

ON THE GEOGRAPHY OF GLOBAL VALUE CHAINS

POL ANTRÀS

Department of Economics, Harvard University and NBER

ALONSO DE GORTARI

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 and estimate 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.

KEYWORDS: Global value chains, sequential production, specialization, gains from trade, trade costs.

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.

This paper studies how the comparative advantage of countries in specific segments of GVCs is determined in a world with barriers to international trade. The role of trade barriers on the geography of GVCs has been relatively underexplored in the literature, largely 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. The aim of this paper is to develop tools to operationalize the study

Pol Antràs: pantras@fas.harvard.edu

Alonso de Gortari: alonso.degortari@gmail.com

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 five anonymous referees for valuable input. We thank 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, Wisconsin, CREI-UPF, Autònoma in Barcelona, Toulouse, CEMFI, Toronto, HKU, HKUST, Rochester, Waseda, University of Tokyo, New Economic School, Notre Dame, Stockholm, Copenhagen, and UC Davis for useful feedback. Evgenii Fadeev, Le Kang, Daniel Ramos, BooKang Seol, Maria Voronina, and Diana Zhu provided outstanding research assistance at various stages of this paper's production chain. Antràs acknowledges support from the NSF (proposal #1628852). de Gortari thanks Harvard University, Princeton University, and Dartmouth College, for their support while writing this paper. All errors are our own.

of the geography of GVCs in both partial-equilibrium and general-equilibrium environments.

We start off our analysis in Section 2 by developing a simple partial-equilibrium framework of multi-stage production in which a *lead firm* optimally chooses 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 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 simple Ricardian model of trade in which the combination of labor productivity and trade cost differences across countries shapes the equilibrium position of countries in GVCs. More specifically, we adapt Eaton and Kortum’s (2002) Ricardian model of trade 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), Fally and Hillberry (2018), Johnson and Moxnes (2019)) have focused on the quantification of relatively low-dimensional mod-

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 Braccaccio, Kalouptsi, and Papageorgiou (2017), search frictions in the shipping industry allow shipping companies to extract rents from exporters by charging shipping fees that are increasing in the value of the goods in transit.

els with two stages or a small number of countries. Indeed, as we describe in Section 3, it is not obvious how to exploit the extreme-value distribution results developed 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 three alternative approaches to restore the tractability of the Eaton and Kortum (2002) framework in a Ricardian model with multi-stage production. The first approach consists in simply treating the *overall* (i.e., chain-level) 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.² 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). Finally, we develop an alternative decentralized approach inspired by the work of Oberfield (2018), in which technology is again specified at the stage level (rather than at the chain level), but in which productivity is buyer-seller-specific. By appropriate choice of functional forms, we follow Oberfield (2018) in showing that this formulation can also deliver a Fréchet distribution of productivity *at the chain level*.

Interestingly, we find that these three alternative approaches are isomorphic, in the sense that they yield the exact same equilibrium equations. More specifically, regardless of the microfoundation one chooses to invoke, 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. 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, Costinot, and Rodríguez-Clare (2012). Although our 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). We close Section 4 by providing suggestive empirical evidence for this centrality-downstreamness nexus and for a key mechanism of the model—namely, the fact that the elasticity of trade flows to distance is larger for downstream stages than for upstream stages.

²A recent paper by Allen and Arkolakis (2019) adopts a similar path-specific representation of productivity in a very different setting.

In Section 5, we generalize our framework along several dimensions, which permits our model to nest and better compare to various other Ricardian models of trade. More specifically, we introduce multiple industries and rich input-output links across these industries, and we allow for technologies that depend on whether an industry's output is used as final consumption or as an intermediate input for different industries. Exploiting properties of the distribution of final-good and input prices produced by the model, we show that the various versions of our model deliver 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 World 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.

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, when aggregated at the country level.³ 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 bilateral 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 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 all the entries of a WIOT. We perform this exercise using 2014 data from the World Input–Output Database, a source which is deemed to provide relatively high-quality and 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. We find that the model is able to match the data very well.

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 trace 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, on average, 60% 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—akin to the work of [Alexander \(2017\)](#)—estimated to match all the entries of the WIOT. Similarly, we find that the real income losses from a hypothetical U.S.-China trade war are, for those two countries, almost 50% higher with sequential production than without it, though such a war may well benefit third countries.

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, Vogel, and Wang \(2013\)](#), [Antràs and Chor \(2013\)](#), [Kikuchi, Nishimura, and Stachurski \(2018\)](#), [Fally and Hillberry \(2018\)](#), and [Tyazhelnikov \(2017\)](#).⁴ A key limitation of this body of theoretical work is that it either completely ab-

³In Section 7.4, we estimate a two-sector version of our model. Estimating our full multi-industry model is not feasible with current computational constraints, as explained in Section 6 (especially, footnote 34).

⁴This 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 \(2004\)](#).

stracts from modeling trade costs or it introduces such barriers in highly stylized ways, for instance assuming common trade costs across all country-pairs as in Baldwin and Venables (2013) or Section 6.1 in Costinot, Vogel, and Wang (2013). On the empirical front, a growing body of work, starting with the seminal work of Johnson and Noguera (2012a), 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, Wang, and Wei (2014), Johnson (2014), Timmer, Erumban, Los, Stehrer, and de Vries (2014), de Gortari (2019)).⁵ A parallel empirical literature has developed indices of the relative positioning of industries and countries in GVCs (see Fally (2012), Antràs, Chor, Fally, and Hillberry (2012), Alfaro, Antràs, Conconi, and Chor (2019)). There is also a prior body of work estimating the differential sensitivity of input trade and final-good trade to trade barriers, as exemplified by Bergstrand and Egger (2010) and Baldwin and Taglioni (2014), among others. On the quantitative side, and as mentioned above, our work builds on and expands on previous work by Yi (2003, 2010), Fally and Hillberry (2018), and Johnson and Moxnes (2019). Other authors, including Caliendo and Parro (2015), Alexander (2017), Antràs and Chor (2018), and Baqaee and Farhi (2019), have developed quantitative frameworks with Input–Output linkages across countries, but in models with a roundabout production structure without an explicit sequentiality of production. The connection between our framework and these previous contributions has been further explored in de Gortari (2019), who blended several strands of this literature by generalizing the formulas on value-added content and downstreamness within the context of a multi-sector Armington model with sequential production. Finally, some implications of the rise of offshoring and GVCs for trade policy have been studied by Antràs and Staiger (2012) and Blanchard, Bown, and Johnson (2017), 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 and provides suggestive empirical evidence for some of the key features of our model. Section 5 develops several extensions of our framework. 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 Supplemental Material (Antràs and de Gortari (2020)) (Appendix A) and to the unpublished Appendix B (downloadable from the journal's webpage as part of the replication file).

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 following a strictly sequential process. We defer a discussion of more general processes and of the general-equilibrium aspects of the model to Sections 3 and 5.

⁵An important precursor to this literature is Hummels, Ishii, and Yi (2001), who combined international trade and Input–Output data to construct indices of vertical specialization.

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 final 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 factor (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 factor. The cost of the composite factor varies across countries and is denoted by c_i in country i .⁶ 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 $p_{\ell(n)}^n \tau_{\ell(n)\ell(n+1)}$.

Although the main results of this section extend to more general specifications of technology, in the main text we will restrict the analysis to the following functional form for the sequential cost function associated with a path of production $\ell = \{\ell(1), \ell(2), \dots, \ell(N)\}$:

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

where $\alpha_n \in (0, 1)$ denotes the cost share of the composite factor at stage n and $a_{\ell(n)}^n$ is the unit factor requirement at stage n in country $\ell(n)$. Because the initial stage of production uses solely the local composite factor, we have $\alpha_1 = 1$. Notice that the cost function in (1) is *Ricardian* in nature, in the sense that cross-country differences in technology are associated with differences in the efficiency with which the local composite factor is used in different stages. The choice of this specification will permit a particularly sharp characterization of the key results of this section, and is important for tractability in the general-equilibrium version of the model in Section 3.

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

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

⁶For now we take this cost as given, but in the general-equilibrium analysis in Section 3, we will break c_i into the cost of labor and of a bundle of intermediate inputs we call *materials*. This will allow our model to encompass previous Ricardian models—and most notably Eaton and Kortum (2002)—featuring roundabout production. On the downside, this assumption precludes a study of the effects of GVC integration on wage inequality, as in the work of Basco and Mestieri (2019) or Lee and Yi (2018).

country j . Using $p_j^F(\ell) = p_{\ell(N)}^N(\ell)\tau_{\ell(N)j}$ and iterating (1), this problem reduces to

$$\ell^j = \arg \min_{\ell \in \mathcal{J}^N} p_j^F(\ell) = \arg \min_{\ell \in \mathcal{J}^N} \left\{ \prod_{n=1}^N (a_{\ell(n)}^n c_{\ell(n)})^{\alpha_n \beta_n} \times \prod_{n=1}^{N-1} (\tau_{\ell(n)\ell(n+1)})^{\beta_n} \times \tau_{\ell(N)j} \right\}, \quad (2)$$

where

$$\beta_n \equiv \prod_{m=n+1}^N (1 - \alpha_m), \quad (3)$$

and where we use the convention $\prod_{m=N+1}^N (1 - \alpha_m) = 1$. Thus, α_n is the share of the stage- n composite factor in stage- n production, while $\alpha_n \beta_n$ is the share of the stage- n composite factor in the production of the finished good (i.e., its share in the global value chain). Note that $\sum_{n=1}^N \alpha_n \beta_n = 1$.

It is worth highlighting two important features of program (2). First, notice that when trade costs are identical for all country-pairs (i.e., $\tau_{ij} = \tau$ for all i and j), the last 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 (2) 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 (2) can be dramatically reduced using dynamic programming.

A second noteworthy aspect of the minimand in equation (2) 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 (3) that, because $\alpha_n > 0$ for all n , we have $\beta_1 < \beta_2 < \dots < \beta_N = 1$. The reason for this compounding effect of trade costs stems from the fact that the costs of transporting goods have been modeled (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.⁷

An implication of this compounding effect is that, in choosing their optimal path of production, firms will be more concerned about reducing trade costs in relatively downstream stages than in relatively upstream stages. As we will demonstrate when exploring the general equilibrium of our model, this feature of the cost function will tend to generate a centrality-downstreamness nexus by which, *ceteris paribus*, relatively more central

⁷For 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 trade-cost elasticity equals $\beta_n = n/N$ and thus increases linearly with the downstreamness n of a stage.

countries will tend to gain comparative advantage and specialize in relatively downstream stages. At this point, however, this result may not appear to be entirely trivial because if α_n were to rise sufficiently fast with n , one could envision firms being more concerned with reducing the cost of the composite factor than with reducing trade costs for the most downstream stages. This would work against the centrality-downstreamness nexus if, plausibly, central locations tend to feature a higher price c_i for their local composite factor.

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.1.1 that the same result applies for arbitrary constant-returns-to-scale technologies. 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 more concerned about minimizing trade costs in downstream stages than in upstream stages is quite general.⁸ We summarize this result as follows:

PROPOSITION 1: *The elasticity of the overall cost of production with respect to trade costs incurred at stage n increases along the chain (i.e., increases with n).*

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 worldwide producers of each stage n . 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 composite factor 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 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 $n - 1$ from the location $\ell(n - 1) \in \mathcal{J}$ that solves

$$p_{\ell(n)}^n = \min_{\ell(n-1) \in \mathcal{J}} \left\{ (a_{\ell(n)}^n c_{\ell(n)})^{\alpha_n} (p_{\ell(n-1)}^{n-1} \tau_{\ell(n-1)\ell(n)})^{1-\alpha_n} \right\}, \tag{4}$$

where $p_{\ell(n-1)}^{n-1}$ is the minimum (free-on-board) price charged by producers of stage $n - 1$ in country $\ell(n - 1)$. Importantly, notice that the problem in (4) amounts to minimizing sourcing costs $p_{\ell(n-1)}^{n-1} \tau_{\ell(n-1)\ell(n)}$ regardless of the composite factor cost $c_{\ell(n)}$ and productivity $a_{\ell(n)}^n$, and of the future path of the good after flowing through $\ell(n)$ at stage n . Producers of the initial stage only use their local composite factor and thus $p_{\ell(1)}^1 = a_{\ell(1)}^1 c_{\ell(1)}$. In sum, in this decentralization, we have $J \times N$ agents computing the cost of J possible sources of goods and picking the minimum one.

⁸For 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.

Once we have solved these $J \times N$ simple problems, it is then straightforward to use forward induction to solve for the resulting equilibrium path $\ell^j = \{\ell^j(1), \ell^j(2), \dots, \ell^j(N)\} \in \mathcal{J}^N$ for each destination market $j \in \mathcal{J}$. To see this, begin by computing $p_{\ell(1)}^1 = a_{\ell(1)}^1 c_{\ell(1)}$ for each possible initial location of stage $n = 1$. Then, for each of the J potential locations of stage $n = 2$, we can use (4) to solve for corresponding optimal stage-1 location and associated cost up to stage $n = 2$. Importantly, the same solution applies regardless of the final destination of consumption j . We can then use these J resulting costs (up to stage $n = 2$) to proceed to stage $n = 3$ and solve for the optimal location of stage $n = 2$ for each potential location of $n = 3$. We can then move to $n = 4$, to $n = 5$, and so on, until reaching final consumption, which delivers the optimal location of final assembly $\ell^j(N)$ for each destination market j . More formally, if we 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 —with a particular sequence given by $\ell_n^j = \{\ell^j(1), \dots, \ell^j(n)\}$ —this sequence can be found recursively for all $n = 2, \dots, N$ by simply solving

$$\ell^j(n) = \arg \min_{k \in \mathcal{J}} \{p_k^n(\ell_{n-1}^k) \tau_{kj}\},$$

with $\ell_n^j = \{\ell_{n-1}^j, \ell^j(n)\}$. For $n = 1$, we have $\ell_0^k = \emptyset$ for all $k \in \mathcal{J}$ and the price depends only the composite factor cost: $p_k^1(\emptyset) = a_k^1 c_k$.

A natural question is then how the solution to this decentralization problem relates to that obtained from the *lead-firm* problem in equation (2). It turns out that as long as technology features constant returns to scale (as in our particular Cobb–Douglas formulation in (1)), these two solutions will in fact coincide (see Appendix A.1.2). Intuitively, with constant returns to scale, the identity of the specific firms making these decisions along the chain is immaterial, so the recursive formulation of the problem in (4) 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 (2) in a brute force manner, the *lead firm* breaks the problem into a series of stage- and country-specific optimal sourcing problems, and then solves the problem via forward induction (starting in the most upstream stage). Equation (4) constitutes the Bellman equation associated with this problem, and invoking the principle of optimality, we can then establish that the resulting optimal path of production $\ell^j = \ell_N^j = \{\ell^j(1), \ell^j(2), \dots, \ell^j(N)\} \in \mathcal{J}^N$ that minimizes the cost $p_j^f(\ell)$ in this decentralized formulation of the problem will coincide with the one obtained solving the *lead-firm* problem in (2) by exhaustive search.

A key advantage of this decentralized or 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 (see Tyazhelnikov (2017) for a contemporaneous derivation of this result). For example, with 200 countries and 5 stages, this amounts to only 200,000 computations rather than 64 trillion computations.⁹ Although it might be clear from our discussion above, it is worth stressing that the isomorphism between the lead-firm problem and the decentral-

⁹Though the dimensionality of the lead-firm’s problem is huge, for the particular case with Cobb–Douglas technologies, in Appendix A.1.2 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>).

ized problem holds true for *any* constant-returns-to-scale technology, and not only for the Cobb–Douglas one in (1).¹⁰ We summarize this result as follows:

PROPOSITION 2: *The solution ℓ^j to the lead-firm cost-minimization problem in (2) coincides with the equilibrium path of a decentralized production chain in which $J \times N$ agents solve the stage cost-minimization problem in (4).*

Before turning to the general equilibrium of our framework, it is important to stress that Proposition 2 relies crucially on the assumption that technology features constant returns to scale. With variable returns to scale, matters would become significantly more complicated. First, the *lead-firm* problem could not be solved independently for each destination market j , because whether a location ℓ constitutes a cost-minimizing location for stage n in a particular chain ending in j will be a function of the scale of this production node, and the latter is shaped by the overall level of production flowing through this node (potentially involving chains ending in destination markets other than j). As a result, dynamic programming ceases to be a powerful tool to simplify the problem (see de Gortari (2020)). Second, and relatedly, in a decentralized equilibrium, agents will fail to internalize the effect of their sourcing decisions on marginal costs (via scale effects), and the equilibrium paths might deviate from the ones the *lead firm* would choose.

3. GENERAL-EQUILIBRIUM MODEL

We next embed our model of firm behavior into a general-equilibrium model of trade along the lines of the multi-country Ricardian model of trade of Eaton and Kortum (2002).

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

$$u(\{y_i^N(z)\}_{z=0}^1) = \left(\int_0^1 (y_i^N(z))^{(\sigma-1)/\sigma} dz \right)^{\sigma/(\sigma-1)}, \quad \sigma > 1. \tag{5}$$

Production of each of the final-good varieties is as described in the previous section: all production processes entail N sequential stages (indexed by n) and are characterized by the Ricardian, Cobb–Douglas specification in (1). We let countries differ in three key aspects: (i) their technological efficiency, as determined by the unit composite factor 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). In Section 5, we will demonstrate that our framework remains tractable even when relaxing many of these assumptions.

The local composite factor used at each stage comprises labor and an aggregator of final-good varieties that corresponds exactly to the CES aggregator in (5). This amounts

¹⁰In Web Appendix B.1, we illustrate some of the salient and distinctive features of our partial model via a simple example featuring four countries ($J = 4$) and four stages ($N = 4$).

to assuming that part of final-good production is not absorbed by consumers, but rather by firms that use those goods as a bundle of materials. More specifically, we let the cost c_i of the composite factor in country i be captured by a Cobb–Douglas aggregator $c_i = (w_i)^\gamma (P_i)^{1-\gamma}$, where P_i is the ideal price index associated with the CES aggregator in (5). This roundabout structure of production is standard in recent quantitative Ricardian models (see Eaton and Kortum (2002), Alvarez and Lucas (2007), or Caliendo and Parro (2015)), and constitutes a particularly convenient way for these models to introduce intermediate input trade flows and a gross-output to value-added ratios higher than 1. Although our model features intermediate-input flows across countries even in the absence of these production “loops,” we still adopt this formulation for comparability (especially when evaluating our model quantitatively in Section 7).¹¹

This completes the discussion of the structure of our general-equilibrium model. In principle, given values for the unit composite factor 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 computationally feasible and transparent manner. With that in mind, we next explore a particularly convenient parameterization of the unit factor 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 factor 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 (2019), for instance, considered 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 in our Cobb–Douglas cost function in (2) is *not* distributed Fréchet.¹² As a result, these papers need to resort to numerical methods to approximate the solution of their models, even when restricting the analysis to two-stage chains. We instead develop three alternative approaches that all permit a sharp and exact characterization of some of the features of the equilibrium for an arbitrary number of stages, and that will be amenable to structural, generalized method of moments estimation using World Input–Output Tables.

¹¹In recent work, Fally and Sayre (2018) and Baqaee and Farhi (2019) have argued that the elasticity of substitution between value added and materials is likely to be lower than 1, especially in the short run. For the same reason, one could argue that our Cobb–Douglas sequential production process in (1) features too much substitution across stages. Nevertheless, like most trade models, ours is a model of the long run, so it is much less clear to us that a Cobb–Douglas assumption biases our quantitative estimates in any particular direction.

¹²Assuming a Leontief cost function (i.e., perfect complementarity) does not provide tractability either because the sum of Fréchet random variables is not distributed Fréchet either.

3.2.1. Lead-Firm Approach With Chain-Level Productivity

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 requirements $a_i^n(z)$ for each stage $n \in \mathcal{N}$ and each country $i \in \mathcal{J}$. The key conceptual innovation we propose, relative to Eaton and Kortum (2002) and previous multi-stage extensions of that framework, is to think about productivity at the *chain level* rather than at the stage level. More specifically, a given production path $\ell = \{\ell(1), \ell(2), \dots, \ell(N)\} \in \mathcal{J}^N$ will be associated with a *chain-level* production cost that is naturally a function of trade costs, composite factor costs, and the state of technology of the various countries involved in the chain. Yet, two chains flowing across the same countries in the exact same order may not achieve the same overall productivity due to (unmodeled) idiosyncratic factors, such as compatibility problems, production delays, or simple mistakes. More formally, and building on the cost function in (2), we assume that the overall productivity of a given chain ℓ is characterized by

$$\Pr\left(\prod_{n=1}^N (a_{\ell(n)}^n(z))^{\alpha_n \beta_n} \geq a\right) = \exp\left\{-a^\theta \prod_{n=1}^N (T_{\ell(n)})^{\alpha_n \beta_n}\right\}, \tag{6}$$

which amounts to assuming that $\prod_{n=1}^N (a_{\ell(n)}^n(z))^{-\alpha_n \beta_n}$ 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 $p_j^F(\ell, z)$ the price paid by consumers in j for a good z produced following the path ℓ , we have

$$\begin{aligned} &\Pr(p_j^F(\ell, z) \geq p) \\ &= \exp\left\{-p^\theta \times \prod_{n=1}^N ((c_{\ell(n)})^{-\theta} T_{\ell(n)})^{\alpha_n \beta_n} \times \prod_{n=1}^{N-1} (\tau_{\ell(n)\ell(n+1)})^{-\theta \beta_n} \times (\tau_{\ell(N)j})^{-\theta}\right\}, \end{aligned} \tag{7}$$

independently of the final good z under consideration. This result will be key for neatly characterizing the equilibrium, as we will show in Section 4.

3.2.2. Decentralized Approaches With Stage-Level Productivity

Although we adopt the “Fréchet-in-the-chain” specification in (6) for tractability, we next outline two alternative approaches that deliver an identical set of equilibrium conditions to those we will derive under the specification in (6). In both cases, we explore an environment akin to the decentralized equilibrium developed in Section 2.3, in which stage-specific producers at each stage $n \in \mathcal{N}$ simply attempt to minimize the sourcing cost $p_{\ell(n-1)}^{n-1}$ of the good completed up to the prior stage $n - 1$. To save space, we relegate the technical details of these alternative approaches to Appendix A.2.

The first approach is close in spirit to the stage-specific productivity randomness in Yi (2010) and Johnson and Moxnes (2019), and assumes that $1/a_{\ell(n)}^n(z)$ is drawn independently (across goods and stages) from a Fréchet distribution. If stage-specific producers were to decide on their source of inputs with precise information on the alternative sourc-

ing costs $p_{\ell(n-1)}^{n-1}$ available to them, the isomorphism result in Proposition 2 would apply, and the resulting distribution of final-good prices $p_j^F(\ell, z)$ would be a function of the product of Fréchet distributions, and thus hard to characterize.¹³ In order to make this alternative approach tractable, we assume that producers at each stage n and location $\ell(n)$ do not observe realized upstream prices $p_{\ell(n-1)}^{n-1}$ 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. Yet, ex post, a supplier's marginal cost might be higher or lower than expected because this supplier may face unexpectedly high or low sourcing costs itself.¹⁴

Although the pre-commitment to buy from a particular source naturally affects the nature of ex post competition, we assume that buyers have all the bargaining power and continue to be able to source upstream inputs at marginal cost. As we show in the Appendix, in this decentralized economy with incomplete information, the equilibrium price $p_j^F(\ell, z)$ paid by consumers in j for a good z produced following the path ℓ is Fréchet distributed. Furthermore, by appropriate choice of the shape and scale parameter of the stage-specific distributions of $1/a_{\ell(n)}^n(z)$, this alternative formulation delivers an expression for the equilibrium distribution of prices $p_j^F(\ell, z)$ that is identical to that obtained in equation (7) in our *lead-firm* approach.

The second decentralized approach is inspired by recent work of Oberfield (2018).¹⁵ This approach again begins by specifying a distribution of productivity $1/a_{\ell(n)}^n(z)$ at the stage level, but assumes that this productivity is now buyer-seller-specific (or *match-specific*) and characterized by a Pareto distribution with shape parameter θ . Furthermore, each producer (or buyer) chooses the best match among a pool of potential suppliers at $n - 1$, with the number of available potential suppliers in each source country k being characterized by a Poisson distribution with a mean related to the parameter $T_{\ell(n)}$. For the case of a closed economy with an infinite loop of stages (i.e., when stage-1 producers buy inputs from other producers), Oberfield (2018) showed that this formulation delivers a Fréchet distribution for chain-level productivity. In Appendix A.2, we show that his result generalizes to our multi-country environment with an arbitrary finite number of stages, provided that the distribution of productivity in the initial stage ($n = 1$) is Fréchet distributed. Furthermore, as long as all transactions are priced at marginal cost, by appropriate choice of the parameters governing the Pareto and Poisson distributions at $n > 1$ and the Fréchet distribution at $n = 1$, this alternative decentralized approach again delivers a distribution for the final-good prices $p_j^F(\ell, z)$ that is Fréchet and given by the same expression (7) as in our *lead-firm* approach (see Appendix A.2 for details).

¹³By Proposition 2, tractability would evidently be restored if one assumed the chain-level technology in (6). The goal of this section, however, is to explore variants of our model in which technology is specified at the stage level.

¹⁴The 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 (2017) and Antràs, Fort, and Tintelnot (2017).

¹⁵We are grateful to Sam Kortum for suggesting this alternative approach to us.

4. CHARACTERIZATION OF THE EQUILIBRIUM

Having presented and discussed our assumptions on technology, in this section we characterize the general equilibrium of our model. We proceed in five 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, and for the aggregate price index in each country. Second, we study the existence and uniqueness of the general equilibrium. Third, we derive an expression for the gains from trade in our model and compare it to the one obtained by Eaton and Kortum (2002). Fourth, we formalize the link between downstreamness and centrality that we hinted at in Section 2. Finally, we present some suggestive evidence for this relationship and for a key mechanism in the model. Throughout this section, we focus on deriving our results for the “Fréchet-in-the-chain” specification in (6), with the understanding (formalized in Appendix A.2) that our two alternative decentralized approaches lead to the exact same equilibrium conditions.¹⁶

4.1. Relative Prevalence of Different GVCs and Equilibrium Prices

Consider 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 (7). In such a case, we can readily invoke some key 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

$$\pi_{\ell j} = \frac{\prod_{n=1}^{N-1} ((T_{\ell(n)})^{\alpha_n} ((c_{\ell(n)})^{\alpha_n} \tau_{\ell(n)\ell(n+1)})^{-\theta})^{\beta_n} \times (T_{\ell(N)})^{\alpha_N} ((c_{\ell(N)})^{\alpha_N} \tau_{\ell(N)j})^{-\theta}}{\Theta_j}, \tag{8}$$

where

$$\Theta_j = \sum_{\ell \in \mathcal{J}^N} \prod_{n=1}^{N-1} ((T_{\ell(n)})^{\alpha_n} ((c_{\ell(n)})^{\alpha_n} \tau_{\ell(n)\ell(n+1)})^{-\theta})^{\beta_n} \times (T_{\ell(N)})^{\alpha_N} ((c_{\ell(N)})^{\alpha_N} \tau_{\ell(N)j})^{-\theta}, \tag{9}$$

and where remember that $c_i = (w_i)^\gamma (P_i)^{1-\gamma}$. 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.¹⁷

Second, and as already anticipated in equation (7), the distribution of final-good prices $p_j^F(\ell, z)$ paid by consumers in j satisfies $\Pr(p_j^F(\ell, z) \leq p) = 1 - \exp\{-\Theta_j p^\theta\}$ and is thus independent of the path of production ℓ . Consequently, the probabilities $\pi_{\ell j}$ in equation (8) above also constitute the shares of country j 's income spent on final goods produced under all possible paths $\ell \in \mathcal{J}^N$.

¹⁶As pointed out by a referee, it is also possible to recast our one-sector model into a multi-industry model with input-output links à la Caliendo and Parro (2015), in which different stages are treated as separate industries. Nevertheless, because the output of these stage-specific industries is not observed in the data, this isomorphism is of limited use in estimating our model and performing counterfactuals.

¹⁷Note that when $N = 1$, we necessarily have $\alpha_N = 1$, and the formulas (8) and (9) collapse to the well-know trade share formulas in Eaton and Kortum (2002).

As is clear from equation (8), GVCs that involve countries with higher states of technology T_i or lower composite factor costs c_i will tend to feature disproportionately in production paths leading to consumption in j . Furthermore, and consistently with Proposition 1, 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 (5):

$$P_j = \kappa(\Theta_j)^{-1/\theta}, \tag{10}$$

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 just described how to adapt the Eaton and Kortum (2002) toolbox to derive an explicit formula for the relative prevalence of various production paths (or GVCs) in delivering consumption goods to a given country. We will next demonstrate that this characterization is sufficient to derive interesting implications for the gains from trade (and the costs of trade wars) from our framework, and also to formalize a link between centrality and GVC position. Nevertheless, it is important to emphasize that these “GVC trade shares” are not observable in the data, thus creating a challenge for the estimation of the model’s parameters. In Section 5, we will demonstrate that it is relatively straightforward to map objects derived from the model (e.g., final-good and intermediate input flows across countries) to the type of information available in World Input–Output Tables.

4.2. General Equilibrium

We have thus far characterized the relative prevalence of different GVCs as a function of the vectors of equilibrium wages $\mathbf{w} = (w_1, \dots, w_j)$ and input bundle costs $\mathbf{P} = (P_1, \dots, P_j)$ in all countries. We next describe how these vectors are pinned down in general equilibrium.

Notice first that, invoking (10) and $c_i = (w_i)^\gamma(P_i)^{1-\gamma}$, we can solve for the vector \mathbf{P} as a function of the vector \mathbf{w} from the system of equations:

$$P_j = \kappa \left(\sum_{\ell \in \mathcal{J}^N} \prod_{n=1}^{N-1} ((T_{\ell(n)})^{\alpha_n} ((c_{\ell(n)})^{\alpha_n} \tau_{\ell(n)\ell(n+1)})^{-\theta})^{\beta_n} \times (T_{\ell(N)})^{\alpha_N} ((c_{\ell(N)})^{\alpha_N} \tau_{\ell(N)j})^{-\theta} \right)^{-1/\theta}, \tag{11}$$

for all $j \in \mathcal{J}$.

To solve for equilibrium wages, notice that for all GVCs, stage- n value added (or labor income) accounts for a share $\gamma\alpha_n\beta_n$ 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_jL_j) and spending in the intermediate-input bundle ($w_jL_j \times (1 - \gamma)/\gamma$). The share of that spending by j 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 $\Lambda_i^n = \{\ell \in \mathcal{J}^N \mid \ell(n) = i\}$ and $\pi_{\ell j}$ is given in equation (8). It thus follows that the equilibrium wage vector is determined by the solution of the following system of equations:

$$w_iL_i = \sum_{j \in \mathcal{J}} \sum_{n \in \mathcal{N}} \alpha_n\beta_n \times \Pr(\Lambda_i^n, j) \times w_jL_j. \tag{12}$$

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 $\Pr(\Lambda_i^n, j) = \pi_{ij} = (\tau_{ij} c_i)^{-\theta} T_i / \sum_k (\tau_{kj} c_k)^{-\theta} T_k$. The equilibrium then reduces to the general equilibrium in [Eaton and Kortum \(2002\)](#) and [Alvarez and Lucas \(2007\)](#).

In [Web 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 (11) delivers a unique vector of input bundle costs \mathbf{P} even when $N > 1$. In that appendix, we also demonstrate the existence of a solution $\mathbf{w}^* \in \mathbb{R}_{++}^J$ to the system of equations in (12)—with (11) plugged in—and we derive a set of sufficient conditions that ensure that this solution is unique.

4.3. 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 j to serve consumers in the same country j . Let us denote this domestic chain by $\mathbf{j} = (j, j, \dots, j)$. From equation (8), such a value chain would capture a share of country j ’s spending equal to

$$\pi_{jj} = \frac{(\tau_{jj})^{-\theta \sum_{n=1}^N \beta_n} \times (c_j)^{-\theta} T_j}{\Theta_j},$$

where we have used the fact that $\sum_{n=1}^N \alpha_n \beta_n = 1$. Combining this equation with (10) and $c_j = (w_j)^\gamma (P_j)^{1-\gamma}$, we can express real income in country j as

$$\frac{w_j}{P_j} = (\kappa (\tau_{jj})^{\sum_{n=1}^N \beta_n})^{-1/\gamma} \left(\frac{T_j}{\pi_{jj}} \right)^{1/(\theta\gamma)}.$$

Consider now a change in trade costs, while holding all other technological parameters ($\alpha_n, \beta_n, T_j, \gamma, \theta, \sigma$) and domestic trade costs τ_{jj} unchanged. Letting “hats” denote gross changes in variables (i.e., $\hat{x} = (x + dx)/x$), the change in welfare following any change to the matrix of trade barriers equals

$$\hat{W}_j = (\hat{\pi}_{jj})^{-1/(\theta\gamma)}. \tag{13}$$

This formula is particularly useful when considering a prohibitive increase in trade barriers. Because under autarky $\pi_{jj} = 1$, we can conclude that the (percentage) real income losses of going to autarky, relative to an initial trade equilibrium, are given by $\hat{G}_j = 1 - (\pi_{jj})^{1/(\theta\gamma)}$. 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, Costinot, and Rodríguez-Clare \(2012\)](#)). An important difference, however, is that π_{jj} is *not* the aggregate share of spending on domestic intermediate or final goods (which we denote Π_{jj}^X and Π_{jj}^F , respectively, and which are readily available in input-output data sets), but rather the share of spending on goods that are produced *entirely* through domestic supply chains. That is, unlike Π_{jj}^X and Π_{jj}^F , π_{jj} cannot be directly observed in the data, and as a result, the sufficient statistic approach advocated by [Arkolakis, Costinot, and Rodríguez-Clare \(2012\)](#) is not feasible in our setting. Instead, one needs a model to structurally back out π_{jj} from available data. For a similar reason, the hat algebra approach to counterfactual analysis proposed by [Dekle, Eaton, and Kortum \(2008\)](#) is not feasible in our setting either.

Importantly, the share π_{jj} is necessarily lower than Π_{jj}^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 model, with stages in our model playing the role of sectors in his framework. We will elaborate on this comparison below. Before doing so, however, it should be noted that the values of γ 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 , 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 invoke estimates from standard gravity equation specifications to back out that parameter, an issue we will return to in Section 6. Overall, whether our model generates larger or smaller gains from trade than models without multi-stage production is an empirical question, and one which we will explore in Section 7.

Having said this, we next briefly study two special parameterizations of our model that might help better understand the implications of sequential production for the gains from trade. In our first exercise, we focus on the impact of heterogeneity in comparative advantage across stages of the supply chain. Specifically, suppose that the state of technology T_j of countries is now stage-specific and denoted by T_j^n , and countries only vary along this dimension. Although our baseline model does not feature such heterogeneity, it is evident that equation (13) would continue to apply in such a case. Our first result (proved in Appendix A.1.5) is then the following:

PROPOSITION 3: *If $\alpha_n \beta_n = 1/N$ for all n , $\frac{1}{N} \sum_{n \in \mathcal{N}} (T_j^n)^{1/N} = \frac{1}{J} \sum_{j \in \mathcal{J}} (T_j^n)^{1/N} = \bar{T}$ for all j and n , and all other parameters are symmetric across countries, then starting from free trade, the losses of reverting to autarky ($\hat{G}_j = 1 - \hat{W}_j$) are given by*

$$\hat{G}_j = 1 - \left(\frac{1}{J} \times \frac{\text{GeometricMean}_n[(T_j^n)^{1/N}]}{\text{ArithmeticMean}_n[(T_j^n)^{1/N}]} \right)^{N/\gamma\theta} .$$

This proposition illustrates that higher dispersion in productivity across stages of production leads to higher gains from trade because countries gain more from specializing along specific segments of the value chain. To see this, note that a mean preserving spread of $(T_j^n)^{1/N}$ leaves the arithmetic mean of these parameters constant (by construction) but decreases their geometric mean. While this result is not entirely novel—similar insights are found in multi-sector models with cross-industry heterogeneity in productivity (Ossa (2015))—this precise characterization in terms of arithmetic and geometric means is new, and it showcases how international trade enables countries to leverage their comparative advantage in specific segments of the value chain *within* a given industry.

In our second parameterization, we focus on the interplay between geography and production fragmentation. In particular, we find the following (see Appendix A.1.6 for a proof):

PROPOSITION 4: *If all countries are symmetric in all respects, and $\tau_{ij} = \tau$ for $i \neq j$ and $\tau_{ij} = 1$ for $i = j$, then the losses of reverting to autarky are given by*

$$\hat{G}_j = 1 - \left(\prod_{n=1}^N (1 + (J - 1)\tau^{-\theta\beta_n}) \right)^{-1/\gamma\theta}.$$

This result is useful for showing that our model echoes Melitz and Redding’s (2014) result that the gains from trade may become unboundedly large as production is sliced into more and more stages of production, that is, $\lim_{N \rightarrow \infty} \hat{G}_j = 1$. There are, however, two important differences relative to the result in Melitz and Redding (2014). First, while they assumed that all value is added at the most upstream stage and that each subsequent production stage entails only re-shipping the upstream input to the next stage without adding further value, our model lets us generalize this insight to richer production settings, with arbitrary value-added shares along the value chain (as long as $\alpha_n < 1$ for all $n > 1$). Second, under the assumptions in Melitz and Redding (2014), the value-added to gross output ratio is given by $\text{GDP/GO} = 1/N$ and thus goes to zero as $N \rightarrow \infty$. In our model, this need not be the case. For example, let $\Psi \in (0, 1)$ be some target GDP to gross output ratio and define $\alpha_n = (n - 1)/N$ for $n > 1$, so $\beta_n = (N - n)!/N^{N-n}$. Because $\sum_{n=1}^N \beta_n \in [1, 2)$ and $\lim_{N \rightarrow \infty} \sum_{n=1}^N \beta_n = 1$, we can just set $\gamma = \Psi$ and this ensures that $\lim_{N \rightarrow \infty} \text{GDP/GO} = \Psi$. Crucially, one can verify that $\lim_{N \rightarrow \infty} \hat{G}_j = 1$. In plain words, the fragmentation of the value chain into arbitrarily many stages of production can drive the gains from trade to infinity without necessarily pushing the GDP/GO ratio to zero.

4.4. The Centrality-Downstreamness Nexus

We now exploit the tractability of our framework to explore the role of a country’s geography (and, in particular, its centrality) in shaping its average position in GVCs. In order to formalize a centrality-downstreamness nexus, let us define the average upstreamness of production of a given country i as

$$U(i) = \sum_{n=1}^N (N - n + 1) \times \frac{\sum_{j \in \mathcal{J}} \alpha_n \beta_n \times \Pr(\Lambda_i^n, j) \times w_j L_j}{w_i L_i}. \tag{14}$$

Recall that the term inside the summation in the numerator— $\alpha_n \beta_n \times \Pr(\Lambda_i^n, j) \times w_j L_j$ —represents how much of country j ’s total consumption of finished varieties, both by firms and by final consumers, is value-added produced directly in country i at stage n in GVCs. The summation across all countries j thus represents the worldwide amount of i ’s value-added produced at stage n , while the denominator is country i ’s GDP. Thus, the index $U(i)$ in (14) delivers the average stage at which country i produces value, with the general-equilibrium equation (12) implying that the weights add up to 1. This is thus a closely related measure to the index of upstreamness proposed by Antràs et al. (2012).

Centrality is defined similarly as the average distance to all countries weighted by each country’s aggregate consumption of finished varieties

$$C(i) = \sum_{j \in \mathcal{J}} \tau_{ij} \times \frac{w_j L_j}{\sum_{k \in \mathcal{J}} w_k L_k}. \tag{15}$$

We seek to establish a connection between the measures of upstreamness $U(i)$ and centrality $C(i)$. As in Section 2, the structure of equation (8) 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. We begin by showing that the centrality-downstreamness nexus holds perfectly under specific, restrictive assumptions, and then show that it remains a strong force in more general scenarios.

In order to develop a precise formulation of this result, let us assume that the ease of trade between any two countries i and j can be decomposed as $\tau_{ij} = \rho_i \rho_j$, where we take $\rho_i > 1$ to be an index negatively related to country i 's centrality. Notice that if country i is more central than country j , that is, $\rho_i < \rho_j$, then it is cheaper to ship from country i to any other country in the world than it is to ship from country j . Under these strong conditions, ρ_i entirely summarizes a country's centrality, that is, if $\rho_i < \rho_j$, then $C(i) < C(j)$. 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.1.7):

PROPOSITION 5: *If trade costs are log-separable and $\alpha_n \beta_n = 1/N$ for all n , then the more central a country is, the lower is the average upstreamness of this country in global value chain production. In other words, $\rho_i < \rho_j$ implies that $U(i) < U(j)$.*

The condition in the proposition, $\alpha_n \beta_n = 1/N$, serves to isolate the role of geography in shaping GVC positioning. 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*.¹⁸

The assumption $\alpha_n \beta_n = 1/N$ implies, however, that $\alpha_n = 1/n$, and thus the value-added intensity of production decreases along the chain. This condition might then appear not to be entirely innocuous for our result. More specifically, because more central countries tend to feature higher wages in general equilibrium, it would appear that if value-added intensity instead rose sufficiently fast along the value chain, the link between downstreamness and centrality might be broken. Nevertheless, the perfect correlation between centrality and downstreamness continues to apply for *any* path of α_n , as long as trade costs are log-separable and countries are symmetric with $T_j = T$ and $L_j = L$. While we have been unable to show this result formally, we have run millions of simulations that strongly indicate that this statement is true.¹⁹

Of course, in the real world, trade costs are not log-separable and countries are not symmetric. With that in mind, we now explore via simulations whether the centrality-downstreamness nexus remains an important force even in more general scenarios.

¹⁸For instance, if downstream stages contributed more to overall value added than upstream stages, we would obtain a prediction analogous to that in Costinot, Vogel, and Wang (2013), namely, that countries with better technologies T_i have comparative advantage in downstream stages. There is an obvious connection between our centrality-downstreamness result and that key result in Costinot, Vogel, and Wang (2013). In both cases, if a country characteristic translates into a share of output being "wasted," this proves to be more costly downstream than upstream. Conceptually, however, there is an important distinction stemming from the fact that centrality affects the ability of countries to produce for export but also to import for consumption. More generally, beyond the log-separable case $\tau_{ij} = \rho_i \rho_j$, centrality is not just a country-characteristic, but it depends in a complex manner on country-pair trade costs, as further explored below.

¹⁹At the end of the proof of Proposition 5, we show that a sufficient condition for the result to hold true is that $c_i^{-\theta} \rho_i$ be increasing in centrality ρ_i . Our simulations indeed show that, regardless of the value of θ , this condition is met for *all* possible paths of α_n whenever countries are symmetric in all respects except for their centrality.

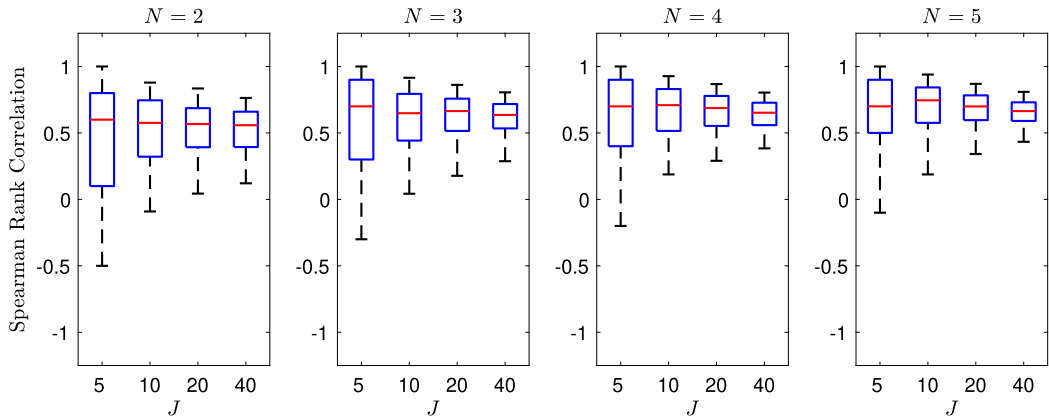


FIGURE 1.—Centrality downstreamness simulations.

Specifically, conditional on a number of countries J and stages of production N , we take random draws for all of the model's parameters, including trade costs which are only restricted to be symmetric ($\tau_{ij} = \tau_{ji}$) and zero domestically ($\tau_{ii} = 1$), and compute the Spearman rank correlation between centrality $C(i)$ and average upstreamness $U(i)$.²⁰ That is, in each simulation we compute the equilibrium and rank countries from 1 to J according to both centrality and average upstreamness and compute the correlation between these two rankings. Under the strong conditions of Proposition 5, this correlation is perfect and the correlation coefficient equals 1. In contrast, if the downstreamness ranking were the reverse of the centrality ranking, the correlation coefficient would equal -1 .

Figure 1 presents the distribution of correlation coefficients across 10,000 simulations of the model for various combinations of J and N . In order to keep results tidy, we summarize the distributions using box plots in which the middle line corresponds to the median, the box to the range 25–75% of values, and the whiskers to the range 5–95% of values. As can be seen, the rank correlation is positive in the overwhelming majority of cases, indicating that more central countries tend to be more downstream. Furthermore, the association is quite strong with the median simulation featuring an average correlation well above 0.5. Finally, the centrality-downstreamness nexus becomes stronger when focusing on longer sequences of production (higher N) and the distribution of values tends to become tighter when including more countries (higher J). We will shortly provide suggestive empirical evidence consistent with this prediction.

4.5. Suggestive Evidence

We complete this section by exploring the empirical relevance of two of our model's key ingredients, namely, the fact that the trade elasticity is larger for downstream than for upstream stages, and the centrality-downstreamness nexus discussed above. These empirical tests are reduced-form in nature and not structurally related to our model, but we nonetheless deem them to be informative.

²⁰We sample $\alpha_n \beta_n$ from a uniform distribution while normalizing $\sum_{n=1}^N \alpha_n \beta_n = 1$, T_j and L_j from lognormal distributions, and sample $\tau_{ij} = (1 + u_{ij})(1 + u_{ji})$ where u_{ij} is a uniform random variable. The variances of these distributions are chosen so that the standard deviations of all variables are roughly comparable.

TABLE I
TRADE COST ELASTICITIES FOR FINAL GOODS AND INTERMEDIATE INPUTS^a

	Total Flows		Final Good Flows		Input Flows	
	(1)	(2)	(3)	(4)	(5)	(6)
Distance	-1.111 (0.019)	-0.718 (0.013)	-1.229 (0.021)	-0.810 (0.015)	-1.058 (0.018)	-0.681 (0.013)
Contiguity		1.149 (0.090)		1.161 (0.098)		1.154 (0.088)
Language		0.402 (0.024)		0.499 (0.028)		0.358 (0.023)
Domestic		5.320 (0.168)		5.610 (0.191)		5.061 (0.159)
Observations	32,400	32,400	32,400	32,400	32,400	32,400
R ²	0.980	0.983	0.966	0.970	0.980	0.983

^aNotes: Standard errors are clustered at the country-pair level. All regressions include exporter and importer fixed effects.

We begin by studying empirically the compounding effect of trade costs. 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

$$\ln \tau_{ij} = \ln \kappa + \delta_{\text{dist}} \ln \text{Distance}_{ij} + \delta_{\text{con}} \text{Contiguity}_{ij} + \delta_{\text{lang}} \text{SameLanguage}_{ij}.$$

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 indicates a differential sensitivity of “up-stream” 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 this 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 I using bilateral trade flows. Columns (1) and (2) 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 are contiguous, share a common language, or are the same country (i.e., domestic trade). Our shipments data are from 2011 and correspond to the Eora MRIO database. 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 data set for the year 2006 (the most recent one available), and the merge between these two data sources leaves us with information on 180 countries (and thus $180 \times 180 = 32,400$ observations).²¹

²¹To proxy for distance between i and j , we use the variable “dist” from the CEPII database. For $i \neq j$, this corresponds to the geodesic distance between the most populated cities in these countries. For internal

Our results in columns (1) and (2) are fairly standard. Distance reduces trade flows with an elasticity of around -1 , while contiguity, common language, and domestic trade have a sizable 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 re-run the specifications in columns (1) and (2) but restricting the analysis to final-good flows, while in columns (5) and (6), we do the same focusing on intermediate input flows. As is apparent from the table, the elasticity of trade flows to distance is significantly larger for final-good trade (-1.229) than for intermediate-input trade (-1.058). The difference is sizable, highly statistically significant, and robust to controlling for alternative determinants of trade flows. Furthermore, comparing columns (4) and (6), we see that the positive effect of contiguity, common language, and within-border trade on trade flows is attenuated when focusing on the intermediate-input component of trade, with the effects being highly statistically significant for the case of common language and domestic trade.

Taken together, the results in Table I are highly suggestive of trade barriers impeding trade more severely in downstream stages than in upstream stages.²² In Web Appendix B.3, we further show that our results are not materially affected when pooling data from all years (1995–2013) for which the Eora data set is available (instead of just using 2011 data).²³ We further show that we obtain very similar results when restricting the analysis to manufacturing trade flows. We also repeat our tests using data from the two releases of the WIOD database, which cover a smaller and more homogeneous 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 data set, we only find support for the second differential effect (see Web Appendix B.3 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, strictly speaking, are inconsistent with our theoretical framework. Specifically, in Section 5.3, we will show that bilateral trade flows will typically be affected by trade costs associated with third countries (see Morales, Sheu, and Zahler (2019), and Adao, Costinot, and Donaldson (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 5. For that purpose, we build on Antràs et al. (2012) who proposed a measure of the positioning of countries in GVCs and studied how this measure correlates with various country-level variables. More specifically, Antràs et al. (2012) proposed a measure of industry “upstreamness” (or average distance of an industry’s output from final use) similar to that in equation (14), and then computed the average upstreamness of a country’s export vector using trade flow data from the BACI data set for the year 2002. Column (1) of Table II reproduces exactly their baseline specification, which includes 120 countries, and correlates a country’s upstreamness with its

distance ($i = j$), the CEPII database follows the bulk of literature in constructing an area-based measure as follows: $\text{dist}_{ii} = 0.67\sqrt{\text{area}/\pi}$ (see Mayer and Zignago (2006)).

²²Our results are consistent with the findings of Baldwin and Taglioni (2014), who estimated an elasticity of trade flows to trade costs—as implied by the ratio of cif to fob trade flows—that is almost twice as large for final goods than for intermediate. Using a significantly different specification, Bergstrand and Egger (2010) estimated fairly similar distance elasticities for final goods and for inputs.

²³In fact, we obtain extremely stable results when running these regressions year-by-year for this same period.

TABLE II
EXPORT UPSTREAMNESS AND CENTRALITY^a

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Centrality (GDP weighted)			-0.173 (0.065)		-0.233 (0.061)		-0.155 (0.044)	
Centrality (pop. weighted)				-0.228 (0.084)		-0.239 (0.081)		-0.238 (0.081)
$\log(Y/L)$	0.083 (0.142)	0.082 (0.142)	0.102 (0.138)	0.046 (0.148)	0.083 (0.046)	0.006 (0.037)		
Rule of Law	-0.029 (0.103)	-0.026 (0.104)	0.010 (0.105)	0.010 (0.105)				
Credit/ Y	-0.437 (0.136)	-0.440 (0.137)	-0.375 (0.130)	-0.407 (0.135)				
$\log(K/L)$	0.156 (0.131)	0.159 (0.132)	0.163 (0.129)	0.188 (0.132)				
Schooling	-0.085 (0.031)	-0.085 (0.031)	-0.083 (0.030)	-0.094 (0.029)				
Observations	120	118	118	118	118	118	118	118
R^2	0.154	0.153	0.194	0.199	0.083	0.056	0.054	0.054

^aNotes: Robust standard errors reported.

GDP per capita, rule of law, financial development, capital-labor ratio, and human capital (schooling).²⁴ Only financial development and schooling have a statistically significant partial correlation with upstreamness.

In order to assess the relationship between upstreamness and centrality, we simply add a measure of centrality to the core specification in column (1). In particular, for each country j , we compute $\text{Centrality}_j^{\text{GDP}} = \sum_i (\text{GDP}_i / \text{Distance}_{ji})$ and $\text{Centrality}_j^{\text{pop}} = \sum_i (\text{Population}_i / \text{Distance}_{ji})$, 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: columns (5) and (6) show that it persists when only controlling for GDP per capita, and columns (7) and (8) demonstrate that it holds even unconditionally. In Web Appendix B.3, 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 5.

5. EXTENDING THE MODEL AND MAPPING IT TO INPUT-OUTPUT DATA

While our discussion so far has been centered on a stylized model of sequential production, we now demonstrate the flexibility and applicability of our tools by extending our

²⁴The source for each of these variables is discussed in Antràs et al. (2012).

framework in ways that incorporate other key features present in standard international trade models. In fact, we show that this extended model nests many of the literature’s existing Ricardian models. At the end of this section, we also show how to map our model to the observable datapoints contained in input-output tables.

5.1. Extending the Model

We generalize our model along three main dimensions. First, we assume there are K industries indexed by $k \in \mathcal{K}$. Second, we assume that the production function depends not only on the country, industry, and stage of production but also on the use of output. For example, we assume that the production of electronics may vary depending on whether these will be used by firms in the form of intermediate inputs or by consumers in the form of final goods. Further, we allow the production of intermediate inputs to vary depending on the purchasing industry—so, for example, the production of electronics sold to the textile industry might differ from the production of electronics sold to the machinery industry. As we will show below, this flexibility will be useful for comparing our sequential production model to the literature’s “state-of-the-art” roundabout production models.

With these two properties in mind, define the n -stage price of an industry k intermediate-input variety z produced for industry k' through a supply chain ℓ as

$$p_{\ell(n)}^{n,k,k'}(\ell, z) = (a_{\ell(n)}^{n,k,k'}(z)c_{\ell(n)}^k)^{\alpha_n^{k,X}} (p_{\ell(n-1)}^{n-1,k,k'}(\ell, z)\tau_{\ell(n-1)\ell(n)}^k)^{1-\alpha_n^{k,X}}.$$

Producing an industry k good thus requires N stages of sequential production with $1 - \alpha_n^{k,X}$ denoting the expenditure share on upstream sequential inputs at stage n (with $1 - \alpha_1^{k,X} = 0$, $\beta_{n-1}^{k,X} = (1 - \alpha_n^{k,X})\beta_n^{k,X}$, and the superscript X denoting that these shares correspond to the sequential production of intermediate inputs).²⁵ Crucially, note that we assume that the production function depends on the use of output (i.e., the industry k' that will purchase this intermediate input) only through the productivity shifters $a_j^{n,k,k'}(z)$ for all n . Following the lead-firm approach, we assume that the overall productivity of chain ℓ in industry k in the production of intermediate inputs sold to k' is characterized as

$$\Pr\left(\prod_{n=1}^N (a_{\ell(n)}^{n,k,k'}(z))^{\alpha_n^{k,X}\beta_n^{k,X}} \geq a\right) = \exp\left\{-a^{\theta^k} \prod_{n=1}^N (T_{\ell(n)}^{k,k'})^{\alpha_n^{k,X}\beta_n^{k,X}}\right\}. \tag{16}$$

Chain-level productivity is thus governed by the industry-specific comparative advantage parameter θ^k and the producing-industry- and buying-industry-specific absolute advantage parameter $T_j^{k,k'}$. To aid the reader’s intuition, Appendix Section A.4 depicts this production structure graphically.

Final-good production is analogous with $p_{\ell(n)}^{n,k,F}(\ell, z)$ denoting the n -stage price of an industry k final-good variety z produced through supply chain ℓ . As in the case for intermediate inputs, producing each final-good variety requires N stages of production governed by the stage expenditure shares $\alpha_n^{k,F}$ and the country-industry-stage-specific productivity

²⁵Note that an industry k can effectively require only $N^k \leq N$ stages of sequential production if $1 - \alpha_n^{k,X} = 0$ for $n = 1 + (N - N^k)$. Hence, instead of making notation heavier and assuming that N^k varies at the industry-level, we use a single common N and assume that the effective number of required stages of sequential production are determined by the expenditure shares $\alpha_n^{k,X}$.

shifters $a_{\ell(n)}^{n,k,F}(z)$ shaped by the absolute technology parameters $T_j^{k,F}$ in a way analogous to (16).

Finally, as a third extension of our framework, we introduce input-output linkages. Specifically, we assume that the unit cost of production in industry k' equals

$$c_j^{k'} = (w_j)^{\gamma_j^{k'}} \prod_{k \in \mathcal{K}} (P_j^{k,k'})^{\gamma_j^{k,k'}}, \tag{17}$$

where $\gamma_j^{k'}$ denotes the value-added share and $\gamma_j^{k,k'}$ the expenditure on industry k finished intermediate inputs (with $\gamma_j^{k'} + \sum_{k \in \mathcal{K}} \gamma_j^{k,k'} = 1$ for all k'). This unit cost depends on wages w_j and on the price indices $P_j^{k,k'}$ associated with the unit cost of a CES bundle of industry k varieties purchased by industry k' in country j .²⁶ Analogously, we denote by $P_j^{k,F}$ the price index of a CES bundle of industry k varieties purchased by final consumers in country j . As is standard, we also assume that consumer preferences are Cobb–Douglas across industries with ζ_j^k denoting the expenditure share on industry k final goods, and $P_j^F = \prod_{k \in \mathcal{K}} (P_j^{k,F})^{\zeta_j^k}$ the consumer price index.²⁷

Invoking the famous terminology developed by Baldwin and Venables (2013), these three extensions turn our benchmark model of “snakes” (i.e., of purely sequential value chains) into a richer production structure that incorporates “spider”-like features, with each production stage sourcing parts and components from several industries. Note, however, that the “spider legs” in our model are composite bundles of snakes, rather than individual, customized snakes.

The multi-industry equilibrium can be characterized using derivations similar to the single-industry case. Following the lead-firm approach, the share of industry k intermediate-input varieties sourced through ℓ by industry k' in country j equals

$$\pi_{\ell j}^{k,k'} = \frac{\prod_{n=1}^{N-1} ((T_{\ell(n)}^{k,k'})^{\alpha_n^X} ((c_{\ell(n)}^k)^{\alpha_n^X} \tau_{\ell(n)\ell(n+1)}^k)^{-\theta^k})^{\beta_n^X} \times (T_{\ell(N)}^{k,k'})^{\alpha_N^X} ((c_{\ell(N)}^k)^{\alpha_N^X} \tau_{\ell(N)j}^k)^{-\theta^k}}{\Theta_j^{k,k'}}.$$

The share of industry k final-good varieties sourced through ℓ by consumers in j is defined analogously and denoted by $\pi_{\ell j}^{k,F}$.

5.2. Relationship to Previous Ricardian Trade Models

Let us now describe the relationship between our extended sequential production model and other Ricardian models of trade through the lens of the the “gains-from-trade” formula. Analogous to (13), the change in real income following any change to trade barriers in the extended model can be written in terms of a set of changes in the domestic expenditure shares as

$$\hat{W}_j = \prod_{k \in \mathcal{K}} (\hat{\pi}_{jj}^{k,F})^{-\frac{1}{\theta^k} \zeta_j^k} \times \prod_{k \in \mathcal{K}} \prod_{k' \in \mathcal{K}} \prod_{k'' \in \mathcal{K}} (\hat{\pi}_{jj}^{k,k'})^{-\frac{1}{\theta^k} \gamma_j^{k,k'} \delta_j^{k',k''} \zeta_j^{k''}}, \tag{18}$$

²⁶Note that industry k' in j buys each variety z from the supply chain ℓ that minimizes $P_{\ell(N)}^{N,k,k'}(\ell, z) \tau_{\ell(N)j}^k$.

²⁷While our model is fairly general, it is straightforward to generalize even further and let the production parameters vary across stages and depending on the use of output by defining $\gamma_j^{n,k,k'}$, $\alpha_n^{k,k',X}$, $T_j^{n,k,k'}$, $T_j^{n,k,F}$.

where, as before, a “hat” indicates the (gross) change in a variable and where j is a purely domestic chain of size N . Additionally, the auxiliary parameters $\delta_j^{k,k'}$ equal the gross value of industry k inputs used in the production of one dollar of industry k' goods across all stages of production.²⁸

The intuition for this formula is straightforward and depends on three forces. First, the gains from trade depend on the expenditure shares on domestic supply chains since these are an inverse measure of how much a country exploits comparative advantage through international trade. Second, the powers on the expenditure shares indicate the importance of each supply chain for domestic consumption in autarky. That is, relative to one dollar of aggregate final consumption, ζ_j^k indicates the gross value of industry k finished final goods that need to be produced, while $\sum_{k',k'' \in \mathcal{K}} \gamma_j^{k,k'} \delta_j^{k',k''} \zeta_j^{k''}$ captures the gross value of industry k finished intermediate inputs that need to be produced for industry k' across all stages of the supply chain (including all input-output linkages) to satisfy that dollar of final consumption. Third, the trade elasticities θ^k regulate the sensitivity of the losses from increases in trade barriers due to the interaction between how much a country is already producing along a specific supply chain in the trade equilibrium (the domestic shares) and the importance of a given supply chain for final consumption (the gross value weights).

This extension nests many of the previously developed Ricardian models (this can be seen graphically in the flow charts in Appendix Section A.4). On the one hand, this model nests the single-industry sequential production model developed above when $K = 1$ and when restricting production technology T_j and α_n to be the same for intermediate inputs and final goods, in which case equation (18) reduces to (13). On the other hand, our model nests “state-of-the-art” Ricardian models featuring roundabout production. Specifically, letting $N = 1$, the gains from trade are given exactly by the formula in (18) but depend on the aggregate domestic expenditure on industry k final goods given by $\Pi_{jj}^{k,F}$ and the aggregate domestic expenditure on industry k intermediate inputs by each industry k' given by $\Pi_{jj}^{k,k'}$, as in the recent work by Alexander (2017). Further restricting the model to $T_j^{k,k'} = T_j^{k,F} = T_j^k$ for all k' delivers exactly the influential roundabout model of Caliendo and Parro (2015), which includes multiple industries and input-output linkages but which features a single aggregate domestic expenditure share Π_{jj}^k for each industry. In addition, restricting to the single-industry case and a common value-added share across countries delivers the classic Eaton and Kortum (2002) model. In sum, the formula in (18) illustrates how the sequential production, multi-industry, and input-output linkages elements interact and shape the gains from trade.²⁹

5.3. Mapping the Model to Observables

So far, we have described how to extend standard Ricardian models to a multi-stage environment that delivers sharp predictions for the “trade shares” $\pi_{\ell j}^{k,k'}$ and $\pi_{\ell j}^{k,F}$ in terms of the relative prevalence of specific production paths (or GVCs) rather than in terms of gross trade volumes. Unfortunately, these “GVC trade shares” are not directly observable

²⁸Formally, $\delta_j^{k,k'}$ is an element of $\delta_j = [\mathbb{I} - \gamma_j]^{-1}$ where γ_j is the $K \times K$ matrix of elements $\gamma_j^{k,k'}$.

²⁹Arkolakis, Costinot, and Rodríguez-Clare (2012) showed that analogous formulas apply under Armington frameworks with roundabout production. This is also the case in our family of models with sequential production. For example, de Gortari (2019) developed a supply-chain Armington model which delivers a gains from trade formula identical to (18) when imposing $N = 2$.

		Input use & value added			Final use			Total use
		Country 1	...	Country <i>J</i>	Country 1	...	Country <i>J</i>	
Output supplied	Country 1	X_{11}	...	X_{1J}	F_{11}	...	F_{1J}	GO_1
	Country <i>J</i>	X_{J1}	...	X_{JJ}	F_{J1}	...	F_{JJ}	GO_J
Value added		GDP_1	...	GDP_J				
Gross output		GO_1	...	GO_J				

FIGURE 2.—A schematic World Input–Output Table.

in the data, so we next describe how to map the model to the type of information available in World Input–Output Tables (WIOTs).

First, let us describe the datapoints contained in WIOTs while aiding the exposition with the schematic representation in Figure 2. This matrix is split into two blocks with the left one containing data on bilateral intermediate-input trade flows across countries and the right one containing the data for final-good trade flows. Specifically, each $X_{ij} = [X_{ij}^{kk'}]$ is the $K \times K$ matrix of bilateral intermediate-input flows from industry k in country i sold to industry k' in country j . Analogously, $F_{ij} = [F_{ij}^k]$ is the $K \times 1$ vector of bilateral final-good sales of industry k in country i to consumers in country j . Hence, rows across the WIOT represent the sales of a country-industry to other country-industries and to consumers, with the sum across a row equaling a country-industry’s gross output. Meanwhile, columns in the left block represent the purchases of intermediate inputs by country-industries (with the sum across a column equaling gross output minus value added) and columns in the right block represent the final-good purchases of final-good consumers in each country (with the sum across a column equaling aggregate final-good consumption).³⁰

Second, consider the implications of our model for final-good consumption. Notice that for final goods to flow from a given source country i to consumers in a given destination country j , it must be the case that country i is in position N in a chain producing final goods for country j . Remember from Section 4.2 that Λ_i^n represents the set of GVCs flowing through i at position n so that Λ_i^N corresponds to the set of chains in which final assembly is carried out in i . With this notation, the dollar value of final goods shipped from industry k in country i to consumers in country j equals

$$F_{ij}^k = \sum_{\ell \in \Lambda_i^N} \pi_{\ell j}^{k,F} \times \zeta_j^k w_j L_j, \tag{19}$$

where $\zeta_j^k w_j L_j$ is country j ’s total consumption of industry k final goods.³¹

Third, note that computing intermediate-input flows between any two countries i and j is more intricate, but equally straightforward. The difficulty lies in the fact that country j potentially buys different shares of inputs from i at different stages of the supply chain, yet WIOTs only report the aggregate amount of intermediate inputs traded across all stages

³⁰Note that the difference between aggregate final consumption and value added is the trade deficit or surplus. These deficits are nontrivial for certain countries and are taken into account in both our estimation and our counterfactual exercises, as discussed below.

³¹Notice that we are interpreting the probability function $\pi_{\ell j}^{k,F}$ as an expenditure share. This is justified by the fact that all finished varieties in country j command the same expected price regardless of the actual chain of production, exactly as in Eaton and Kortum (2002). For the same reason, below we will also treat $\pi_{\ell j}^{k,k'}$ as an expenditure share.

of the supply chain. As we now show, this challenge can be overcome by introducing a set of auxiliary variables ϑ_j^k capturing the value-added created in each country j and industry k and then solving jointly for these variables and the bilateral intermediate-input flows.

To begin, note there are two broad types of traded inputs. The first type are finished N -stage intermediate inputs which are traded both within and across industries—in other words, the intermediates traded through the input-output linkages characterized by $\gamma_j^{k,k'}$. The second type are n -stage sequential inputs (with $n < N$) which are only traded within industries—and disciplined by the sequential production expenditure shares $\alpha_n^{k,X}$ and $\alpha_n^{k,F}$.

Cross-industry input flows are made entirely of inputs of the first type such that the input flow sold from industry k in country i to industry $k' \neq k$ in country j equals

$$X_{ij}^{kk'} = \sum_{\ell \in \Lambda_i^N} \pi_{\ell_j}^{k,k'} \times \frac{\gamma_j^{k,k'}}{\gamma_j^{k'}} \vartheta_j^{k'}. \tag{20}$$

Since $\vartheta_j^{k'}$ is a value-added measure, then $\gamma_j^{k,k'} / \gamma_j^{k'}$ imputes the aggregate amount of N -stage inputs purchased from industry k used in this production and then $\sum_{\ell \in \Lambda_i^N} \pi_{\ell_j}^{k,k'}$ imputes the share of inputs produced through supply chains in which i produces the N th stage.

Within-industry input flows, on the other hand, equal the sum of the finished N -stage inputs plus the sequential inputs used to produce both finished N -stage intermediate inputs and finished N -stage final goods:

$$\begin{aligned}
 X_{ij}^{kk} &= \underbrace{\sum_{\ell \in \Lambda_i^N} \pi_{\ell_j}^{k,k} \times \frac{\gamma_j^{k,k}}{\gamma_j^k} \vartheta_j^k}_{N\text{-stage intermediates}} \\
 &+ \sum_{n=1}^{N-1} \sum_{\ell \in \Lambda_{i \rightarrow j}^n} \sum_{h \in \mathcal{J}} \left(\underbrace{\beta_n^{k,X} \sum_{k' \in \mathcal{K}} \pi_{\ell_h}^{k,k'} \times \frac{\gamma_h^{k,k'}}{\gamma_h^{k'}} \vartheta_h^{k'}}_{n\text{-stage intermediates}} + \underbrace{\beta_n^{k,F} \pi_{\ell_h}^{k,F} \times \zeta_h^k w_h L_h}_{n\text{-stage final goods}} \right). \tag{21}
 \end{aligned}$$

The first term on the right-hand side is analogous to the cross-industry flow in (20) and captures the N -stage inputs traded within industry k across countries i and j . The second term is similar, but instead captures the n -stage inputs traded within industry k across countries i and j used in the production of N -stage intermediate-input varieties consumed in all markets $h \in \mathcal{J}$. In terms of notation, remember that $\Lambda_{i \rightarrow j}^n = \{\ell \in \mathcal{J}^N | \ell(n) = i \text{ and } \ell(n+1) = j\}$ accounts for all supply chains through which i and j produce sequentially the n th and $(n+1)$ th stages of production, respectively. This term also includes a $\beta_n^{k,X}$ which accounts for the fact that for every dollar of N -stage varieties bought by some country, only $\beta_n^{k,X}$ cents are traded at stage n . Finally, the third term captures the total amount of i 's output shipped to j at stage n in final-good varieties that are eventually consumed as N -stage final goods across all countries $h \in \mathcal{J}$.³²

³²Note the difference between finished and final goods. Finished refers to stage- N goods which are traded both as intermediate inputs used to produce further downstream goods and as final goods sold to consumers.

Having defined input flows in terms of value added, the latter can in turn be defined in terms of the former as the difference between gross output and aggregate intermediate-input purchases:

$$\vartheta_i^k = \sum_{j \in \mathcal{J}} \left(\sum_{k' \in \mathcal{K}} X_{ij}^{k,k'} + F_{ij}^k \right) - \sum_{j \in \mathcal{J}} \left(\sum_{k' \in \mathcal{K}} X_{ji}^{k',k} \right). \tag{22}$$

Equations (20), (21), and (22) thus pin down both the bilateral intermediate-input flows and value-added terms as a function of wages. Finally, the general-equilibrium equation is such that aggregate value added is equal to the return on the factors of production: $w_j L_j = \sum_{k \in \mathcal{K}} \vartheta_j^k$.

In sum, and although computing these flows has proved somewhat cumbersome, a key feature of our model is that it provides explicit formulas that link the various entries of a WIOT to fundamental parameters of the model, as well as to endogenous variables determined by these same parameters.³³

6. ESTIMATION

We now describe how we leverage the fact that our model delivers explicit analytical counterparts to the entries of a WIOT in order to structurally estimate the model’s key parameters. For computational reasons, we will mostly focus on an application in which the various industries in a WIOT are collapsed into a single sector. Nevertheless, following the lead of Alexander (2017), we will allow for asymmetries in the technology to produce final goods and inputs. It is undeniable that an estimated multi-industry version of our model could deliver estimates with significantly different implications for the counterfactuals studied in Section 7, so readers are entitled to treat the results in the next two sections as a proof of concept, rather than as a definitive quantitative exercise. Having said this, in Section 7, we will briefly describe the empirical results of an extension of our model featuring two sectors, manufacturing and services.³⁴

Before delving into the details of the estimation, we briefly describe our data source. Note that building a WIOT of the type in Figure 2 is a formidable endeavor because it requires collecting trade and production data from many different sources, including

³³We have demonstrated how trade flows within our model’s GVCs aggregate into the bilateral input-output trade flows observed in the data. de Gortari (2019) studied the inverse question of how to use the aggregate data to disentangle the shape of the GVCs underlying these flows.

³⁴We do not go beyond this two-industry model for four main reasons. First, the number of parameters to estimate increases quadratically with the number of industries. The parameters to estimate in each industry are $\{\alpha_n^{k,X}, \beta_n^{k,X}, T_j^{k,k'}, \alpha_n^{k,F}, \beta_n^{k,F}, T_j^{k,F}, \gamma_j^{k,k'}, \gamma_j^k\}$ implying $\sum_k 2(N^{k,X} - 1 + N^{k,F} - 1) + 2JK(K + 1)$ parameters in total. Second, computing the model’s general equilibrium requires solving for $\{w_j, P_j^{k,k'}, P_j^{k,F}, \vartheta_j^k\}$ or $J(1 + K(K + 1))$ endogenous variables, and thus becomes increasingly complicated as more industries are added. Third, our MPEC estimation procedure requires deriving and computing the model’s Jacobian matrix which is itself of size $[\sum_k 2(N^{k,X} - 1 + N^{k,F} - 1) + JK(J + 3)(K + 1) + J] \times [\sum_k (N^{k,X} - 1 + N^{k,F} - 1) + JK(J + 1)(K + 1)]$. Each dimension of the Jacobian matrix increases quadratically with K , thus implying that the number of elements increases at the rate of K^4 . These three reasons imply that increasingly more computational power is needed in each iteration of the estimation algorithm when more industries are added. Finally, a fourth reason is that the problem appears to behave increasingly nonlinearly when more industries are added. This nonlinearity both increases the number of iterations required for the estimation algorithm to converge and decreases the reliability of the solution being at a global minimum (as is discussed further in the Supplemental Material). We thus focus on simple quantitative explorations of our model, and hope that future work will uncover quicker and more efficient computational procedures for unleashing the model’s quantitative power in more complex environments.

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 that was carried out by a consortium of 12 research institutes headed by the University of Groningen in the Netherlands (see [Timmer, Dietzenbacher, Los, Stehrer, and de Vries \(2015\)](#)). We choose this data set for our 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 two countries in Latin America (Brazil and Mexico).³⁵

The data contained in WIOTs (when aggregated at the country level) are a set of aggregate bilateral intermediate-input flows denoted by \bar{X}_{ij} and a set of aggregate bilateral final-good flows denoted by \bar{F}_{ij} . Our notation is such that “bars” define observed data-points. It will often be useful to describe the data using aggregate expenditure shares, which we define as

$$\bar{\Pi}_{ij}^X = \frac{\bar{X}_{ij}}{\sum_{i' \in \mathcal{J}} \bar{X}_{i'j}} \quad \text{and} \quad \bar{\Pi}_{ij}^F = \frac{\bar{F}_{ij}}{\sum_{i' \in \mathcal{J}} \bar{F}_{i'j}}. \quad (23)$$

In order to describe how we use the data to discipline our model in our estimation procedure, let us first outline the parameters needed 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} , which are assumed common for inputs and for final goods. Production depends, first, on the labor value-added shares γ_j , which we allow to be country-specific but common across stages. Second, production depends on the Cobb–Douglas input expenditure shares α_n^X and α_n^F , which are allowed to vary across the production of intermediate-input and final-good varieties, and which are stage-specific but common across countries. Third, labor productivity depends on the country-specific state-of-technology levels T_j^X and T_j^F , which are also allowed to differ across intermediate input and final-good varieties (as in [Alexander \(2017\)](#)). The parameter θ governing the strength of comparative advantage is instead assumed common across stages and countries. Overall, production thus depends on $J + 2(N - 1) + 2J$ parameters together with the N -stage trade elasticity θ . Finally, 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 absolute productivity and of equilibrium wages, but have no bearing on the counterfactuals discussed below. With that in mind, we simply normalize equipped labor as $L_j = (\text{capital}_j)^{\frac{1}{3}} (\text{population}_j)^{\frac{2}{3}}$, with capital and population drawn from the Penn World Tables, while we normalize T_j^X and T_j^F to sum up to 100 across countries.

With regard to the length of the sequentiality of production N , we will argue that this parameter cannot be properly identified off WIOTs because bilateral flows represent highly aggregated counterparts of the rich supply chain flows underlying world trade flows (as can be seen in (19), (20), and (21)). Hence, our goal will be to use all of the variation contained in the data in order to pin down the model’s parameters conditional on a given N . We will then verify that increasing N necessarily improves the model’s fit,

³⁵Two 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 in 2013 contained information for 40 countries and the rest of the world, for the period 1995–2011. See <http://www.wiod.org>.

since this increases its flexibility and permits it to better match the variation contained in the data.³⁶ This approach will then let us compare the model’s parameterizations across different N ’s in order to study the effects of sequential production on our counterfactual experiments.

Let us now turn to a more detailed description of our estimation approach. 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, that is, $\tau_{jj} = 1$ 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 (19) and our empirical analogs in (23):

$$\tau_{ij}^{-\theta} = \sqrt{\frac{\overline{\Pi}_{ij}^F \overline{\Pi}_{ji}^F}{\overline{\Pi}_{ii}^F \overline{\Pi}_{jj}^F}} \tag{24}$$

Trade costs are symmetric by construction, that is, $\tau_{ij} = \tau_{ji}$, 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 this approach to backing out trade costs is that the calibrated values for $\tau_{ij}^{-\theta}$ 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 renormalization of T_j^X and T_j^F that yields the same equilibrium (conditional on the same set of parameters γ_j , α_n^X , and α_n^F). 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 $\theta\beta_1^X, \theta\beta_2^X, \dots, \theta$ and $\theta\beta_1^F, \theta\beta_2^F, \dots, \theta$).³⁷

To reiterate, our empirical strategy will be to fix N and find the remaining set of parameters that best fit the data. Specifically, having pinned down $\tau_{ij}^{-\theta}$ and θ , we fix N and estimate the set of parameters γ_j , α_n^X , α_n^F , T_j^X , and T_j^F by targeting all of the WIOD’s flows via the generalized method of moments, thus minimizing the difference between the observed and the model’s simulated WIOT flows.³⁸ We estimate our model using a constrained optimization approach based on the MPEC algorithm of [Su and Judd \(2012\)](#) described in Web Appendix Section B.4. The next section will then show how quantitative

³⁶The only case in which this does not hold is when our model represents the true data generating process—in which case increasing N beyond its true value bears no more fruits. While our model might be a good approximation to the real world, it is certainly not the real world’s exact data generating process and so this case never holds empirically (nonetheless, we explore this idea further in Section 6.2).

³⁷[Simonovska and Waugh \(2014\)](#), in a widely cited study, found a range for the elasticity of trade between 2.47 and 5.51. Using U.S. import data, [Antràs, Fort, and Tintelnot \(2017\)](#) estimated an elasticity of trade of 4.54. Finally, [Caliendo and Parro \(2015\)](#) found an aggregate elasticity of 4.49 using the previous release of the WIOD data. By simulating data from our model, we have indeed confirmed that the elasticity of aggregate bilateral trade flows to trade resulting from our model ranges from 4.47 for $N = 2$ to 4.51 for $N = 5$. By construction, when focusing on final-good flows, we recover a value of 5.

³⁸Specifically, for a given N , we minimize the weighted squared sum of differences between the observed and simulated WIOT flows. In 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 trade flows. Hence, we weigh each moment using the square root of the targeted moment and minimize $\sum_{i,j \in \mathcal{J}} \sqrt{\overline{X}_{ij}} (X_{ij} - \overline{X}_{ij})^2 + \sum_{i,j \in \mathcal{J}} \sqrt{\overline{F}_{ij}} (F_{ij} - \overline{F}_{ij})^2$.

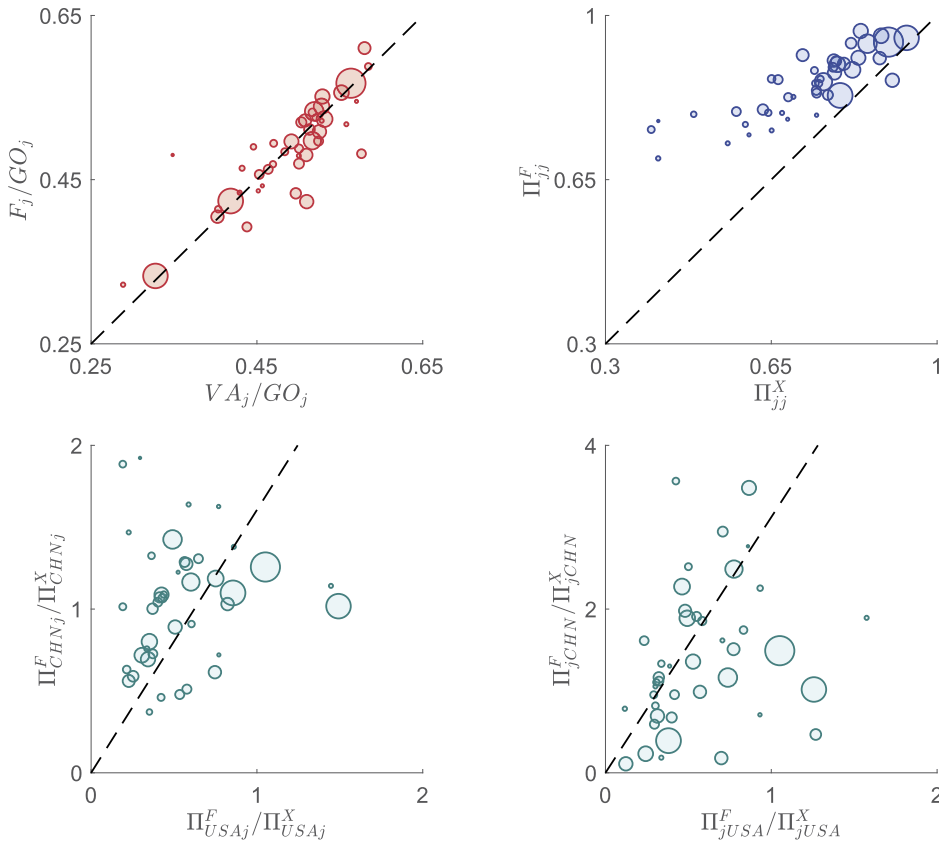


FIGURE 3.—Some key features of the World Input-Output Database.

counterfactual experiments vary when increasing N and thus will showcase the effects of incorporating multi-stage production. In particular, note that our estimated model with $N = 1$ corresponds to the roundabout model of Alexander (2017), so that this parameterized model will be a useful benchmark for comparing our sequential production model with $N > 1$ to roundabout production models.

Having described the estimation procedure, let us now discuss the variation in the data that allows us to identify each set of parameters. In particular, Figure 3 depicts four salient characteristics of the 2014 WIOD data that we use to pin down the production function parameters (note that circle size is proportional to country j 's GDP). As we shall now show, the variation in each type of chart will let us reject specific elements found in previous roundabout production models.

First, the northwest panel shows the relationship between a country's ratio of aggregate value added to gross output (VA/GO) and its share of final output relative to its gross output (F/GO). 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 VA/GO , a high share of F/GO indicates that a country is relatively downstream in GVCs. Similarly, for a given F/GO , a low VA/GO ratio indicates that a country uses a relatively large amount of foreign inputs in production, which again suggests a relatively downstream position of this country in GVCs. Hence, in a world in which countries are in

markedly different segments of GVCs, the ratios VA/GO and F/GO might be expected to be negatively correlated. With this background in mind, 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. To understand the implications of this for our estimation, envision the simplest version of our model (akin to Eaton and Kortum (2002)) with $N = 1$, $T_j^X = T_j^F$, and a common γ across countries. In this case, it is easy to see that $VA/GO = F/GO = \gamma$ for all countries and so this variation rejects this particular parameterization. Indeed, subsequent iterations of the Eaton and Kortum (2002) model, such as Caliendo and Parro's (2015) work, address this issue by letting γ_j vary across countries. This has the virtue of making VA/GO exactly equal to γ_j and of generating a positive correlation between VA/GO and F/GO. In sum, when $N = 1$, γ_j is identified by the cross-country variation in the northwest panel of Figure 3. While analytical expressions for VA/GO and F/GO are more complex when $N > 1$, it turns out that γ_j is also well suited for matching this variation in the sequential production model.³⁹

Second, the northeast panel shows the relationship between a country's domestic expenditure share on intermediate inputs and final goods. As is clear from the graph, most observations lie above the 45 degree line, indicating that the observed input $\bar{\Pi}_{jj}^X$ and output $\bar{\Pi}_{jj}^F$ matrices are asymmetric, and that countries tend to rely on foreign sources more prevalently for inputs than for final goods.⁴⁰ Since standard roundabout models such as Eaton and Kortum (2002) and Caliendo and Parro (2015) assume that all output is produced in the same way regardless of its use (i.e., $\Pi_{jj}^X = \Pi_{jj}^F$), the WIOD rejects these *symmetric* parameterizations. There are at least two alternative ways in which this variation can be rationalized. The first, adopted by Alexander (2017), is that absolute productivity levels $T_j^X \neq T_j^F$ vary across the production of intermediate inputs and final goods so that $\Pi_{jj}^X \neq \Pi_{jj}^F$. A second explanation is instead based on the central insight of this paper, namely, the notion that trade costs are more detrimental for downstream versus upstream stages of production. To see this, note that $\Pi_{jj}^X < \Pi_{jj}^F$ even if $T_j^X = T_j^F$ as long as there is sequential production with $N > 1$. The reason for this is that the sequential input expenditure shares α_n^X and α_n^F determine how fast the trade elasticity increases along GVCs, and are thus crucial in shaping the observed differences between these two domestic input shares. Thus, our model incorporates both channels so that the variation in the northeast panel of Figure 3 helps identify T_j^X , T_j^F , α_n^X , and α_n^F . In addition, note that while this variation identifies the relative size of T_j^X and T_j^F within countries, the levels across countries are pinned down by the size of each country's aggregate value added and trade flows.

Third, and finally, the bottom two panels in Figure 3 depict the relation between the ratio of import shares of final goods relative to intermediate inputs from two source countries (China and the U.S., in this particular case) across all countries (southwest panel), and the relation between the ratio of export shares of final goods relative to intermediate inputs to those two destination countries (southeast panel). This is the type of variation

³⁹To see this, imagine for simplicity $N = 2$ and that $\alpha_2^F = \alpha_2^X = \alpha_2$. As discussed, when $\alpha_2 \rightarrow 1$, the upstream stage of production is irrelevant for production and the VA/GO ratio equals γ_j . Conversely, when $\alpha_2 \rightarrow 0$, the downstream stage of production adds very little value, and so the VA/GO ratio is close to $\gamma_j/2$, since the same output is shipped twice but value is added essentially only once. In practice, for a general N , the VA/GO ratio features variation both because countries have different labor value-added shares but also because they find themselves at different degrees of upstreamness along the GVC; the interaction of both forces determines this statistic.

⁴⁰This asymmetry is consistent with the difference in the regression coefficients on the domestic dummy of columns (4) and (6) in Table I.

that roundabout models cannot handle—not even the highly flexible model of Alexander (2017). The reason for this is that when $N = 1$, then

$$\frac{\Pi_{ij}^F}{\Pi_{ij}^X} = \frac{T_i^F}{T_i^X} \times \frac{\Theta_j^X}{\Theta_j^F},$$

and so the model predicts that the relative ratio of import shares from any two given countries or the relative ratio of export shares to a given pair of countries must be proportional across all importing or exporting countries. In other words, the model with $N = 1$ predicts that all observations in both of the two bottom panels depicted in Figure 3 should lie along the single ray depicted in these graphs.⁴¹ Hence, this last variation in the data can only be fitted by the sequential production channel with, as before, the sequential expenditure input shares α_n^X and α_n^F delivering the degrees of freedom needed to pin this heterogeneity down. Crucially, increasing N delivers additional flexibility, and this is why the model’s fit improves as N grows.

Before turning to a discussion of our estimation results, we briefly comment on our treatment of trade imbalances. As mentioned above (see footnote 30), 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 Supplemental Material Section A.3).

6.1. Estimation Results

Table III presents our estimation results. The first five rows present the estimates of the sequential production input expenditure shares α_n^X and α_n^F and the mean value-added share γ_j across all countries for $N = 1, 2, 3, 4, 5$. We will refer to this set of estimates as asymmetric parameterizations since both the sequential input expenditure shares and

TABLE III
ESTIMATION RESULTS

N	Obj. Func.	α_{N-4}^X	α_{N-3}^X	α_{N-2}^X	α_{N-1}^X	α_N^X	α_{N-4}^F	α_{N-3}^F	α_{N-2}^F	α_{N-1}^F	α_N^F	Mean γ_j
<i>Asymmetric Parameterization</i>												
1	39.44					1.00					1.00	0.51
2	12.07				1.00	0.25				1.00	1.00	0.62
3	11.68			1.00	0.03	0.17			1.00	1.00	1.00	0.71
4	11.57		1.00	0.02	0.04	0.14		1.00	1.00	1.00	1.00	0.77
5	11.50	1.00	0.02	0.04	0.05	0.10	1.00	1.00	1.00	1.00	1.00	0.82
<i>Symmetric Parameterization</i>												
1	55.28					1.00					1.00	0.55
2	37.88				1.00	0.21				1.00	0.21	0.88
3	37.88			1.00	1.00	0.21			1.00	1.00	0.21	0.88

⁴¹Note that the model predicts that the ray in the southwest panel is given by $(T_{\text{CHN}}^F/T_{\text{CHN}}^X)/(T_{\text{USA}}^F/T_{\text{USA}}^X)$, while the ray in the southeast panel is endogenous and given by $(\Theta_{\text{CHN}}^X/\Theta_{\text{CHN}}^F)/(\Theta_{\text{USA}}^X/\Theta_{\text{USA}}^F)$.

absolute productivity levels depend on whether the N -stage goods are used as intermediate inputs or final goods. The last three rows of Table III instead present our estimation results when restricting $\alpha_n^X = \alpha_n^F$ and $T_j^X = T_j^F$, and we refer to this set of estimates as the symmetric parameterizations of our model. The full set of estimates can be found in Section A.5 of the Supplemental Material.

Let us first discuss the set of asymmetric parameterizations. Interestingly, we find that sequential production is only relevant for the production of intermediate inputs. That is, we find that $\alpha_N^F = 1$ for all N so that, effectively, the production of final-good varieties requires a single stage of production even when we let the model have potentially more stages. In other words, our estimation procedure indicates that the model's fit does not improve when moving from roundabout to sequential production in final goods (i.e., from $N = 1$ to $N > 1$).

In contrast, we find that increasing the number of stages of sequential production used to produce intermediate inputs does improve the model's fit as $\alpha_2^X, \dots, \alpha_N^X$ are all strictly less than 1 for $N > 1$. Crucially, it turns out that making use of this added flexibility has sizable effects on the levels of the estimated value-added shares γ_j . Intuitively, Table III shows that increasing the production length of intermediate inputs comes together with a high dependence on the upstream sequential inputs (i.e., high $1 - \alpha_n^X$) which tends to decrease the VA/GO ratio. To continue to match the data, γ_j thus needs to increase in order to offset this effect. As we discuss below, this has important implications for the comparison of our counterfactual experiments across the model parameterized with different values of N . Having said this, we should note that the numerical fit of the model improves only very slightly when moving from $N = 2$ to larger N (i.e., the minimized objective function becomes increasingly flat when N increases beyond 2 as can be seen in Table III's second column). The reason is that, while variation in the country-level parameters γ_j , T_j^X , and T_j^F significantly enhances the model's flexibility, incorporating additional stage-level parameters α_n^X only enhances the model's flexibility marginally since this added flexibility has to be averaged across all countries. The fit of the model with $N = 2$ is, however, significantly better than with $N = 1$.

On the other hand, and perhaps surprisingly, the symmetric parameterizations indicate no need of additional flexibility beyond the sequential $N = 2$ case. That is, while the model's fit improves substantially between the roundabout and the two-stage model, estimating the symmetric model with $N > 2$ leads to a set of parameters that effectively shut down the more upstream stages (i.e., $\alpha_{N-1} = 1$ for all $N \geq 2$). As in the asymmetric case, the intuition for this result can be obtained from studying the estimated value-added shares γ_j . Since the symmetric roundabout model imposes $\Pi_{jj}^X = \Pi_{jj}^F$, the improvement in the model's fit achieved when moving from $N = 1$ to $N = 2$ is driven by the parameter α_2 , which is estimated to be relatively low ($\alpha_2 = 0.21$) in order to fit the sharp differences across the observed domestic expenditure shares in final goods and inputs. A consequence of this is that the VA/GO ratio falls relative to the roundabout case and so the value-added shares γ_j need to increase in order for the model to also fit this dimension of the data. As a result, imposing symmetry in the sequential production of intermediate inputs and final goods imposes so much structure on the estimated model that the average value-added share γ_j with $N = 2$ is high at 0.88 and close to the upper bound of 1 for many countries (see Appendix Table A.II). Increasing N to three or more stages of production no longer improves the model's fit in the symmetric case because there is no further room for the value-added shares to adjust upwards in order to fit the VA/GO ratios. Given the limitations of this symmetric parameterization, hereafter we focus the empirical exercises on

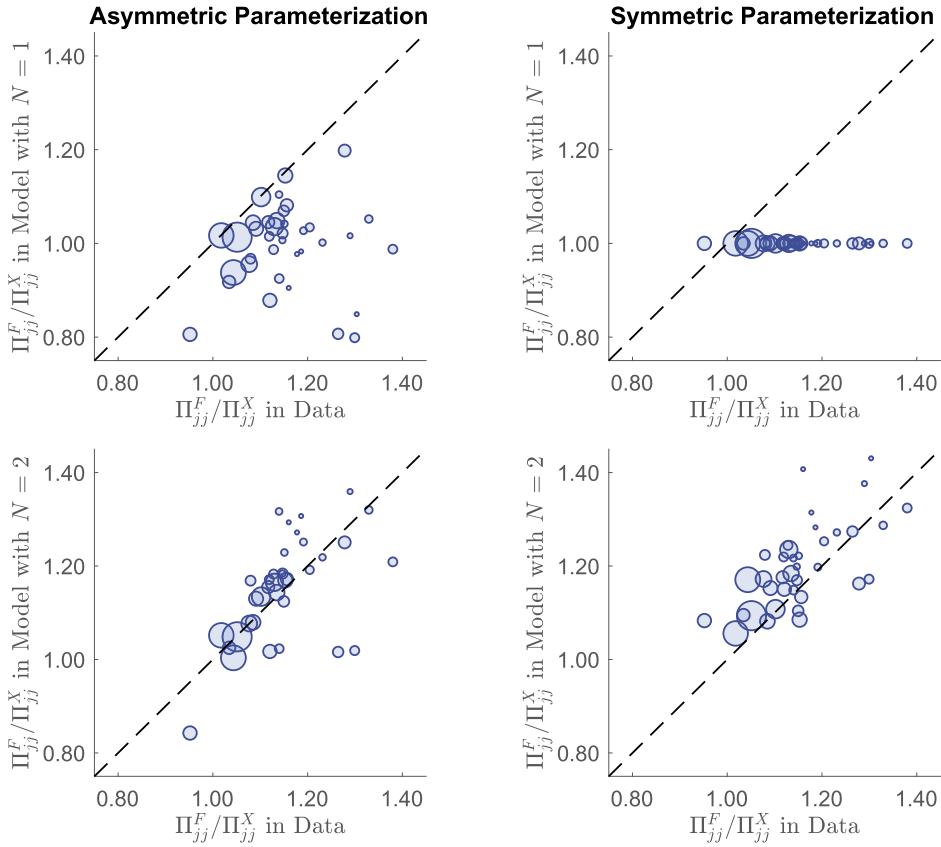


FIGURE 4.—Model fit.

the asymmetric case. Furthermore, Table III shows clearly that, for a given N , the asymmetric parameterizations fit the data much better than the symmetric parameterization.

We summarize our model’s fit in Figure 4. Recall that since we are trying to fit the total variation in the data, there are no untargeted moments to compare our estimated model to. Rather, we use Figure 4 to show how the model’s fit to the targeted moments improves when including different margins of adjustment. Since our model is fairly flexible, it can fit very accurately (in all parameterizations) both the distribution of country sizes (i.e., GDP shares) as well as the VA/GO ratio in each country. However, the heterogeneity in aggregate domestic expenditure shares Π_{jj}^F and Π_{jj}^X can only be accurately matched with the full flexibility of our parameters. Thus, we study the precision of our model’s fit through the lens of these statistics.

Specifically, Figure 4 compares the domestic expenditure share ratio Π_{jj}^F / Π_{jj}^X in the data to the estimated model with the two panels on the left representing the asymmetric parameterization and the two panels on the right representing the symmetric parameterization. The two upper panels present the estimated model fit when $N = 1$, while the two lower panels present the estimated model when $N = 2$. As discussed above, the fit with $N > 2$ in the asymmetric case improves slightly but the effect is quite marginal so we do not include these plots; in the symmetric case, the fit does not improve at all and so there is nothing to show. As can be seen, the roundabout model’s fit is much better in the asymmetric case than in the symmetric case (the top two panels), indicating that

heterogeneity across T_j^X and T_j^F is useful for matching the data. Further, the fit improves when moving from $N = 1$ to $N = 2$, indicating that the sequentiality of production is an additional margin that is useful for matching the heterogeneity in input shares observed in the data. Overall, the best fit is that shown in the southwest panel, which presents the $N = 2$ estimated model with heterogeneity in both absolute productivity T_j^X and T_j^F and input expenditure shares α_n^X and α_n^F across intermediate inputs and final goods.

6.2. Revisiting the Calibration of N

Up to now, we have fixed the number of stages N and estimated the model conditional on this parameter. While increasing N makes our estimation computationally more demanding, note that in terms of the parameters to estimate, this only amounts to estimating a longer vector of input shares α_n^X and α_n^F . Perhaps surprisingly, our structural estimation shuts down the upstream sequential stages in the production of finished final goods in our asymmetric parameterizations, while the number of stages in intermediate-input production does not appear to be larger than 2 in the symmetric parameterizations. Why does this happen?

First, note that we are estimating an average chain-length N for the *whole* world economy, including industries in which producing goods might require long or short chains. If we were to estimate our model at the industry level for various industries, we might well recover heterogeneous values of N for these industries, reflecting variation in sequentiality across industries. As discussed before, however, an estimation of a full-fledged multi-industry model poses unsurmountable computational challenges (at present).

Second, while increasing N always improves the fit of the model, as N increases, the value-added shares γ_j also need to adjust upwards to make sure our model is able to match the observed aggregate VA/GO ratio. As N increases more and more, the required adjustment in γ_j becomes less and less feasible, and the gain in the model's fit is smaller and smaller. This tight relation between the average number of stages of production and the aggregate VA/GO ratio resonates with the theoretical results in the input-output model of Fally (2012). Nevertheless, our estimated model yields a more subtle prediction: while the average number of stages of production is pinned down by the VA/GO ratio, our estimated model fits the data better when the total number of stages of production N increases as long as the volume of gross output traded at each stage adjusts accordingly.

Some readers might still object that observing no improvement in the model's fit for $N > N'$ is not synonymous with correctly identifying $N = N'$. In Web Appendix B.5, we show through simulations that there is a precise sense in which recovering the same parameters for $N \geq N'$ implies that the true N is indeed equal to N' and that $N > N'$ can be rejected. We focus on simulated economies with $J = 5$ countries and fix N at either 1, 2, or 3. For each economy, we simulate a set of primitives of the model and compute the general equilibrium. We then take the resulting simulated WIOT entries and estimate the parameters using our MPEC algorithm. Our results confirm that we are able to recover the "true" value of N .

Before concluding this section on estimation, it is important to stress that our identification of N relies heavily on our assumption that the matrix of trade costs τ_{ij} is common (and symmetric) for inputs and final goods. For example, one can show that an extension of our framework without multi-stage production (i.e., $N = 1$) could be calibrated to *exactly* match a WIOT, provided that one allows for arbitrary and asymmetric trade costs τ_{ij}^X and τ_{ij}^F for intermediate inputs and final goods. Thus, admittedly, our data cannot reject $N = 1$ if one allows enough flexibility in the modeling of trade costs.

7. COUNTERFACTUALS

Having estimated the fundamental parameters of the model, we next explore how counterfactual changes in trade costs, holding other parameters constant, alter the entries of WIOTs, thereby affecting the real income and positioning of countries in GVCs.

7.1. *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 for understanding some distinctive features of our framework. From equation (18), the loss in welfare from reverting back to autarky in a single-industry world is given by

$$\hat{W}_j^{\text{aut}} = \left(\pi_{jj}^F (\pi_{jj}^X)^{\frac{1-\gamma_j}{\gamma_j}} \right)^{\frac{1}{\theta}}, \quad (25)$$

where π_{jj}^F and π_{jj}^X are the share of final-good and intermediate-input varieties, respectively, produced entirely through domestic supply chains.⁴² In general, these shares are not observable in the data and so have to be estimated through the lens of our model. However, in the special case of roundabout production, this formula depends on the aggregate domestic expenditure shares, that is, $\hat{W}_j^{\text{aut}} = (\Pi_{jj}^F (\Pi_{jj}^X)^{(1-\gamma_j)/\gamma_j})^{1/\theta}$, and so can be implemented directly with the observed domestic shares (as in Alexander (2017)). We choose instead to follow the alternative approach of studying \hat{W}_j^{aut} in the roundabout case using our estimated model with $N = 1$ in order to be internally consistent when comparing to the estimated model with $N > 1$ (in practice, this is largely immaterial for our results).

Figure 5 presents the losses from reverting to autarky across all countries using the asymmetric parameterizations in Table III (note each panel has different y-axis). Three features are particularly salient. First, the gains from trade (or costs of autarky) are considerably larger with sequential production than in the roundabout model. On average, the losses with $N = 2$ are 60% higher than with $N = 1$. This evidence is consistent with Yi's (2003) result that deep vertical specialization magnifies the costs of protectionism. A reversal to autarky when supply chains are highly fragmented is costly since countries can no longer exploit the comparative advantage of foreign countries in the various stages of production in value chains.

Second, while increasing N could, in principle, either increase or decrease the losses from going to autarky, in practice it either does not affect the losses (for countries such as the U.S. and China) or decreases the losses (for countries such as the Netherlands and Taiwan). To see what drives these results, note from (25) that three elements account for the autarky losses. Since final goods are produced in a single stage, the estimated domestic final-good share π_{jj}^F roughly equals the observed aggregate domestic expenditure share Π_{jj}^F and so is constant across all N . On the other hand, remember from Table III that the value-added parameters γ_j are increasing with N . This force decreases the losses from trade because it implies that finished intermediate-input varieties become less important

⁴²This formula still measures the real income gains from trade in the presence of trade imbalances. The implications for real spending, however, may be quite different since autarky implies a closing of trade imbalances.

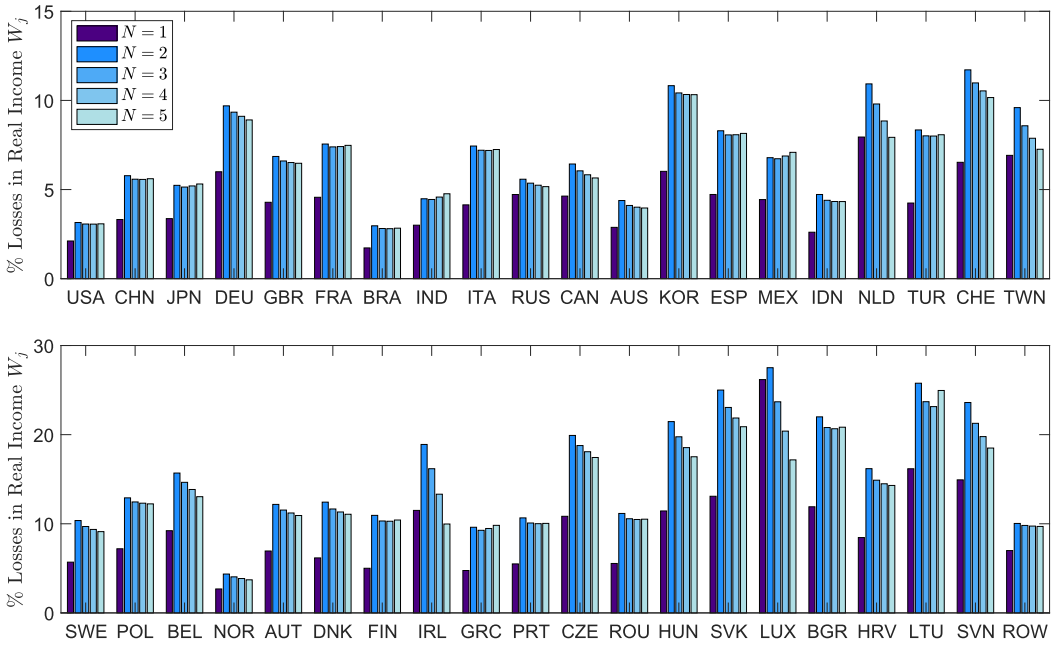


FIGURE 5.—The welfare losses from autarky.

for final consumption. In other words, leveraging international trade to deliver cheap inputs is less important when γ_j is high. Finally, as N increases, each country gets more opportunities to leverage other countries' comparative advantage through additional stages of production. Since upstream value added can be sourced from a richer combination of foreign countries, this decreases π_{jj}^X . Overall, the increase in γ_j either cancels out or dominates the decrease in π_{jj}^X : while disrupting intermediate input supply chains becomes more costly when N increases, it is also true that these supply chains become less important in terms of the value they contribute to final consumption.

As a corollary of this discussion, a third feature of our model is that the decrease in π_{jj}^X rarely dominates the increase in γ_j , and thus the gains from trade do not become unboundedly large as $N \rightarrow \infty$. As discussed in Section 4.3, while this might seem at odds with the results of Melitz and Redding (2014), this occurs because our estimation exercise targets the GDP/GO ratios observed in the data. Furthermore, while Proposition 4 generalized their results and showed that it is possible to obtain both gains from trade that become unboundedly large with a constant GDP/GO ratio, Table III and Figure 5 show that, in practice, this is not the case.

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 more pronounced. Figure 6 shows that without sequential production, that is, with $N = 1$, these gains range from 100% for the United States and 160% for China, to a staggering 2500% for Luxembourg. In contrast to the autarky case, increasing N does raise the potential gains from trade by sizable amounts for some countries: increasing N from 1 through 5 raises the U.S. gains up to 150% and the China gains up to 470%. More generally, since the potential gains increase strongly with N for rich countries, the GDP-weighted potential gains from free trade across countries also increase from 225% with $N = 1$ to 276%

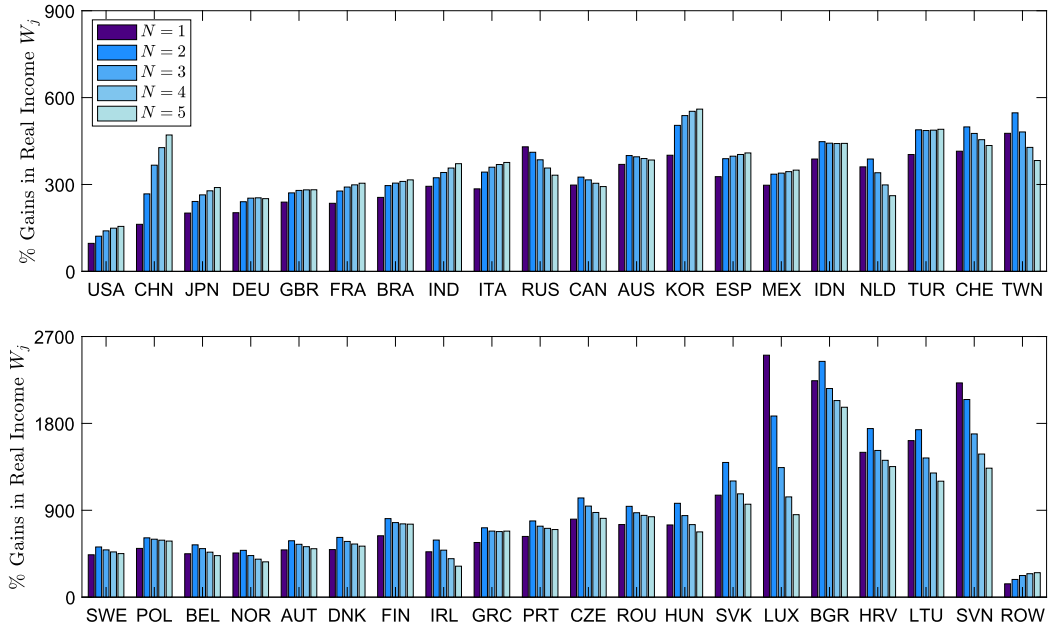


FIGURE 6.—The welfare gains from free trade.

and 317% when $N = 2$ and $N = 5$, respectively. Hence, the roundabout model vastly understates the average potential gains from free trade.

7.2. *The U.S.–China Trade War*

We now study how our model’s distinctive elements shape a real-world inspired counterfactual, namely, a trade war between the United States and China in which trade barriers increase by 25% on each other’s exports. The three panels in Figure 7 present the change in welfare, the change in U.S. intermediate-input import shares, and the change in U.S. final-good import shares. In the last two cases, we focus on the import shares from China, from three close competitors in Asian supply chains (Japan, Korea, and Taiwan), and from the other two NAFTA countries (Canada and Mexico).

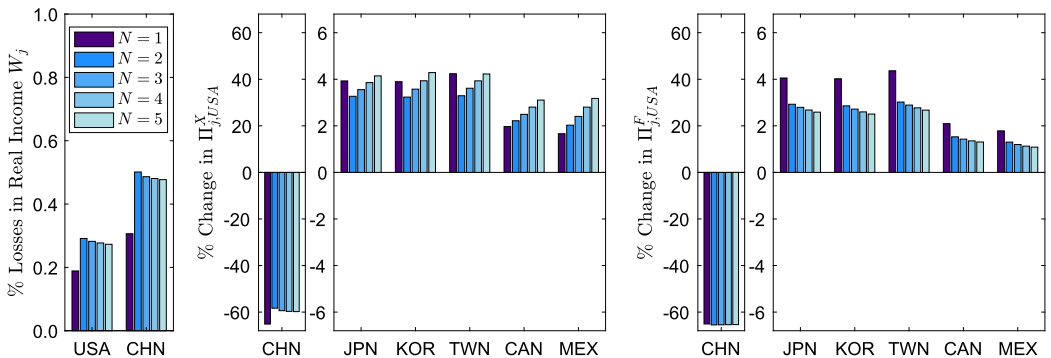


FIGURE 7.—The U.S.–China trade war.

In terms of welfare, Figure 7 reveals three features. First, both the United States and China suffer substantial welfare losses. Second, China's losses are about 50% higher than for the U.S. These results echo the patterns in Figure 5 and are driven by China being more exposed than the U.S. to international trade, and also to China being more exposed to the U.S. as a specific trade partner than the U.S. is to China. Third, the losses from a trade war are almost 50% higher in a world with sequential production, once again indicating that GVC linkages amplify the effects of trade-cost shocks.

Unsurprisingly, in terms of trade diversion, the last two panels show that China's U.S. import share falls dramatically while other countries' import share goes up. More interestingly, however, are the differential patterns observed across trading partners and across intermediate-input and final-good imports. Overall, Japan, Korea, and Taiwan are close substitutes for China in Asian supply chains and thus enjoy a strong export boom. The NAFTA countries also increase their U.S. import share, but at a smaller rate since they are competing less directly with Chinese exports.

Finally, note that the effects on intermediate-input and final-good exports differ strikingly in the context of multi-stage production. When $N = 1$, the growth of intermediate-input and final-good import shares is common across countries since there is a common trade elasticity for both types of goods. However, when $N > 1$, the export boom in intermediate inputs is increasing in N for the NAFTA countries, while the export boom of final goods is decreasing in N for all countries. Since the U.S. imports a large share of final goods from China, the fall in Chinese imports is partially offset by an increase in final-good imports from other countries. Remember, though, that the increasing trade elasticity implies that the increase in import costs is disproportionately large for final goods. This drives the U.S. to import an even larger amount of intermediate inputs from its main trade partners, Canada and Mexico, so that it can domestically finish producing final goods. This last force is strongest when sequential production is pervasive (high N) since there is more flexibility when allocating segments of the supply chain across locations. Thus, when N increases, we see a substitution from final-good export growth towards intermediate-input export growth.

7.3. Regional versus Global Integration

As a third set of exercises, we contribute to the debate arguing that much of the global fragmentation in GVCs is actually regional in nature (see Johnson and Noguera (2012b)). As should be intuitive, while the relative importance of global GVC integration monotonically increases when trade costs fall, the relative importance of regional GVC participation initially rises but eventually falls when trade costs are lowered sufficiently.

We explore the relative importance of domestic, regional, and global value chains across several trade equilibria defined by a value of Δ_τ such that $\tau'_{ij} = 1 + \Delta_\tau(\tau_{ij} - 1)$, with τ_{ij} being our calibrated trade costs for the WIOD sample in 2014. Focusing on our estimated global economy with $N = 2$, we define a domestic GVC as $\ell_d = \{\text{USA}, \text{USA}\}$ and associate the prevalence of *domestic* value chains in overall U.S. consumption with the share $\pi_{\ell_d, \text{USA}}^X$. Similarly, we capture the relative prevalence of *regional* (or NAFTA) value chains in overall U.S. consumption by $\sum_{\ell_r} \pi_{\ell_r, \text{USA}}^X$, where ℓ_r are all chains that only include the United States, Canada, or Mexico, with the exception of the chain ℓ_d . Finally, we define the relative prevalence of *global* value chains in U.S. consumption as $\sum_{\ell_g} \pi_{\ell_g, \text{USA}}^X$, where ℓ_g are all the possible chains that involve at least one country outside of NAFTA. Naturally, the sum of these three relative measures is 1. An important caveat is that, due to the use of a bundle of materials at each stage, what we label as domestic and regional value

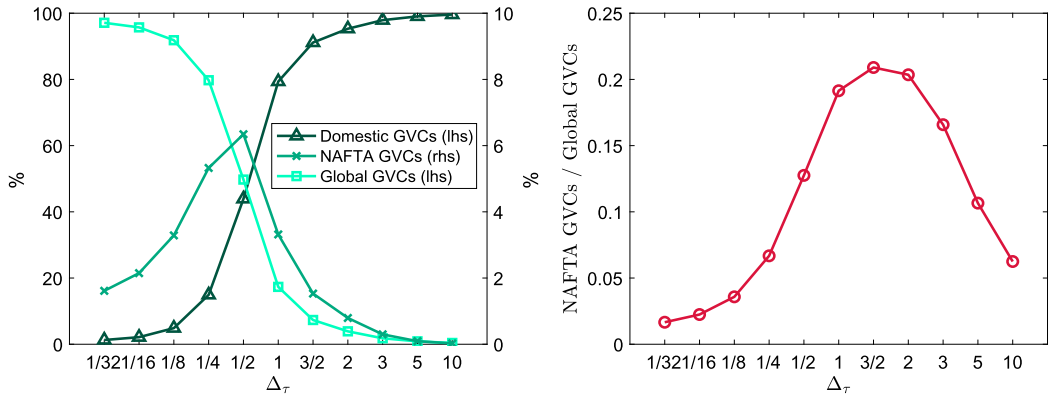


FIGURE 8.—Regional versus global integration.

chains actually embody value added from countries outside NAFTA. In fact, for bounded trade costs, our model features no purely domestic value chains. Yet, the above taxonomy is useful for understanding the broad orientation of value chains serving U.S. consumers for different levels of trade costs.

The left panel of Figure 8 plots these three measures for various values of Δ_τ between $1/32 \approx 0.031$ and 10, and provides evidence supporting our intuition for the evolution of regional and global integration. Furthermore, the right panel of Figure 8 plots the ratio of the relative importance of regional (NAFTA) versus global value chains for the same values of Δ_τ . Interestingly, our benchmark equilibrium, $\Delta_\tau = 1$, is very close to the point at which the relative importance of regional value chains is maximized. Thus, further reductions in trade costs would reduce the relative importance of regional value chains in U.S. consumption, and increase that of global value chains.

7.4. A Multi-Industry Application

As a final exercise, we provide a preliminary empirical exploration of the multi-industry, multi-stage production model with input-output linkages described in Section 5. For computational reasons, we focus on the case of two industries and aggregate the WIOD’s industries into manufacturing and services blocks. We then re-estimate our model using a combination of the multi-industry extension of our MPEC algorithm and a simulated annealing algorithm. Again for the sake of simplicity, we focus on estimating the roundabout model with $N = 1$ and the GVC model with $N = 2$.

Table IV contains our estimation results. The first row corresponds to the benchmark roundabout model, while the last four rows correspond to our GVC model. We study four different GVC estimations for the following two reasons. First, as in the single-industry case, we always find that the production of both manufacturing and services final goods requires a single stage of sequential production—that is, $\alpha_2^{M,F} = \alpha_2^{S,F} = 1$ —with the intuition being analogous to the discussion in Section 6.1. However, we also find that the loss function is relatively flat in the parameter governing the sequential production of intermediate inputs in the services industry and this leads to four multiple local minima that fit the data quite well.⁴³ This can be seen in the heatmap in Web Appendix Section B.6

⁴³We estimate the multi-industry model with only the largest $J = 11$ countries in order to tractably compute the loss function across the parameter space. This implies that the values of the optimized loss function in

TABLE IV
MULTI-INDUSTRY COUNTERFACTUALS

	Obj. F.	Estimation Results				World Weighted Average \hat{W}_j					
		$\alpha_2^{M,X}$	$\alpha_2^{S,X}$	$\alpha_2^{M,F}$	$\alpha_2^{S,F}$	$\hat{\tau}^M = 2$	$\hat{\tau}^S = 2$	$\hat{\tau} = 2$	$\hat{\tau}^M = \infty$	$\hat{\tau}^S = \infty$	$\hat{\tau} = \infty$
Roundabout	49	1.00	1.00	1.00	1.00	3.7	1.0	4.6	4.9	1.1	6.0
GVC 1	34	0.00	0.18	1.00	1.00	3.6	1.6	5.1	4.8	1.9	6.7
GVC 2	40	1.00	0.40	1.00	1.00	3.7	2.4	5.9	4.9	3.0	7.7
GVC 3	37	0.00	0.93	1.00	1.00	3.6	3.0	6.5	4.8	23.1	26.8
GVC 4	39	1.00	0.91	1.00	1.00	3.8	3.4	7.0	5.0	21.7	25.5

which plots the optimized loss function across $\alpha_2^{M,X}$ and $\alpha_2^{S,X}$. The loss function is quite flat in $\alpha_2^{S,X}$ because services intermediate inputs are relatively rare (54% of world services output is in final goods). On the other hand, the loss function is substantially nonlinear in $\alpha_2^{M,X}$ since manufactured intermediate inputs are quite prevalent (only 31% of world manufacturing output is in final goods). Second, we also compare these four GVC estimations in order to shed more light on how the model’s main mechanisms interact in a multi-industry setting.

The last six columns of Table IV present the world’s weighted-average change in welfare across six counterfactuals. The first three counterfactuals correspond to a doubling of iceberg trade costs.⁴⁴ In the first case we assume that trade costs double only for manufactures, in the second only for services, and in the third for both industries. The last three counterfactuals are analogous but impose autarky-level trade costs.

Our multi-industry counterfactuals reveal four important insights. First, as is well known, the multi-industry gains from trade are higher than the single-industry gains. Second, shutting down trade in manufactures is, in most cases, much more costly than shutting down trade in services even though services account for a much larger share of world output. This occurs because the services industry is much less open than manufactures. Third, multi-stage production in services can lead to extremely high gains from trade. While this might sound counterintuitive, explaining this further is useful for understanding the role of the model’s increasing trade elasticity. In the last two estimations, GVC 3 and GVC 4, finished services inputs are built almost exclusively with downstream production (over 90% of stage-2 value-added). This leads to a tiny upstream trade elasticity and, thus, in the trade equilibria the upstream services inputs are sourced largely from abroad. Since the tiny trade elasticity leads to a tiny domestic share, this implies that countries are benefiting highly from other countries’ comparative advantage and thus reverting back to autarky is extremely costly. Thus, a reversal to autarky in the estimations of GVC 3 and GVC 4 generates very high welfare losses that are driven mostly by services rather than by manufactures. Fourth, and relatedly, it should be made clear that these former statements are only true in the extreme case of autarky. When trade costs in services only double, the welfare losses across all GVC estimations are quite similar and the losses from manufactures dominate in all cases. Again, this is driven by the tiny upstream trade elasticity in services: While going back to autarky is extremely costly, the tiny trade elasticity also implies that these losses only kick in if trade costs change drastically.

Tables III and IV are not comparable since the bilateral flows of the “Rest-of-World” aggregate are larger in the latter case.

⁴⁴More precisely, we define the counterfactual trade costs as $\tau'_{ij} = 1 + 2(\tau_{ij} - 1)$.

8. CONCLUSION

In this paper, we have studied how trade barriers shape the location of production along GVCs. Relative to the case of free trade, trade costs generate interdependencies in the sourcing decisions of firms. Specifically, when deciding on the location of production of a given stage, firms necessarily take into account where the good is coming from and where it will be shipped next. As a result, instead of solving N location decisions (where N is the number of stages), firms need to solve the much more computationally burdensome problem of finding an optimal *path* of production. Despite these complications, we have proposed tools to feasibly solve the model in high-dimensional environments.

After deriving these results in partial equilibrium, we have developed a multi-stage general-equilibrium model in which countries specialize in different segments of GVCs. We have demonstrated that, due to the compounding effect of trade costs along value chains, relatively central countries gain comparative advantage in relatively downstream stages of production. We have also borrowed from the seminal work of [Eaton and Kortum \(2002\)](#) to develop a tractable quantitative model of GVCs in a multi-country environment with costly trade. Relative to previous quantitative models of multi-stage production, our suggested approach maps more directly to World Input–Output Tables, and allows the use of a mix of calibration and structural estimation to back out the model's parameters. We finally illustrated some distinctive features of the model by performing counterfactual analyses.

Our framework is admittedly stylized and abstracts from many realistic features that we hope will be explored in future work. For instance, a potentially interesting avenue for future research would be to introduce scale economies (external or internal) into our analysis. In a previous version of the paper, we explored a variant of our model with external economies of scale featuring a proximity-concentration tradeoff. The interaction of trade costs and scale economies substantially enriches—but also complicates—the analysis. Although our dynamic programming approach is no longer feasible in that setting, the integer linear programming approach developed in Supplemental Material Appendix A.1.2 is still quite powerful in that environment. We believe that variants of that approach could also prove useful in extending our framework to include internal economies of scale and imperfect competition. It would also be interesting to incorporate contractual frictions into our framework and study the optimal governance of GVCs in a multi-stage, multi-country environment.

Beyond these extensions of our framework, we view our work as a stepping stone for a future analysis of the role and scope of man-made trade barriers in GVCs. Although we have focused on an analysis of the implications of exogenously given trade barriers, our theoretical framework should serve as a useful platform to launch a study of the role of trade policies, and of policies more broadly, in shaping the position of countries in value chains. Should countries actively pursue policies that foster their participation in GVCs? Should they implement policies aimed at moving them to particular stages of those chains? If so, what are the characteristics of these particularly appealing segments of GVCs? These are the types of questions we hope to tackle in future research.

REFERENCES

- ADAO, R., A. COSTINOT, AND D. DONALDSON (2017): “Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade,” *American Economic Review*, 107 (3), 633–689. [1574]
- ALEXANDER, P. (2017): “Vertical Specialization and the Gains From Trade,” Bank of Canada Staff Working Paper 2017-17. [1556,1557,1578,1581,1582,1584-1586,1590]

- ALFARO, L., P. ANTRÀS, P. CONCONI, AND D. CHOR (2019): "Internalizing Global Value Chains: A Firm-Level Analysis," *Journal of Political Economy*, 127 (2), 508–559. [1557]
- ALLEN, T., AND C. ARKOLAKIS (2019): "The Welfare Effects of Transportation Infrastructure Improvements," NBER Working Paper No. 25487. [1555]
- ALVAREZ, F., AND R. E. JR. LUCAS (2007): "General Equilibrium Analysis of the Eaton–Kortum Model of International Trade," *Journal of monetary Economics*, 54 (6), 1726–1768. [1555,1563,1568]
- ANTRÀS, P., AND D. CHOR (2013): "Organizing the Global Value Chain," *Econometrica*, 81 (6), 2127–2204. [1556]
- (2018): "On the Measurement of Upstreamness and Downstreamness in Global Value Chains," in *World Trade Evolution: Growth, Productivity and Employment*. Taylor & Francis Group, 126–194. [1557]
- ANTRÀS, P., AND A. DE GORTARI (2020): "Supplement to 'On the Geography of Global Value Chains'," *Econometrica Supplemental Material*, 88, <https://doi.org/10.3982/ECTA15362>. [1557]
- ANTRÀS, P., AND R. W. STAIGER (2012): "Offshoring and the Role of Trade Agreements," *American Economic Review*, 102, 3140–3183. [1557]
- ANTRÀS, P., D. CHOR, T. FALLY, AND R. HILLBERRY (2012): "Measuring the Upstreamness of Production and Trade Flows," *American Economic Review Papers & Proceedings*, 102 (3), 412–416. [1557,1570,1574,1575]
- ANTRÀS, P., T. FORT, AND F. TINTELOT (2017): "The Margin of Global Sourcing: Theory and Evidence From U.S. Firms," *American Economic Review*, 107 (9), 2514–2564. [1565,1583]
- ARKOLAKIS, C., A. COSTINOT, AND A. RODRÍGUEZ-CLARE (2012): "New Trade Models, Same Old Gains?" *The American Economic Review*, 102 (1), 94–130. [1555,1568,1578]
- BALDWIN, R., AND D. TAGLIONI (2014): "Gravity Chains: Estimating Bilateral Trade Flows When Parts and Components Trade Is Important," *Journal of Banking and Financial Economics*, 2 (2), 61–82. [1557,1574]
- BALDWIN, R., AND A. VENABLES (2013): "Spiders and Snakes: Offshoring and Agglomeration in the Global Economy," *Journal of International Economics*, 90 (2), 245–254. [1556,1557,1577]
- BAQAEI, D., AND E. FARHI (2019): "Networks, Barriers, and Trade," NBER Working Paper No. 26108. [1557, 1563]
- BASCO, S., AND M. MESTIERI (2019): "The World Income Distribution: The Effects of International Unbundling of Production," *Journal of Economic Growth*, 24 (2), 189–221. [1558]
- BERGSTRAND, J. H., AND P. EGGER (2010): "A General Equilibrium Theory for Estimating Gravity Equations of Bilateral FDI, Final Goods Trade and Intermediate Goods Trade," in *The Gravity Model in International Trade: Advances and Applications*. New York: Cambridge University Press. [1557,1574]
- BLANCHARD, E., C. BOWN, AND R. JOHNSON (2017): "Global Supply Chains and Trade Policy," Report, Dartmouth College. [1557]
- BRANCACCIO, G., M. KALOUPTSIDIS, AND T. PAPAGEORGIOU (2017): "Geography, Search Frictions and Trade Costs," Report, Harvard University. [1554]
- CALIENDO, L., AND F. PARRO (2015): "Estimates of the Trade and Welfare Effects of NAFTA," *Review of Economic Studies*, 82 (1), 1–44. [1557,1563,1566,1578,1583,1585]
- COSTINOT, A., AND A. RODRÍGUEZ-CLARE (2015): "Trade Theory With Numbers: Quantifying the Consequences of Globalization," in *Handbook of International Economics*, Vol. 4, 197–261. [1586]
- COSTINOT, A., J. VOGEL, AND S. WANG (2013): "An Elementary Theory of Global Supply Chains," *Review of Economic Studies*, 80 (1), 109–144. [1556,1557,1571]
- DE GORTARI, A. (2019): "Disentangling Global Value Chains," NBER Working Paper 25868. [1557,1578,1581]
- (2020): "Global Value Chains and Increasing Returns." [1562]
- DEKLE, R., J. EATON, AND S. KORTUM (2008): "Global Rebalancing With Gravity: Measuring the Burden of Adjustment," *IMF Economic Review*, 55 (3), 511–540. [1568]
- DIXIT, A., AND G. GROSSMAN (1982): "Trade and Protection With Multistage Production," *Review of Economic Studies*, 49 (4), 583–594. [1556]
- EATON, J., AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70 (5), 1741–1779. [1554-1556,1558,1562-1568,1578,1579,1585,1596]
- FALLY, T. (2012): "Production Staging: Measurement and Facts," Report, University of Colorado Boulder. [1557,1589]
- FALLY, T., AND R. HILLBERRY (2018): "A Coasian Model of International Production Chains," *Journal of International Economics*, 114, 299–315. [1554,1556,1557]
- FALLY, T., AND J. SAYRE (2018): "Commodity Trade Matters," NBER Working Paper No. 24965. [1563]
- HARMS, P., O. LORZ, AND D. URBAN (2012): "Offshoring Along the Production Chain," *Canadian Journal of Economics*, 45 (1), 93–106. [1556]
- HEAD, K., AND J. 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. [1556,1583]

- HUMMELS, D., J. ISHII, AND K.-M. YI (2001): "The Nature and Growth of Vertical Specialization in World Trade," *Journal of International Economics*, 54 (1), 75–96. [1557]
- JOHNSON, R. (2014): "Five Facts About Value-Added Exports and Implications for Macroeconomics and Trade Research," *Journal of Economic Perspectives*, 28 (2), 119–142. [1557]
- JOHNSON, R., AND A. MOXNES (2019): "GVCs and Trade Elasticities With Multistage Production," Report. [1554,1557,1563,1564]
- JOHNSON, R. C., AND G. NOGUERA (2012a): "Accounting for Intermediates: Production Sharing and Trade in Value Added," *Journal of International Economics*, 86 (2), 224–236. [1557]
- (2012b): "Proximity and Production Fragmentation," *American Economic Review: Papers & Proceedings*, 102 (3), 407–411. [1593]
- KIKUCHI, T., K. NISHIMURA, AND J. STACHURSKI (2018): "Span of Control, Transaction Costs, and the Structure of Production Chains," *Theoretical Economics*, 13 (2), 729–760. [1556]
- KOHLER, W. (2004): "International Outsourcing and Factor Prices With Multistage Production," *Economic Journal*, 114 (494), C166–C185. [1556]
- KOOPMAN, R., Z. WANG, AND S.-J. WEI (2014): "Tracing Value-Added and Double Counting in Gross Exports," *American Economic Review*, 104 (2), 459–494. [1557]
- KREMER, M. (1993): "The O-Ring Theory of Economic Development," *Quarterly Journal of Economics*, 108 (3), 551–575. [1556]
- LEE, E., AND K.-M. YI (2018): "Global Value Chains and Inequality With Endogenous Labor Supply," *Journal of International Economics*, 115, 223–241. [1558]
- MAYER, T., AND S. ZIGNAGO (2006): "Notes on CEPII's Distances Measures," Report available at https://mpra.ub.uni-muenchen.de/26469/1/MPRA_paper_26469.pdf. [1574]
- MELITZ, M. J., AND S. J. REDDING (2014): "Missing Gains From Trade?" *American Economic Review*, 104 (5), 317–321. [1569,1570,1591]
- MORALES, E., G. SHEU, AND A. ZAHLER (2019): "Extended Gravity," *The Review of Economic Studies*, 86 (6), 2668–2712. [1574]
- OBERFIELD, E. (2018): "A Theory of Input–Output Architecture," *Econometrica*, 86 (2), 559–589. [1555,1565]
- OSSA, R. (2015): "Why Trade Matters After All," *Journal of International Economics*, 97 (2), 266–277. [1569]
- SANYAL, K. K., AND R. W. JONES (1982): "The Theory of Trade in Middle Products," *American Economic Review*, 72 (1), 16–31. [1556]
- SIMONOVSKA, I., AND M. E. WAUGH (2014): "The Elasticity of Trade: Estimates and Evidence," *Journal of International Economics*, 92 (1), 34–50. [1583]
- SU, C.-L., AND K. L. JUDD (2012): "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80 (5), 2213–2230. [1583]
- TIMMER, M. P., E. DIETZENBACHER, B. LOS, R. STEHRER, AND G. 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. [1582]
- TIMMER, M. P., A. A. ERUMBAN, B. LOS, R. STEHRER, AND G. J. DE VRIES (2014): "Slicing up Global Value Chains," *Journal of Economic Perspectives*, 28, 99–118. [1557]
- TINTELNOT, F. (2017): "Global Production With Export Platforms," *Quarterly Journal of Economics*, 132 (1), 157–209. [1565]
- TYAZHELNIKOV, V. (2017): "Production Clustering and Offshoring," Report, UC Davis. [1556,1561]
- YI, K.-M. (2003): "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy*, 111 (1), 52–102. [1554,1556,1557,1590]
- (2010): "Can Multistage Production Explain the Home Bias in Trade?" *American Economic Review*, 100 (1), 364–393. [1554,1557,1563,1564]

Co-editor Fabrizio Zilibotti handled this manuscript.

Manuscript received 23 May, 2017; final version accepted 21 January, 2020; available online 30 March, 2020.