\Pr \left(a_j^n(z)^{\alpha_n \beta_n} \geq a \right) = \exp \left\{ -a^\theta (T_j)^{\alpha_n \beta_n} \right\}.$$

In particular, we obtain

$$\Pr \left(\ell_z^j(n) = i \right) = \frac{\left((T_i)^{\alpha_n} ((c_i)^{\alpha_n} \tau_{ij})^{-\theta} \right)^{\beta_n} \mathcal{E}_i^{n-1} \left((1 - \alpha_n) (1 - \alpha_{n+1}) \right)^{-\beta_{n+1} \theta}}{\sum_{l \in \mathcal{J}} \left((T_l)^{\alpha_n} ((c_l)^{\alpha_n} \tau_{lj})^{-\theta} \right)^{\beta_n} \mathcal{E}_l^{n-1} \left((1 - \alpha_n) (1 - \alpha_{n+1}) \right)^{-\beta_{n+1} \theta}}.$$

These probabilities can now be leveraged in order to compute expected input prices. Define $\tilde{a}_{ij} = (c_i)^{\alpha_n s} \mathcal{E}_i^{n-1} [(1 - \alpha_n) s] (\tau_{ij})^s$ so that $1/(a_i^{\alpha_n s} \tilde{a}_{ij}) \sim \text{Frechet} \left(T_i^{\alpha_n \beta_n} \tilde{a}_{ij}^{-\frac{\beta_n \theta}{s}}, \frac{\beta_n \theta}{s} \right)$ (note that the above distribution is the special case in which $s = 1 - \alpha_{n+1}$).

These probabilities can now be leveraged in order to compute expected input prices. Define $\tilde{a}_{ij} = (c_i)^{\alpha_n s} \mathcal{E}_i^{n-1} [(1 - \alpha_n) s] (\tau_{ij})^s$ so that $1/(a_i^{\alpha_n s} \tilde{a}_{ij}) \sim \text{Frechet} \left(T_i^{\alpha_n \beta_n} \tilde{a}_{ij}^{-\frac{\beta_n \theta}{s}}, \frac{\beta_n \theta}{s} \right)$ (note that the above distribution is the special case in which $s = 1 - \alpha_{n+1}$). Then using the moment generating formula for the Fréchet distribution, it can be verified that

$$\mathcal{E}_j^n [s] = \left[\sum_{l \in \mathcal{J}} T_l^{\alpha_n \beta_n} \tilde{a}_{lj}^{-\frac{\beta_n \theta}{s}} \right]^{-\frac{s}{\beta_n \theta}} \Gamma \left(1 + \frac{\beta_n \theta}{s} \right),$$

where Γ is the gamma function. From this equation it should also be clear why we are denoting $E_j^n [s]$ as a function of s , since as we move down the value chain we need to compute the upstream expectations at different 'moments'.

We are now ready to determine the equilibrium variables: (1) composite prices and (2) the distribution of GVCs. Composite prices can be derived recursively using our expectations:

$$\begin{aligned} P_j &= (\mathcal{E}_j^N [1 - \sigma])^{\frac{1}{1-\sigma}} = \left[\sum_{l \in \mathcal{J}} (T_l)^{\alpha_N} ((c_l)^{\alpha_N} \tau_{lj})^{-\theta} \mathcal{E}_l^{N-1} [(1 - \alpha_N) (1 - \sigma)]^{-\frac{\theta}{1-\sigma}} \right]^{-\frac{1}{\theta}} \Gamma \left(1 + \frac{1 - \sigma}{\theta} \right) \\ &= \left[\sum_{\ell \in \mathcal{J}} \prod_{n=1}^N \left((T_{\ell(n)})^{\alpha_n} ((c_{\ell(n)})^{\alpha_n} \tau_{\ell(n)\ell(n+1)})^{-\theta} \right)^{\beta_n} \right]^{-\frac{1}{\theta}} \prod_{n=1}^N \Gamma \left(1 + \frac{1 - \sigma}{\beta_n \theta} \right)^{\frac{1}{1-\sigma}} \end{aligned}$$

Finally, since input decisions from n are independent from the decisions that firms at $n - 1$ made

then

$$\begin{aligned}
\pi_{\ell j} &= \Pr \left(\ell_z^j(N) = \ell(N) \mid \ell_z^{\ell(N)}(N-1) = \ell(N-1) \right) \times \\
&\quad \times \prod_{n=2}^{N-1} \Pr \left(\ell_z^{\ell(n+1)}(n) = \ell(n) \mid \ell_z^{\ell(n)}(n-1) = \ell(n-1) \right) \times \Pr \left(\ell_z^{\ell(2)}(1) = \ell(1) \right) \\
&= \Pr \left(\ell_z^j(N) = \ell(N) \right) \times \prod_{n=1}^N \Pr \left(\ell_z^{\ell(n+1)}(n) = \ell(n) \right) \\
&= \frac{\prod_{n=1}^{N-1} \left((T_{\ell(n)})^{\alpha_n} \left((c_{\ell(n)})^{\alpha_n} \tau_{\ell(n)\ell(n+1)} \right)^{-\theta} \right)^{\beta_n} \times (T_{\ell(N)})^{\alpha_N} \left((c_{\ell(N)})^{\alpha_N} \tau_{\ell(N)j} \right)^{-\theta}}{\sum_{\ell' \in \mathcal{J}} \prod_{n=1}^{N-1} \left((T_{\ell'(n)})^{\alpha_n} \left((c_{\ell'(n)})^{\alpha_n} \tau_{\ell'(n)\ell'(n+1)} \right)^{-\theta} \right)^{\beta_n} \times (T_{\ell'(N)})^{\alpha_N} \left((c_{\ell'(N)})^{\alpha_N} \tau_{\ell'(N)j} \right)^{-\theta}},
\end{aligned}$$

which is identical to equation (11) in the main text obtained in the ‘randomness-in-the-chain’ formulation of technology.

A.4 Proof of Centrality-Downstreamness Nexus

Let $(\tau_{ij})^{-\theta} = \rho_i \rho_j$. In such a case, the probability of country j sourcing through ℓ reduces to

$$\pi_{\ell j} = \frac{\prod_{m=1}^N \left(T_{\ell(m)} (c_{\ell(m)})^{-\theta} \right)^{\alpha_m \beta_m} \left(\rho_{\ell(m)} \right)^{\beta_{m-1} + \beta_m}}{\sum_{\ell \in \mathcal{J}} \prod_{m=1}^N \left(T_{\ell(m)} (c_{\ell(m)})^{-\theta} \right)^{\alpha_m \beta_m} \left(\rho_{\ell(m)} \right)^{\beta_{m-1} + \beta_m}}$$

and is thus independent of the destination country j . The aggregate probability of observing country i in location n can thus be expressed as

$$\Pr(\Lambda_i^n) = \sum_{\ell \in \Lambda_i^n} \pi_{\ell j} = \frac{\sum_{\ell \in \Lambda_i^n} \prod_{m=1}^N \left(T_{\ell(m)} (c_{\ell(m)})^{-\theta} \right)^{\alpha_m \beta_m} \left(\rho_{\ell(m)} \right)^{\beta_{m-1} + \beta_m}}{\sum_{k \in \mathcal{J}} \sum_{\ell \in \Lambda_k^n} \prod_{m=1}^N \left(T_{\ell(m)} (c_{\ell(m)})^{-\theta} \right)^{\alpha_m \beta_m} \left(\rho_{\ell(m)} \right)^{\beta_{m-1} + \beta_m}}. \quad (\text{A.3})$$

But note that we can decompose this as

$$\begin{aligned}
\Pr(\Lambda_i^n) &= \frac{\left(T_i(c_i)^{-\theta} \right)^{\alpha_n \beta_n} (\rho_i)^{\beta_{n-1} + \beta_n} \times \sum_{\ell \in \Lambda_i^n} \prod_{m \neq n} \left(T_{\ell(m)} (c_{\ell(m)})^{-\theta} \right)^{\alpha_m \beta_m} \left(\rho_{\ell(m)} \right)^{\beta_{m-1} + \beta_m}}{\sum_{k \in \mathcal{J}} \left(T_k(c_k)^{-\theta} \right)^{\alpha_n \beta_n} (\rho_k)^{\beta_{n-1} + \beta_n} \times \sum_{\ell \in \Lambda_k^n} \prod_{m \neq n} \left(T_{\ell(m)} (c_{\ell(m)})^{-\theta} \right)^{\alpha_m \beta_m} \left(\rho_{\ell(m)} \right)^{\beta_{m-1} + \beta_m}} \\
&= \frac{\left(T_i(c_i)^{-\theta} \right)^{\alpha_n \beta_n} (\rho_i)^{\beta_{n-1} + \beta_n}}{\sum_{k \in \mathcal{J}} \left(T_k(c_k)^{-\theta} \right)^{\alpha_n \beta_n} (\rho_k)^{\beta_{n-1} + \beta_n}}
\end{aligned}$$

where the second line follows from the fact that, for GVCs in the sets Λ_i^n and Λ_k^n , the set of all possible paths excluding the location of stage n are necessarily identical (and independent of the country where n takes place), and thus the second terms in the numerator and denominator of the first line cancel out.

For the special symmetric case with $\alpha_n \beta_n = 1/N$ and $\alpha_n = 1/n$ we obtain that

$$\Pr(\Lambda_i^n) = \frac{\left(T_i(c_i)^{-\theta}\right)^{\frac{1}{N}} (\rho_i)^{\frac{2n-1}{N}}}{\sum_{k \in \mathcal{J}} \left(T_k(c_k)^{-\theta}\right)^{\frac{1}{N}} (\rho_k)^{\frac{2n-1}{N}}}$$

Now consider our definition of upstreamness

$$U(i) = \sum_{n=1}^N (N - n + 1) \times \frac{\Pr(\Lambda_i^n)}{\sum_{n'=1}^N \Pr(\Lambda_i^{n'})}. \quad (\text{A.4})$$

This is equivalent to the expect distance from final-good demand at which a country will contribute to global value chains. The expectation is defined over a country-specific probability distribution over stages, $f_i(n) = \Pr(\Lambda_i^n) / \sum_{n'=1}^N \Pr(\Lambda_i^{n'})$.

Finally, note that for two countries with $\rho_{i'} > \rho_i$ and two inputs with $n' > n$ we necessarily have

$$\frac{f_{i'}(n') / f_{i'}(n)}{f_i(n') / f_i(n)} = \left(\frac{\rho_{i'}}{\rho_i}\right)^{2(n'-n)/N} > 1.$$

As a result, the probability functions $f_{i'}(n)$ and $f_i(n)$ satisfy the monotone likelihood ratio property in n . As is well known, this is a sufficient condition for $f_{i'}(n)$ to first-order stochastically dominate $f_i(n)$ when $\rho_{i'} > \rho_i$. But then it is immediate that $E_{f_{i'}}[n] > E_{f_i}[n]$, and thus the expected value in (A.4), which is simply $N + 1 - E_{f_i}[n]$, will be lower for country i' than for country i when $\rho_{i'} > \rho_i$.

A.5 Further Estimation Results

A.5.1 WIOD for 2014

Table A.1 presents the values of γ_j and T_j for the sample of 44 countries in the WIOD found in our benchmark estimation with $N = 2$.

A.5.2 EORA for 2013

Figure A.2 depicts some salient features of the Eora MRIO dataset for the year 2013. The figure is analogous to Figure 4 in the main text, and depicts the same qualitative patterns. The domestic shares are on average higher for final output than for inputs and there is wide dispersion in gross output to value added ratios and gross output to final output ratios, with both ratios being highly positively correlated.

Table A.1: WIOD Calibration

	γ_j	T_j		γ_j	T_j
Australia	0.93	126.75	Ireland	0.86	1.04
Austria	0.91	18.03	Italy	0.89	30.87
Belgium	0.83	2.04	Japan	0.96	23.56
Bulgaria	0.78	0.01	South Korea	0.72	1.28
Brazil	1.00	0.04	Lithuania	0.95	0.17
Canada	0.95	16.40	Luxembourg	0.51	0.17
Switzerland	0.89	1910.55	Latvia	0.79	0.07
China	0.59	0.24	Mexico	1.00	0.01
Cyprus	0.99	0.60	Malta	0.65	0.04
Czech Republic	0.73	0.30	Netherlands	0.87	2.87
Germany	0.92	139.34	Norway	0.90	0.49
Denmark	0.93	19.24	Poland	0.84	1.34
Spain	0.93	16.22	Portugal	0.95	0.55
Estonia	0.81	0.06	Romania	0.85	0.04
Finland	0.88	4.12	Russia	0.89	0.02
France	0.97	65.99	Slovakia	0.77	0.57
Great Britain	0.97	117.83	Slovenia	0.85	1.52
Greece	1.00	0.10	Sweden	0.94	67.34
Croatia	0.94	0.11	Turkey	0.91	0.14
Hungary	0.83	0.02	Taiwan	0.75	0.02
Indonesia	0.93	0.00	USA	1.00	417.44
India	0.97	0.00	Rest of World	0.79	0.02




Figure A.2: Some Key Features of the Eora MRIO Dataset

We next turn to the estimation results for the EORA 2013 database when our model is calibrated to the same moments as the WIOD and with $N = 2$. As mentioned in the main text, we find $\alpha_2 = 0.19$ in that case. Table A.2 presents the values of γ_j and T_j for the sample of 101 country/regions.

Figures A.3 and A.4 illustrate the goodness of fit of our model. As mentioned in the main text, the correlation between model and data is very high, even when considering untargeted moments.




Figure A.3: EORA Targeted Moments.

Figure A.5 presents the gains from trade with respect to autarky for the GVC and EK models

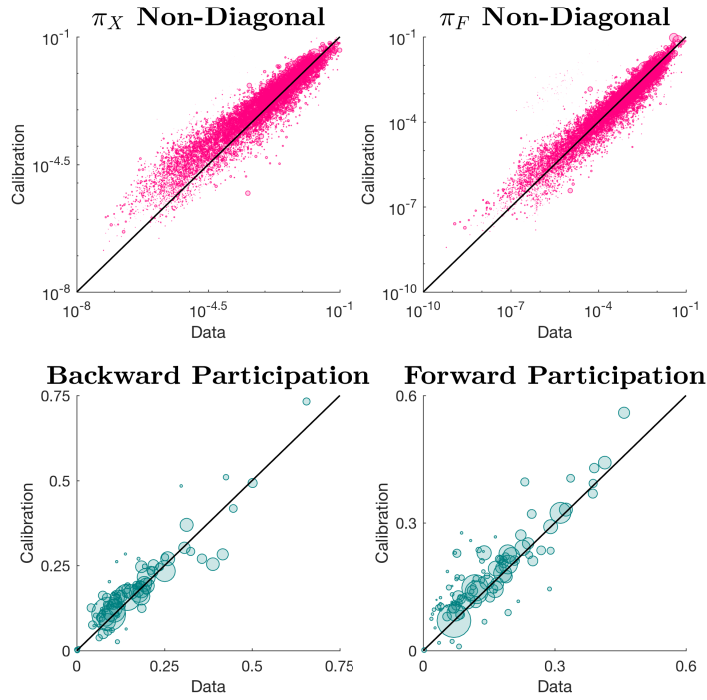


Figure A.4: EORA Untargeted Moments.

for the largest 25 countries/regions. The gains are on average 21% higher for the sample of 20 largest country-regions and 31% for the full sample of 101 country/regions.

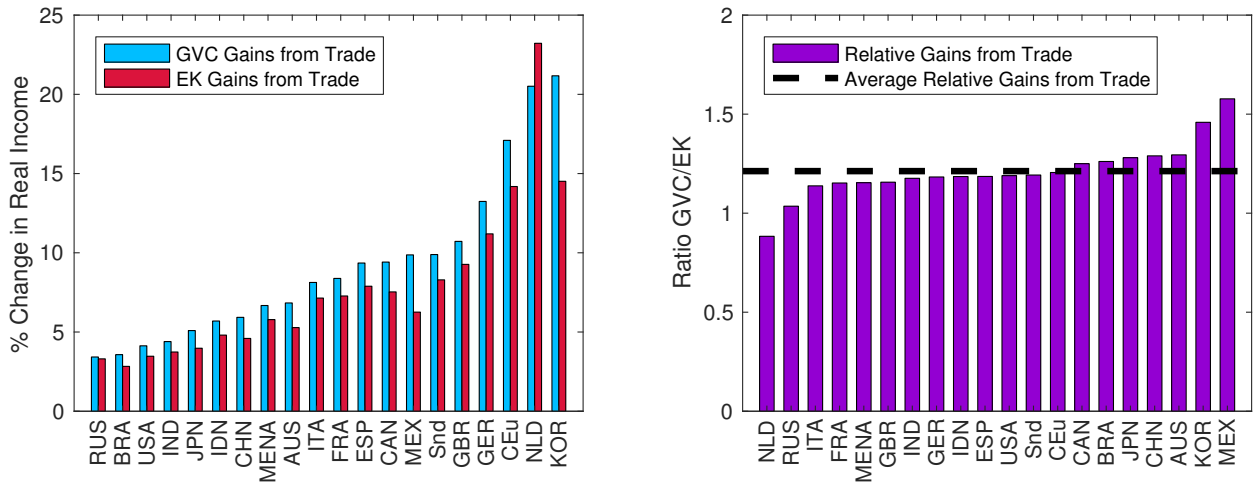


Figure A.5: EORA gains from trade relative to autarky comparison between GVC model ($N = 2$) and EK model ($N = 1$).

Table A.2: EORA Calibration

	γ_j	T_j		γ_j	T_j
Afghanistan	0.70	0.000	Israel	0.96	207.106
Eastern Europe	0.79	0.000	Italy	0.94	2.132
Algeria	0.98	0.000	Japan	0.97	19.349
Western Europe	0.72	0.029	Kazakhstan	0.86	0.001
Angola	0.98	0.015	Kenya	0.93	0.000
Latin America & Caribbean	0.85	0.347	Madagascar	0.87	0.000
Argentina	1.00	0.304	Malawi	0.61	0.000
Australia	0.85	78.774	Malaysia	0.85	0.002
Central Europe	0.89	4.964	Mali	0.64	0.001
Central Asia	0.84	0.000	Mexico	1.00	0.000
Middle East & North Africa	0.97	0.318	Morocco	0.93	0.001
Bangladesh	0.94	0.000	Mozambique	0.89	0.000
Belgium	0.77	0.034	Myanmar	0.69	1.000
Benin	0.59	0.000	Nepal	0.82	0.000
South Asia	0.68	0.007	Netherlands	0.91	0.055
Bolivia	0.88	0.096	Niger	0.97	0.000
Sub-Saharan Africa	0.92	0.000	Nigeria	1.00	0.000
Brazil	0.86	0.110	Pakistan	0.90	0.001
East Asia & Pacific	0.79	2.205	Peru	0.89	0.018
Burkina Faso	0.94	0.000	Philippines	0.99	0.000
Burundi	0.87	0.000	Poland	0.81	0.054
Cambodia	0.62	0.000	Portugal	0.85	0.115
Cameroon	0.82	0.000	South Korea	0.51	0.115
Canada	0.91	2.241	Romania	0.81	0.003
Chad	1.00	0.000	Russia	0.93	0.032
Chile	0.88	0.020	Rwanda	0.94	0.000
China	0.63	0.025	Saudi Arabia	1.00	0.019
Colombia	1.00	0.000	Senegal	0.87	0.000
Cuba	0.95	0.001	Singapore	0.68	0.111
Czech Republic	0.76	0.037	Somalia	0.39	0.000
Cote d'Ivoire	0.80	0.011	South Africa	0.88	0.003
North Korea	0.89	0.001	South Sudan	1.00	0.022
DR Congo	0.85	1.129	Spain	0.84	0.953
Scandinavia	0.89	3.265	Sri Lanka	0.98	0.000
Dominican Republic	0.93	0.001	Sudan	0.91	4.559
Ecuador	0.93	0.004	Syria	0.84	1.000
Egypt	1.00	0.000	Taiwan	0.90	1.943
Eritrea	0.81	0.000	Thailand	0.73	0.089
Ethiopia	0.02	0.000	Tunisia	0.78	0.002
France	0.94	3.155	Turkey	1.00	0.000
Germany	0.82	1.466	Uganda	0.77	0.000
Ghana	0.94	0.000	Ukraine	0.73	0.001
Greece	0.99	0.034	UK	0.90	2.642
Guatemala	0.97	0.124	Tanzania	0.46	0.000
Guinea	0.68	0.000	USA	0.94	28.721
Haiti	0.59	0.000	Uzbekistan	0.93	0.000
Hong Kong	0.54	0.007	Venezuela	0.94	0.013
India	0.87	0.001	Viet Nam	0.59	0.000
Indonesia	0.98	0.001	Yemen	0.91	0.000
Iran	0.97	0.000	Zambia	0.73	0.000
Iraq	1.00	0.000			