Risk-Sharing with Network Transaction Costs

Christian Cox[*] Akanksha Negi[†]
University of Arizona *Monash University*

Digvijay Negi[‡]
Cornell University & Indira Gandhi Institute of Development Research

September 2, 2023

Abstract

Transaction costs can impede transfers, causing consumption to co-vary with endowment. We extend the standard risk-sharing model to include transaction costs and derive a risk-sharing test that uses consumption, production, and trade network data. We incorporate network-robust inference and demonstrate how information from the bilateral transactions can be aggregated in a sufficient way to control for them in estimation. We illustrate the method using data from the global trade in rice, wheat, and maize. Our results indicate that transaction costs impede risk sharing, particularly in maize and rice. We use the estimated model to perform counterfactual analysis, showing how changing transaction costs affects global risk-sharing and how a complete halt of wheat trade between Ukraine-Russia and major importers changes equilibrium networks.

JEL codes: D50, D81, D85, F1

[*]christiancox@arizona.edu

[†]akanksha.negi@monash.edu

[‡]dn326@cornell.edu, digvijay@igidr.ac.in

We thank Christopher Barrett, Mushfiq Mobarak, and Nicola Pavoni for comments. We also thank seminar participants at the Indian School of Business, the Econometric Society Australasian meetings, the Indian Statistical Institute, and the "100 Years of Economic Development Conference" at Cornell University.

1 Introduction

Tests of perfect risk sharing are pervasive in economics. The canonical model of risk sharing considers a social planner costlessly allocating a pool of aggregate endowment across risk-averse economic agents (Mace, 1991; Cochrane, 1991; Townsend, 1994, 1995). The model predicts that under optimality, the social planner allocates a fixed proportion of aggregate endowment to each agent. Therefore, the consumption of an agent varies only with the aggregate pool of endowments and is independent of their own endowment (Mace, 1991; Cochrane, 1991; Townsend, 1994; Obstfeld, 1994; Bardhan and Udry, 1999). This prediction of the complete risk-sharing model has been tested several times and at different levels across panels of countries, states, regions, villages, and households. Although the justification across studies for testing the complete risk-sharing hypothesis may vary, what remains common is the canonical complete markets benchmark dictated by the planner's solution.[1]

In this paper, we study the role of transaction costs in risk sharing. In the canonical model, redistribution by a social planner eliminates any positive correlation between consumption and income. Adding transaction costs to this model limit the planner's ability to fully insure agents (Obstfeld and Rogoff, 2001; Fitzgerald, 2012; Jack and Suri, 2014). We introduce costly transfers using iceberg-type trade costs popular in the trade literature. The subsequent first order conditions do not relate the consumption of an agent to an "aggregate risk" component shared by all agents. An agent's consumption instead depends on the transaction costs and network structure linking agents in the economy.

We take a network theory based approach which allows us to derive an expression that makes an agent's consumption a function of a network-specific "aggregate risk" component. Our method

[1] Even with a large body of literature testing the risk-sharing hypothesis, few have found any evidence of perfect risk sharing (Canova and Ravn, 1996; Lewis, 1996). This has led to alternative explanations for why the data shows a positive correlation between income and consumption and deviates from the complete market benchmark. One strand of literature explains this observed correlation based on information frictions leading to moral hazard and limited commitment (Ligon, 1998; Ligon et al., 2002; Genicot and Ray, 2003; Laczó, 2015; Ábrahám and Laczó, 2018; Attanasio and Krutikova, 2020; Bold and Broer, 2021; Ambrus et al., 2022). Another growing strand of literature has focused on frictions in the form of transaction costs to explain the apparent lack of risk sharing (Obstfeld and Rogoff, 2001; Schulhofer-Wohl, 2011; Fitzgerald, 2012; Jack and Suri, 2014; Clance et al., 2019; Bradford et al., 2022).

takes into account all of the trading links within a network, including between agents who do not directly trade but are indirectly connected through other trading partners. This gives rise to a new risk-sharing estimation method that we implement on the global trade of three major food commodities — rice, wheat, and maize. We show how the information from these network frictions can be aggregated to sufficiently control for them, and ignoring these bilateral frictions in the aggregate analysis can bias estimation. In our application, the aggregate variable is consumption at the country-year level, and the bilateral variable is the trade linkage between two countries. Our framework is micro-founded from a social planner's problem, and is applicable to any environment with bilateral interactions between agents.

Using simulations from the model, we first establish some basic results. We show that a reduction in transaction costs leads to: (i) greater linkages between agents and more connectedness in a network, (ii) reduced consumption inequality across agents, (iii) lower correlation between domestic consumption and endowment, and (iv) reduced consumption variation and higher global utility. Similar results are reported by Jack and Suri (2014) in a three agent risk-sharing model with fixed transaction costs. We show that these results hold in our generalized framework with iceberg costs and with an arbitrary number of players. Next, we derive a structural test of risk sharing which explicitly takes the network structure of transfers into account. Finally, we solve the estimated model and study two relevant issues in global food trade. First, we show how counterfactual changes to transaction costs affect global risk-sharing, and second, we illustrate how the Russia-Ukraine conflict affects the global wheat trade network.

A unique contribution of this paper is the introduction of a network approach to the analysis of risk-sharing models. While other studies have introduced transaction costs in risk-sharing analysis, we are the first to characterize and estimate a generalized model. Bradford, Negi, and Ramaswami (2022) also introduce transaction costs in the risk-sharing model but do not directly incorporate network structure in estimation. We differ from their analysis in multiple ways. First, we generalize their model and fully characterize how transaction costs affect the trade network for any number of agents. Second, we directly incorporate trade networks in estimation and use the structural model for counterfactual predictions. Schulhofer-Wohl (2011) allows for the possibility of costly transfers in the standard social planner's allocation problem. In Schulhofer-Wohl's formulation, however, transaction costs do not vary at the trading pair level. We relax this assumption in our model. Fitzgerald (2012) embeds a gravity model of trade in intermediate goods in a standard DSGE model to propose a gravity based test of risk sharing. Fitzgerald's approach induces dependency across countries which is distinct from our case of perfect substitution; this difference has consequences on network formation and the role of transaction costs in risk sharing. Finally, Laczó (2015) structurally estimates a risk-sharing model, but with frictions that arise in the form of limited commitment.

This paper also relates to the literature on specification issues in risk-sharing tests. A recent strand of literature has established the omission of heterogeneous risk preferences as an important source of bias (Schulhofer-Wohl, 2011; Mazzocco and Saini, 2012; Asdrubali et al., 2019). Likewise, studies have suggested that risk sharing may happen within networks which may or may not span the entire universe of agents (Fafchamps and Lund, 2003; De Weerdt and Dercon, 2006; Bramoullé and Kranton, 2007; Attanasio et al., 2012; Attanasio and Krutikova, 2020). The implication is that aggregate shocks should be network specific (Ambrus et al., 2014, 2022). Adding to this literature, we discuss how ignoring transaction costs and the implied network structure can lead to biased estimates of risk sharing. Incorporating networks into estimation also introduces a complication in the form of network dependence in the error structure. We allow errors in the risk-sharing regression to exhibit such dependence and use results from Kojevnikov et al. (2021) to propose a network heteroskedasticity and autocorrelation robust variance estimator for conducting valid inference on the risk-sharing parameters. The standard errors robust to general network dependence are presented along with the results.

The rest of the paper is laid out as follows. Section 2 presents the augmented model of risk sharing with transaction costs. Section 3 presents simulation based results from the model. Section 4 presents the structural risk-sharing test specification and discusses its connection with the conventional tests. Section 5 presents the empirical application in the global trade network of three commodities alongside counterfactual analysis. Finally, section 6 concludes.

2 Risk Sharing Model with Transaction Costs

In the following section, we modify the standard model to accommodate bilateral transactions costs with transfers among agents. We model transfers as costly bilateral trade flows between agents and solve a social planner's solution to this problem. We study the optimality conditions of the social planner's problem to illustrate the connection between an agent's consumption and their bilateral transfers. We then re-frame the model in terms of formal networks and use this setup to precisely characterize how all agents within a network are indirectly linked to each other. We then illustrate the empirically relevant aspects of the model.

2.1 Model Setup

The world consists of T time periods $t = 0, ..., T$, with a state of the world per period s_t drawn from a known discrete distribution $s_t \sim D_t$. The realization of each state affects the primitives of the model, and thus each is state contingent.

Consider N agents (countries, villages, households, etc.), each indexed with i. Each agent has a commonly known endowment per period (realized after the state) $y_{it}(s_t)$ of a divisible homogeneous good with which they can trade; let y^E_{ijt} be exported endowment (from agent i to agent j) and y^D_{it} be domestic endowment (part of endowment not exported and consumed instead). Thus,

$$y_{it} = y^D_{it} + \sum_{j \neq i}^{N} y^E_{ijt} \tag{1}$$

The state contingent global endowment in period t is then $Y_t(s_t) \equiv \sum_i y_{it} = \sum_{i=1}^N y^D_{it} + \sum_{i=1}^N \sum_{j \neq i}^N y^E_{ijt}$. Likewise, we define total consumption in period t for an agent as:

$$c_{it} = c^D_{it} + c^I_{it} \tag{2}$$

where c_{it}^D is the consumption from domestic endowment and c_{it}^I is the consumption from imports. We abstract away from storage and intertemporal trade as a means to smooth consumption to focus on frictions only in the form of transaction costs. Since there is no storage in this model:

$$c_{it}^D = y_{it}^D \tag{3}$$

Now suppose that imports are costly, meaning that one exported unit from agent j ends up being less than one unit received for importing agent i. The fraction with which exports are converted into imports is allowed to vary by each trading pair. Suppose that for each unit of exports from agent j to agent i, only a fraction $\delta_{jit} \in [0,1]$ is converted into imported consumption:

$$c_{it}^I = \sum_{j \neq i}^{N} y_{jit}^E \delta_{jit} \tag{4}$$

The term $1 - \delta_{jit}$ captures the loss from shipping costs (in terms of the good), "lost at sea" cargo, and wasteful rents extracted by some middleman; all of these prevent the full exported amount from being consumed by importers.[2]

The social planner's program is to maximize the ex-ante lifetime utilitarian social welfare function with discount β^t, per period agent utility $u_i(c_{it})$, and per period Pareto weights α_{it} defined such that $\sum_{i=1}^N \alpha_{it} = 1$. They choose the endowment allocation per agent $(y_{it}^D(s_t), y_{ijt}^E(s_t) \;\; \forall j \neq i)_{\forall i}$ for each period and state, with constraints for agent i governed by the endowments: $c_{it} = y_{it}^D + \sum_{j \neq i}^N y_{jit}^E \delta_{jit}$, $y_{it} = y_{it}^D + \sum_{j \neq i}^N y_{ijt}^E$, and $y_{it}^D, y_{ijt}^E \geq 0$. To simplify the program, we can express it just in terms of exports and total endowments, with an ex-ante expectation $E_0[\cdot]$:[3]

$$\max_{\substack{y_{jit}^E(s_t) \geq 0 \\ \forall s_t \forall t, \forall i, j \neq i}} \sum_{i=1}^N E_0 \left[\sum_{t=0}^T \alpha_{it} \beta^t u_i \left(y_{it}(s_t) - \sum_{j \neq i}^N y_{ijt}^E(s_t) + \sum_{j \neq i}^N y_{jit}^E(s_t) \delta_{jit} \right) \right] \tag{5}$$

[2]In this model, global consumption can be less than global endowment per period: $C_t \equiv \sum_i c_{it} = \sum_i (c_{it}^D + c_{it}^I) = \sum_{i=1}^N (y_{it}^D + \sum_{j \neq i}^N y_{jit}^E \delta_{jit}) \leq Y_t$. Also, these iceberg costs are distinct from fixed costs; the optimal number of links and the amount exchanged per link would change. Fixed costs would likely lead to higher inequality as more agents would be left out of the network as opposed to just having a smaller inflow of trade.

[3]Note that we do not have the ceiling inequality constraint on exports; when you receive imports you can in fact export more than your own endowment.

This setup generalizes the standard risk-sharing model. Let $P_0(s_t)$ be the probability of state s_t. The system of first order conditions (FOC) is as follows, with modified Lagrangian multiplier for the non-negativity constraint $\beta^t P_0(s_t)\lambda_{ijt}(s_t)$ (following Schulhofer-Wohl (2011)):

$$\begin{cases} \alpha_{jt}\frac{\partial u_j}{\partial c_{jt}}\delta_{ijt} - \alpha_{it}\frac{\partial u_i}{\partial c_{it}} - \lambda_{ijt} = 0 \\ \lambda_{ijt}y_{ijt} = 0, \lambda_{ijt} \geq 0, y_{ijt} \geq 0 \end{cases} \quad \forall s_t \forall t \cap \forall i, j \neq i \tag{6}$$

Since our focus is on an empirical setting with a known trade network, we assume knowledge of the set of trading partners for deriving the empirical expressions; we know which $y_{ijt} = 0$.[4] Then, for a known set of trading partners and CRRA utility $u(c_{it}) = \frac{c_{it}^{1-\gamma}}{1-\gamma}$, we can derive a closed-form solution for a given realization of the state variable; we drop the explicit conditioning on s_t for notational ease. The FOC system in equation (6) yields the linear system of equations (7).[5] An important parameter in this system that captures how trading partners are linked is the direct trade weight $\Delta_{ijt} \equiv \left(\frac{\delta_{ijt}\alpha_{jt}}{\alpha_{it}}\right)^{-1/\gamma}$.

$$\Delta_{ijt}\left(y_{jt} - \sum_{k\neq j}^{N} y_{jkt}^{E} + \sum_{k\neq j}^{N} y_{kjt}^{E}\delta_{kjt}\right) = \left(y_{it} - \sum_{k\neq i}^{N} y_{ikt}^{E} + \sum_{k\neq i}^{N} y_{kit}^{E}\delta_{kit}\right) \quad \forall i \text{ exporting to } j \tag{7}$$

The system of equations in (7) can be solved for optimal exports, which then allows us to calculate optimal consumption c_{it}^*. In Appendix A, we show that the optimal consumption per state-period is linear in endowments, meaning $c_{it}^* = \sum_k y_{kt}\omega_{kt}^{it}$, where the weight $\omega_{kt}^{it} \geq 0$ is an agent-k-specific function of parameters $(\Delta_{ijt}, \delta_{ijt})_{\forall i, j \neq i}$. Using this and equation (7) rewritten as $\Delta_{ijt}c_{jt} = c_{it}$, we can then express agent j's optimal consumption in terms of i's solution:

$$c_{jt}^* = \frac{1}{\Delta_{ijt}} \cdot \left(\sum_k y_{kt}\omega_{kt}^{it}\right) \quad \text{if } i \text{ exports to } j \tag{8}$$

Thus if two agents trade, their consumption share a common element that is scaled by the direct

[4]For our simulations, we evaluate all possible trading network shapes, required to solve the the full system in (6).

[5]Note that if there are heterogeneous preferences ($\gamma_i \neq \gamma_j$), then the FOC cannot be linearly transformed.

trade weight. Now consider an agent k that shares a trading partner with i (say k also exports to j) but i and k do not trade. In this case, there is no FOC directly pairing i and k, but they are indirectly linked via an intermediary (agent j). Define the indirect trade weight $\Delta_{i\to k,t}$ for two agents that are not directly linked: $\Delta_{i\to k,t}c^*_{kt} = c^*_{it}$, where $\Delta_{i\to k,t} = \Delta_{ikt}$ if the two are directly linked.[6] To finish characterizing the model and derive a general formula for $\Delta_{i\to k,t}$, we utilize network theory.

2.2 Trade Networks

The extensive part of our trade model, meaning whether two agents have any trade relationship, can be thought of as a directed graph or network. The graph represents each agent as a node (or vertex), and if two agents trade, they are connected to each other with an edge (or link). Since trade can be asymmetric, the graph is directed. If there is an undirected path connecting all agents in our directed graph, then our network is weakly connected and trade is "global". Otherwise there are "trading islands", called components: connected induced-subgraphs that are disjoint from other components. We will refer to each (weak) component of a disconnected graph as the "trade network", and describe the network within the subgraph. Thus, the universe of agents is the graph and each "trade network" is a subgraph, with autarkic agents being isolated nodes within the graph which are not connected to any subgraph. Define the trade network in period t by the ordered pair $\mathcal{T}_t = (\mathcal{N}_t, \mathcal{S}_t)$, where $\mathcal{N}_t = \{1, 2, \ldots, N\}$ is the set of nodes and $\mathcal{S}_t$ is the shape or adjacency matrix of directed links.

To link the previous section with the current discussion, first fix an exporter $\bar{i} \in \mathcal{N}_t$ within $\mathcal{T}_t$, called the index agent for periot t. Then, using equation (8), we can express c^*_{jt} $\forall j \in \mathcal{N}_t$ in terms of $c^*_{\bar{i}t}$ with a common term within network $\mathcal{T}_t$: $\mathcal{A}^*_{\mathcal{T}_t} \equiv c^*_{\bar{i}t} = \sum_{k\in\mathcal{N}} y_{kt}\omega^{\bar{i}t}_{kt}$, as shown in equation (9). We call $\mathcal{A}^*_{\mathcal{T}_t}$ the aggregate risk function, a concept originating from Bradford et al. (2022), which we generalize and integrate into the estimation.

$$c^*_{jt} = \frac{\mathcal{A}^*_{\mathcal{T}_t}}{\Delta_{\bar{i}\to j,t}} \quad \text{if } \bar{i}, j \in \mathcal{N}_t \tag{9}$$

[6] In the example, i is linked to k via $c_i = c_k(\frac{\delta_{ij}\alpha_k}{\delta_{kj}\alpha_i})^{-1/\gamma}$, and $\Delta_{i\to k} = \Delta_{ij}/\Delta_{kj}$. In general $\delta_{ik} \neq \frac{\delta_{ij}}{\delta_{kj}}$.

Thus a network member's consumption is just the ratio of a network-specific function of endowments and parameters and an agent-specific function of parameters.[7] Any consumption risk that is based on an endowment shock is shared at the network level (rather than agent level), with heterogeneity in consumption being driven solely by exogenous parameters such as preferences, Pareto weights, and transaction costs. Furthermore, this risk is shared by everyone throughout the entire network, not just with one's direct trading partners.

Our main empirical specification is based on equation (9), which requires calculating the indirect trade weight; we derive a formula for the shortest-path $\Delta_{\bar{i}\to j}$ utilizing the observed network shape $\mathcal{S}$, dropping the time subscript to reduce clutter. Let $\mathcal{P}_{ik}$ be the collection of links for the geodesic (shortest) undirected path from i to k, meaning the minimum link-distance from node i to node k from the matrix $\tilde{\mathcal{S}} = \mathcal{S} + \mathcal{S}'$; in our case, this symmetric matrix is exactly the undirected adjacency matrix.[8] Thus $\mathcal{P}_{ik}$ is a sequence of elements of $\tilde{\mathcal{S}}$ along the undirected path.[9] For example, to get from i to k, suppose the smallest distance using $\tilde{\mathcal{S}}$ goes through links (i,j), (j,l), and (l,k). The indirect trade weight $\Delta_{i\to k}$, shown in equation (10), is a function of all of the transaction terms δ along such a path. How these intermediary trade costs affect the end result is based on whether the links are export links ($\mathcal{S}(i,j) = 1$) or import links ($\mathcal{S}'(i,j) = 1$).

$$\Delta_{i\to k} = \prod_{(r,s)\in\mathcal{P}_{ik}} \frac{(\Delta_{rs})^{\mathcal{S}(r,s)}}{(\Delta_{sr})^{\mathcal{S}'(r,s)}} = \left(\frac{\alpha_k}{\alpha_i} \prod_{(r,s)\in\mathcal{P}_{ik}} \frac{(\delta_{rs})^{\mathcal{S}(r,s)}}{(\delta_{sr})^{\mathcal{S}'(r,s)}} \right)^{-1/\gamma} \tag{10}$$

The exponent terms are elements of the adjacency matrices and are either 1 or 0. The first [second] intermediaries (in each pair) who export [import] are in the numerator and the opposite intermediaries are in the denominator. See proof of Proposition 2 in Appendix A. Note that $\mathcal{P}_{ik}$ are in general not unique; when calculating them, we order agents by their closeness to the index such that the path chosen is the geodesic path that passes through the most central agents.

[7]Note that it is superfluous to say "agent-network-specific" because if an agent is connected to two separate "networks", then in fact they are all one network.

[8]No symmetric trade is an outcome of our model and so there is no overlap between $\mathcal{S}$ and $\mathcal{S}'$; thus we do not need to divide such entries by 2, which would be required in the general case.

[9]A directed path only follows arrows in their specific direction; for our purposes, the indirect links that "connect" two non-trading partners are undirected via the first order conditions. The direction determines the order of δs.

This formulation comes directly from the first order conditions. Any first order condition can be rewritten in terms of others through substitution; if two agents do not trade, there is no direct way of comparing their consumption via a trade weight; however one can indirectly compare them with pairwise comparisons of each trading pair along any path that connects the two non-trading agents. The geodesic path is chosen as it has the natural interpretation of being the most efficient path in terms of number of links. However, one could use an alternative algorithm like minimum cost by calculating the path that maximizes the sum of δs towards the index. Those two paths are not guaranteed to be unique, but the existing trade network is based on δs, so the shortest path and minimum cost path will be similar in equilibrium. Solving for the minimum cost path is costly as the number of combinations to check is high, but we find that the shortest distance does cross through agents that have relatively low transaction costs compared to neighbors not directly along the path; this indicates high overlap between the two types of paths in reality.

The intuition behind the indirect trade weight is that an agent's degree of centrality within a network is not only dependent on the efficiency of their trading, but also on how isolated it is by the trade efficiency of its neighbors. This iterative approach allows one to characterize every agent's trade efficiency with respect to every other agent even if they do not directly trade in equilibrium; there are indirect transfers via intermediary agents that are not captured by looking at direct bilateral transfers. This is a novel measure based not only on the network shape, but also on the transaction costs that inform the equilibrium shape.

3 Simulations

How do trade flows and the overall network change as transaction costs and endowment shocks change? First, we consider the effects of a decrease in costs by increasing trade efficiency δ. The model of Jack and Suri (2014) is a simplified version of our model with lump-sum transaction costs; they find that a cost reduction results in better consumption smoothing across shocks, increased number of transactions, and increased number of network participants. We show that this result holds in our generalized framework with iceberg costs and with a larger number of players.

In the example shown in Figure 1, the endowments are $[42, 81, 55]$ and all countries have identical transaction costs and Pareto weights. With high transaction costs, there is no trade. With a slightly lower cost, agent 2 exports to 1. With a further decrease in cost, agent 2 increases its amount to 1, and eventually exports to agent 3. In a larger example of five countries with endowments $[100, 50, 40, 90, 72]$, we see the trade networks growing and shifting as more countries start to trade as costs continually decrease. In terms of consumption smoothing, Figure 2 shows that as transaction costs decrease, the social planner is able to smooth the endowment shocks such that consumption equalizes across countries. Increasing trade flows decreases inequality from redistribution, but also decreases total consumption due to transaction costs.

It is also interesting to see how adding new trading partners changes global risk sharing differentially. Consider 5 countries with endowments, $[100, 20, 20, 100, 22]$, all with equal Pareto weights and equal (small) transaction costs. In this case, trade is efficient from the high endowment to low endowment countries. Consider five different trading scenarios displayed in Figure 3. First we have full autarky, then bilateral trade between 1 and 2, then additional isolated trade between 4-5, followed by letting 1 trade with 3 as well, and finally, connecting all five countries by having 4 trade with 3. How do these different scenarios affect welfare?

In Figure 4 we show how the previous trade scenarios affect the standard deviation of consumption [on the left] and global utility [on the right]. Not surprisingly, increasing trade flows reduces consumption inequality and promotes global utility. This is even the case when going from scenario 4 to scenario 5: note that in this scenario, all countries are still trading in scenario 4 but in scenario 5, the two distinct trading islands are linked by agent 3. Even in this case, risk sharing is improved. Thus a single linkage between two disparate networks will lead to risk sharing across all countries involved. Note however that the effect is smaller than when adding a completely autarkic agent to a network: the slope on utility increasing from scenario 4 to 5 is smaller than the rest. Again this is not surprising as we are simply allowing for smoother risk sharing.

Next, we consider the number of trading interactions and players as the transaction costs decrease; this is shown in Figure 5. Both the number of players in the network and the number of total trading links increase. Finally, we consider the correlation between consumption and production, the main measure of risk sharing. As costs decrease, this correlation decreases, as shown in Figure 5.

4 Identification and Estimation

Recall that our goal is to estimate risk sharing controlling for transaction costs and network structure. We can use our model to derive such a test. We discuss the main risk-sharing regression, the bias associated with ignoring transaction costs, model extensions, and network dependent inference.

4.1 Risk-Sharing Equation from Model

Consider the optimal consumption for time period t as defined in equation (9). Assuming that the observed consumption c_{jt} is measured with exponential multiplicative (zero mean) error compared to the model consumption: $c_{jt} = \exp(\varepsilon_{jt})c^*_{jt}$ and taking logs, equation (9) can be expressed as:

$$\ln(c_{jt}) = -\ln(\Delta_{jt}) + \ln(\mathcal{A}^*_{\mathcal{T}_t t}) + \varepsilon_{jt} \tag{11}$$

The first component $\Delta_{jt} \equiv \Delta_{\bar{i}_t \to j,t}$, is an agent-year specific parameter (relative to an index agent which changes by network and year, denoted with $\bar{i}_t$) and is a function of transaction costs, Pareto weights, and preferences. In particular, it contains information about agent j and the index agent $\bar{i}$. If j does not trade directly with $\bar{i}$, then Δ_{jt} also contains information on all agents that indirectly link j and $\bar{i}$ based on the network shape $\mathcal{S}_t$. The indirect trade weight is indexed by j and t as it will be different for each agent based on its position in the network and will also change with time since we allow networks to change overtime.

The second component $\mathcal{A}^*_{\mathcal{T}_t t}$ is the aggregate risk function, which is the common element to consumption allocation of all agents in a given network, and thus varies at the network-year level.

Network $\mathcal{T}$ is indexed by time, $\mathcal{T}_t$, to emphasize the fact that networks can change every period; both membership $\mathcal{N}_t$ and shape $\mathcal{S}_t$ affect consumption each year.

The empirical implication of the model is that, conditional on indirect trade weight and the aggregate risk function, consumption should be independent of own endowment. To arrive at a test, we follow the literature and add endowment as a covariate in equation (11).

$$\ln(c_{jt}) = -\ln(\Delta_{jt}) + \ln(\mathcal{A}^*_{\mathcal{T}_t t}) + \eta \cdot \ln(y_{jt}) + \varepsilon_{jt} \tag{12}$$

Equation (12) defines a misspecification test of our model. Under perfect risk sharing, $\eta = 0$. If $\eta \neq 0$, consumption is correlated with production and therefore the complete risk sharing hypothesis is rejected. Under that scenario, $(1-\eta)$ gives us a measure of the degree of risk sharing achieved in the network (Asdrubali et al., 1996; Crucini, 1999; Crucini and Hess, 1999; Asdrubali et al., 2020; Bradford et al., 2022). It is important to reiterate that Δ_{jt} and $\mathcal{A}^*_{\mathcal{T}_t t}$ are functions of parameters and hence will be estimated simultaneously with η.[10] We also add a constant to equation (12) to better justify the zero-mean assumption for the error term ε_{jt}. As demonstrated in Schulhofer-Wohl (2011), it is more realistic to have agent specific preferences γ_i instead of a constant γ. In Appendix C.1 we discuss and estimate such a specification.

The canonical risk-sharing regression is equivalent to the model with $\delta_{ijt} = 1$ for all trading pairs. In Appendix B.1, we give a detailed discussion of how the regressions compare in terms of potential bias and interpretation. Given that the network shape is an equilibrium function of the endowments, and Δ is a nonlinear function of the shape and transaction costs, the bias does not have a straightforward expression. We illustrate the bias with a simulated example.[11] For the countries in each network per year, we regress the consumption on production, initially alone, and then control for the aggregate risk function and transaction costs. The results are presented in Table 1. The coefficient on production is akin to the risk sharing test; controlling for the aggregate

[10] Recall that y_{it} is contained in $\mathcal{A}^*_{\mathcal{T}_t t}$ and so separate identification of η relies on nonlinearity of $\ln(\mathcal{A}^*_{\mathcal{T}_t t})$ in $\ln(y_{jt})$; this is satisfied as $\mathcal{A}^*_{\mathcal{T}_t t}$ is linear in network participant endowments (see Appendix A). This restriction is implicit in other risk-sharing tests as well when controlling for the aggregate shock.

[11] We simulate the model for a 4 country case with homogeneous Pareto weights and fluctuating transaction costs (random uniform draws averaged around 0.6) across 50 time periods.

risk slightly reduces the coefficient and significance, but controlling for transaction costs causes the production coefficient to be small and noisy. Thus, depending on the context, controlling for transaction costs can change the risk sharing prediction.

4.2 Estimation

The estimation steps are as follows. First we calculate the network shape $\mathcal{S}_t$ for each year using the observed set of bilateral trading partners. We set the network to be undirected and calculate $\mathcal{S}_t$ for each isolated network.[12] For each isolated network, we define the index based on the highest betweenness-centrality measure for all countries in the network. We then calculate the geodesic (shortest) path between all countries in the network and the index. We generate the path of countries that links each country to the index. This allows us to calculate the indirect trade weight for each country j as Δ_{jt}, with the added step of noting which country in each link along the path is an exporter or importer.

To estimate equation (12) we use non-linear least squares with Logistic transaction costs, $\delta_{ijt} = \exp(X_{ijt}\beta)/(1+\exp(X_{ijt}\beta))$. To calculate $\mathcal{A}^*_{\mathcal{T}_t t}$, one can use the index country's consumption or solve the model for each ω_{it} (see Section 2.2). If the data allows, one can also capture $\mathcal{A}^*_{\mathcal{T}_t t}$ with a network-year indicator, and we use this approach in our application as we have a global network structure every year.[13] We set the Pareto weights α_{it} to each country's population, scaled to sum to 1. For a given parameter vector: $(\gamma, \beta, \eta, T_t\ \forall t)$, we calculate Δ_{jt} for all countries and years. Then we plug Δ_{jt} into equation (12), calculate $\sum_{jt} \varepsilon^2_{jt}$, and iterate using modified Newton's method.

The transaction cost network terms capture time-varying changes that affects countries directly and indirectly; for example, the Green Revolution affected grain production capacity in some countries more than others, but this also affected the overall network structure, which is captured in the time-varying Δ and $\mathcal{A}$. While the indirect trade weight, combined with the aggregate risk function, are in some sense "sufficient" to characterize consumption within a network, the specification of

[12] We utilize the *nwcommands* package in Stata (Grund, 2015), which allows the user to format data into a network and calculate various network properties.

[13] In Appendix B.2 we show how to incorporate non-trading countries.

the individual δ are important. We also include time trends and individual fixed effects.

4.3 Network Dependent Inference

Given the network structure of trade, we allow the errors (ε_{it}) in the risk sharing equation to exhibit general network dependence.[14] This affects inference since the usual heteroskedasticity and autocorrelation robust (HAC) standard errors, which assume independence between cross sectional units, will be incorrect for the parameters of the regression equation in (12). To account for network dependence, we use the results in Kojevnikov et al. (2021) (KMS thereafter) to arrive at a consistent network-HAC variance estimator. The latter allows covariances between any two nodes to decay relatively fast as a function of the network-distance, $D_N(i,j) = d \in \mathbb{R}$, which is defined on $\mathcal{T}_N$ to be the length of the shortest-path between nodes, $i, j \in \mathcal{N}_N$.[15]

With network dependent data, the limit theorems needed to obtain asymptotic results including the variance-covariance matrix only hold if the network is not too dense. In particular, we need sufficient conditions which restrict not only the strength of network dependence between nodes but also the number of neighbors at a given distance, d, which depends on the network shape. To limit the strength of dependence, we assume that $\boldsymbol{\varepsilon}_i = (\varepsilon_{i1}, \ldots, \varepsilon_{iT})' \in \mathbb{R}^T$ is conditionally ψ-dependent given $\mathcal{C}_N$ with dependence coefficients given by $\lambda_N = \{\lambda_{N,d}\}_{d \geq 0}$, $\lambda_{N,0} = 1$. The notion of ψ-dependence is the primary way in which KMS measure the strength of covariance between two sets of nodes that are at least d-distance apart. The sequence of dependence coefficients λ_N provide an upper bound on such covariance terms where this sequence approaches zero as d increases.

The exact conditions needed to argue that the Law of Large Numbers and the Central Limit Theorem apply and consistent estimation of the variance is possible are given in Appendix B.3. These conditions have been adapted from KMS to fit the current setup.

Then given the regression equation in (19) of the appendix, a consistent estimator for the net-

[14]A motivation for using network dependent robust standard errors would be a formal statistical test. Network dependence tests exist for specific functional forms or settings not studied in our paper (Liu and Prucha 2018; Lee, Shen, Priebe, and Vogelstein 2019; Su, Lu, Song, and Huang 2019; Fischer and LeSage 2020).

[15]The indexing by N is to emphasize the fact that asymptotic theory follows due to the size of the network getting large, i.e. $N \to \infty$.

work HAC variance is then given by $\hat{\mathbf{Q}}_T\hat{\mathbf{\Omega}}_T\hat{\mathbf{Q}}_T$ where the middle (sandwich term) is constructed to be a weighted average of the covariances of node i with its neighbors, j, who are exactly d-*distance* away denoted by the set $\mathcal{N}_N^{\partial}(i,d)$. Formally, $\mathcal{N}_N^{\partial}(i,d) = \{j \in \mathcal{N}_N : D_N(i,j) = d\}$.

$$\hat{\mathbf{\Omega}}_T = \frac{1}{NT^2}\sum_{d\geq 0}\mathcal{K}_N(d)\left[\sum_{i\in\mathcal{N}_N}\sum_{j\in\mathcal{N}_N^{\partial}(i,d)}\mathbf{h}_i^{(1)}(\hat{\boldsymbol{\theta}})\hat{\boldsymbol{\varepsilon}}_i\hat{\boldsymbol{\varepsilon}}_j'\mathbf{h}_j^{(1)}(\hat{\boldsymbol{\theta}})'\right] \tag{13}$$

where $\mathbf{h}_i^{(1)}(\hat{\boldsymbol{\theta}})$ is the jacobian of $\mathbf{h}_i(\hat{\boldsymbol{\theta}})$ (see equation (19) in the appendix) and $\mathcal{K}_N(d) = \mathcal{K}(d/b_N)$ is the kernel used to weight the sample covariances as a function of the network distance. In particular, we use the Bartlett kernel i.e. $\mathcal{K}(x) = (1-|x|)\cdot\mathbb{1}\{|x| < 1\}$. The term b_N is the bandwidth parameter which depends on N and is used to truncate how many such covariance terms appear in the double sum. In our case, we use $b_N = \lfloor (N)^{1/3}\rfloor$ where $\lfloor\cdot\rfloor$ is the floor function.

Another thing to note is that while KMS do not impose any restriction on $\mathbf{\Omega}_T$, we assume that the off-diagonal elements of $\mathbf{\Omega}_T$ are zero. These terms correspond to $\mathbb{E}(\varepsilon_{it}\varepsilon_{jt'}|\mathbf{w}_{it},\mathbf{w}_{jt'})$, which are the error covariances at lags and leads. In other words, we assume that there is no dependence between node i and its network neighbors across different time points. Therefore, we only allow network dependence between contemporaneous errors, $\mathbb{E}(\varepsilon_{it}\varepsilon_{jt}|\mathbf{w}_{it},\mathbf{w}_{jt})$, and assume that such cross-sectional dependence becomes substantially weaker as one considers increasingly distant neighbors in the network.

5 Empirical Application

5.1 Data

We implement our risk-sharing test on data from global markets of rice, wheat, and maize. Data on consumption, production, exports, imports, and other aggregates at the country level are extracted from the Food and Agriculture Organization's (FAO) 'Food Balance Sheets' database (FAOSTAT, 2014). They provide country-level time series of production, domestic supply, food consumption, stock variations, and trade of major agricultural commodities from 1960 to 2013.

FAO also has data on trading pair-level trade flows for countries. This data is available from 1986 to 2013.[16] The dataset lists origin and destination country pairs with export and import amounts to the destination from the origin country. Our main usage of the trade-flow data is just the binary information on whether two countries have a trading relationship in a given year and in what direction. We use the trade quantities to trim the dataset, as discussed in Section 5.3. To measure trade flows, we use the export quantity for measuring exports from origin to destination (with some exceptions).[17]

For estimation, we combine the FAO trade-flow data with data from the Centre d'Etudes Prospectives et d'Informations (CEPII) gravity dataset. The CEPII contains a set of variables that are generally used by researchers to understand the determinants of trade flows between countries (Conte et al., 2020).[18] We use some of the variables in CEPII dataset to model the per unit bilateral transaction costs between countries. To capture trade frictions in terms of tariffs and other indirect trade barriers, we also use data on Nominal Rates of Assistance from the Distortions to Agricultural Incentives database (Anderson and Signe, 2013). Finally, trade networks may also be influenced by relative productivity and comparative advantage. Data on indicators of a particular crop's productivity based on agro-climatic conditions and natural endowments is available from the FAO's Global Agro-Ecological Zones database (Fischer et al., 2021). This dataset provides country-level simulated potential yields from agronomic models. We use these simulated potential yields to calculate relative productivity for country pairs.

[16] For some year and country pairs, we observe exports not matching imports on a non-trivial level. The FAO reports this can occur for a variety of reasons. One explanation is "exported quantities could be destroyed or lost en route due to accidents, weather conditions". Note that our definition of transaction-cost includes "lost at sea" as explaining part of the gap between imported consumption and exported endowment amount in the model. We do not believe that the transaction cost can simply be identified by comparing the gap between the FAO reported export/import amounts because our definition also includes other possible reasons.

[17] The dataset is not balanced and thus we must use the import amount when measuring exports from destination country to origin country in certain cases. To be specific: there are cases where country A exports to country B but country A is only listed as a destination country with an "import" amount that country B declares from A (which was the export from A to B). Thus for these cases, we generate the missing row for an A to B export amount using the import reported to be received by B from A.

[18] The CEPII provides trade pair level data on population weighted distance, time difference, whether the pairs were common colonies, common language, religious proximity, etc. It also provides data on macroeconomic variables like national income, population, total trade flows, and membership of regional trade agreements and the WTO.

5.2 Global Trade Networks

Figure 6 shows the network plots for the three commodities. The nodes denote countries and the edges denote the trade linkages. Since these linkages would vary over time depending on the state of the world and the production realizations, we show networks for three time periods. As can be seen from Figure 6, the network structure varies both across commodities and across time. Our estimation routine takes the network structure of trade into account while testing for risk sharing.

Another feature of networks is the number of connections called the *degree* of each node in the network. The degree distribution of a network gives us some information about the type of network. In our case, the networks of all three commodities have a left-skewed degree distribution (Appendix Figure A1). This implies that most nodes have a small number of connections, but a few nodes have a very high number of connections. This means that trade in these commodities is dominated by one or a few countries. Appendix Figure A2 plots the average number of trade partners per country for each year in our dataset. Across all three commodities, we observe a rising number of trade partners over time. This indicates that trade linkages have possibly grown on account of greater connectivity and lowering of freight and transaction costs.

5.3 Regression Results

We estimate equation (12) for 1986-2013. Our main sample includes countries with a population of at least 10 million and we ignore small trade quantities (we consider an expanded sample for robustness).[19] Also, we ignore the smaller amount in any reciprocal trade of the same commodity in the sense that we do not consider it as a link in that commodity's trade network; this is an equilibrium outcome of the model and fits the data in most cases. Our standard errors are derived using the network HAC estimator given in (13) which allows for general network dependence in the errors of the risk sharing equation as discussed in Section 4.3. The results are displayed in Table 2 for the commodities maize, wheat, and rice respectively. All specifications include year

[19] We use separate cutoffs per commodity as the mean trading levels vary by an order of magnitude: we cut off trade levels of 800 for maize, 20,000 for wheat, and 500 for rice (each is below 10% of the mean). We tried various cutoffs based on rules like % of total imports, and the relative results (within commodities across models) were not sensitive.

fixed effects to control for aggregate shocks. This is feasible as we have a global network each year. This implies that the aggregate shock will be common to all (trading) countries and be distinct each year. As discussed, this is generally not the case, but happens to be true in this particular sample.

The baseline specifications in Table 2 show the results from the risk-sharing test without accounting for transaction costs. We observe that own production is significant in the baseline regression in all three commodities, implying a rejection of the perfect risk-sharing hypothesis.[20] The wheat market is closest to achieving perfect risk sharing followed by rice and maize. In particular, for maize, the baseline production coefficient η is 0.303 with a standard error of 0.133. For wheat, η is 0.180 (0.042) and for rice, η is 0.241 (0.053).

In the second column for each commodity, we introduce transaction costs. The constant in the transaction cost term varies across each country-network pair as even if all countries have identical costs, the indirect Δ paths connecting them to the index country are different; some are longer/shorter, some have more exporters/importers, and these affect Δ even with constant δ. The pure effects of the network shape are captured here, via the relative number of exporters to importers along the path between a country and the index. The other controls add link specific heterogeneity in transaction costs. We consider log distance between the two countries and various economic/cultural similarity indices for the country pairs, which are standard variables in trade gravity models. Specifically, we include common language, time zone difference, GATT/WTO membership by destination and origin, relative (historical) productivity by country-pair, and rate of assistance from the World Bank "Distortion from Agricultural Incentives" Database. We also include time trends and country fixed effects to the transaction cost specifications.

For maize, the transaction cost specification yields an estimate on η of 0.159 (0.130). This is a 48% decline from the baseline and a 2% decrease in its standard error. For wheat, this specification yields a coefficient of 0.149 (0.042). The coefficient decreases by 17%, whereas the standard error does not change meaningfully. Finally, for rice, the estimate is 0.163 (0.046); this is a 32% decrease in the production coefficient and a 13% decrease in the standard error. Overall, the variables

[20]Relative results do not change if we include non-trading countries, only the coefficient on own production gets shifted down by running regression (16) instead.

included are not individually significant. However, the downstream effects on consumption via the network structure with link specific transaction cost variation do affect the production coefficient across various specifications, in particular for maize and rice.

In Appendix C, we discuss additional specifications, specifically one with heterogeneous risk preferences. Heterogeneity in risk preferences effectively include country fixed effects in the specification. Appendix Table A1 presents estimates, including country fixed effects in the baseline and network based risk sharing regressions. The results from these regressions are broadly comparable to those in Table 2.

5.4 Counterfactual Analysis

The regressions allow us to predict how a change in the transaction cost affects consumption, holding the existing network shape fixed and accounting for possible misspecification in the models. We first calculate counterfactual consumption by changing the estimated transaction cost parameters. We then compare these values with the fitted consumption from the original estimates. For maize, a 10% decrease in transaction costs leads to a 8.0% increase in consumption. For wheat and rice, the consumption increases are similar, at 6.9% and 5.2%, respectively. These predictions have a reduced form interpretation as they do not take into account how changing transaction costs will also change the equilibrium network shape.

Next, we fully solve the model for the equilibrium network shape and study how trade frictions, particularly freezing of trade during wartime, affects global risk sharing and utility.[21] It is not feasible to solve for all network shapes for 74 countries, so we consider global trade at the continental level. We group countries in North America, South America, Europe, Africa, and Asia/Oceania. We then collapse the data to calculate average transaction costs across them.

Our first full-model counterfactual in Table 3 shows the effects of changing the trade efficiency δ (thereby changing costs $1-\delta$) on the degree of risk sharing, measured by the correlation between

[21] Note that our risk-sharing test estimates indicate that the production parameter is non-trivial in most specifications, which does indicate that the intensive margin of trade may differ from the model. Thus our focus is on the extensive margin of trade, which may be less sensitive to this.

production and consumption. The first column shows the effects of fully efficient trade, which is equivalent to zero transaction costs (or $\delta = 1$). In this case, trade flows increase significantly such that the correlation between production and consumption decrease for all commodities. For rice, the negative correlation is a consequence of the highly concentrated exporters from whom the social planner redistributes a significant amount. This highlights how the global rice trade is substantially lower in volume than is optimal due to transaction costs. The second column shows the degree of risk sharing from the model based on the estimates from the actual trade data. Finally, the next two columns take the main estimate and decrease the estimated δ to simulate higher transaction costs. A 50% decrease leads to almost zero risk sharing in maize and rice (correlation close to 1). The decrease in risk sharing from increasing the costs is moderate, whereas the benchmark correlations are substantially lower than the main estimates. This indicates that transaction costs are still a major component in limiting global trade, and policy improvements in decreasing costs may have significant impacts.

The second counterfactual in Figure 7 shows the effects of isolating a single agent; consider the case of Russia after the 2020 invasion of Ukraine. Given that Ukraine and Russia are major producers and exporters of wheat, we simulate how the global trade in wheat is affected by a complete halting of trade between the rest of the world and Russia/Ukraine. This is close to reality, as the conflict led to Ukraine being unable to export wheat to the rest of the world and Russia being sanctioned. The trade embargo also has the effect of breaking down of other trade links. The first main effect is the isolation of Europe, as Russia/Ukraine were their main trade partner, and a second order effect is a shift in Africa's dependence on the US as a trading partner. Our estimates suggest that the overall global risk sharing does not significantly change as Asia's trading increases, but we find that the global utility declines substantially.

6 Conclusion

In this paper, we study the role of transaction costs in risk sharing. The standard model of risk sharing implicitly assumes costless trade across agents and is modeled at the level of consumption. We explicitly model trade flows between agents with iceberg-type transaction costs which vary with each trading pair and are proportional to the volume of trade flows. In this augmented model, link formation is costly, and hence trade networks play an important role in risk sharing. This seemingly simple modification adds a non-trivial estimation hurdle. Although the model gives us an empirically testable specification that mirrors the standard risk-sharing analysis, the existence of the transaction cost parameters requires a network-based approach. We estimate the model using consumption, production, and trade network data from the global food markets of rice, wheat, and maize. To conduct valid inference, we also adjust the standard errors for network dependence.

Simulations from the model show that declining transaction costs lead to expansion of the network and greater connectedness within the network. This in turn, leads to reduced consumption volatility and greater social welfare. We show that the omission of transaction costs and the implied network structure can bias the estimates of risk sharing. Risk sharing estimates and counterfactual analysis show that transaction costs impede risk sharing in all three food commodities, but more so for maize and rice trade than for wheat. This is consistent with the observation that more of the globally produced wheat is traded compared to rice and maize. Finally, our counterfactual analysis highlights the role that transaction costs and embargoes have in shaping global consumption and trading networks.

Our theoretical setup is general enough to fit into different contexts and can be used to study risk sharing in the presence of transaction costs among households in a village or across countries. The estimation method does not require observing the trade flows, but only the network shape alongside data on aggregate consumption and production. Given the availability of such information, the same model can be used to test for risk sharing under different contexts and settings. Extending the framework to incorporate storage, savings, or intertemporal trade are promising avenues for future research.

References

Ábrahám, Á. and Laczó, S. Efficient risk sharing with limited commitment and storage. *The Review of Economic Studies*, 85(3):1389–1424, 2018.

Ambrus, A., Mobius, M., and Szeidl, A. Consumption risk sharing in social networks. *American Economic Review*, 104(1):149–82, 2014.

Ambrus, A., Gao, W., and Milán, P. Informal risk sharing with local information. *The Review of Economic Studies*, 2022. ISSN 0034-6527. doi: 10.1093/restud/rdab091. URL Link.

Anderson, K. and Signe, N. Updated national and global estimates of distortions to agricultural incentives, 1955 to 2011. *The World Bank*, 2013. Retrieved from, worldbank.org.

Asdrubali, P., Sørensen, B. E., and Yosha, O. Channels of interstate risk sharing: United states 1963–1990. *The Quarterly Journal of Economics*, 111(4):1081–1110, 1996.

Asdrubali, P., Tedeschi, S., and Ventura, L. Heterogeneity in risk aversion and risk sharing regressions. *Journal of Applied Econometrics*, 34(5):827–835, 2019.

Asdrubali, P., Tedeschi, S., and Ventura, L. Household risk-sharing channels. *Quantitative Economics*, 11(3):1109–1142, 2020.

Attanasio, O. and Krutikova, S. Consumption insurance in networks with asymmetric information: Evidence from tanzania. *Journal of the European Economic Association*, 18(4):1589–1618, 2020.

Attanasio, O., Barr, A., Cardenas, J. C., Genicot, G., and Meghir, C. Risk pooling, risk preferences, and social networks. *American Economic Journal: Applied Economics*, 4(2):134–67, 2012.

Bardhan, P. and Udry, C. *Development Microeconomics*. OUP Oxford, 1999.

Bold, T. and Broer, T. Risk sharing in village economies revisited. *Journal of the European Economic Association*, 19(6):3207–3248, 2021.

Bradford, S. C., Negi, D. S., and Ramaswami, B. International risk sharing for food staples. *Journal of Development Economics*, 2022. Forthcoming: Retrieved from link.

Bramoullé, Y. and Kranton, R. Risk sharing across communities. *American Economic Review*, 97 (2):70–74, 2007.

Canova, F. and Ravn, M. O. International consumption risk sharing. *International Economic Review*, pages 573–601, 1996.

Clance, M., Ma, W., and Naraidoo, R. International consumption risk sharing and trade transaction costs. *University of Pretoria Department of Economics Working Paper Series*, 2019.

Cochrane, J. H. A simple test of consumption insurance. *Journal of Political Economy*, 99(5): 957–976, 1991.

Conte, M., Cotterlaz, P., and Mayer, T. The CEPII Gravity Database. *CEPII*, 2020.

Crucini, M. J. On international and national dimensions of risk sharing. *Review of Economics and Statistics*, 81(1):73–84, 1999.

Crucini, M. J. and Hess, G. D. International and intranational risk sharing. *CESifo Working Paper Series*, 1999.

De Weerdt, J. and Dercon, S. Risk-sharing networks and insurance against illness. *Journal of Development Economics*, 81(2):337–356, 2006.

Fafchamps, M. and Lund, S. Risk-sharing networks in rural philippines. *Journal of Development Economics*, 71(2):261–287, 2003.

FAOSTAT. Food Balance Sheets, 2014. Retrieved from www.fao.org/faostat.

Fischer, G., Nachtergaele, F., Van Velthuizen, H., Chiozza, F., Franceschini, G., Henry, M., Muchoney, D., and Tramberend, S. *Global Agro-Ecological Zones v4–Model documentation*. Food & Agriculture Org., 2021. Retrieved from, fao.org.

Fischer, M. M. and LeSage, J. P. Network dependence in multi-indexed data on international trade flows. *Journal of Spatial Econometrics*, 1(1):1–26, 2020.

Fitzgerald, D. Trade costs, asset market frictions, and risk sharing. *American Economic Review*, 102(6):2700–2733, 2012.

Genicot, G. and Ray, D. Group formation in risk-sharing arrangements. *The Review of Economic Studies*, 70(1):87–113, 2003.

Grund, T. U. Network analysis using Stata, 2015. Retrieved from nwcommands.org.

Jack, W. and Suri, T. Risk sharing and transactions costs: Evidence from kenya's mobile money revolution. *American Economic Review*, 104(1):183–223, 2014.

Kojevnikov, D., Marmer, V., and Song, K. Limit theorems for network dependent random variables. *Journal of Econometrics*, 222(2):882–908, 2021.

Laczó, S. Risk sharing with limited commitment and preference heterogeneity: Structural estimation and testing. *Journal of the European Economic Association*, 13(2):265–292, 2015.

Lee, Y., Shen, C., Priebe, C. E., and Vogelstein, J. T. Network dependence testing via diffusion maps and distance-based correlations. *Biometrika*, 106(4):857–873, 2019.

Lewis, K. K. What can explain the apparent lack of international consumption risk sharing? *Journal of Political Economy*, 104(2):267–297, 1996.

Ligon, E. Risk sharing and information in village economies. *The Review of Economic Studies*, 65 (4):847–864, 1998.

Ligon, E., Thomas, J. P., and Worrall, T. Informal insurance arrangements with limited commitment: Theory and evidence from village economies. *The Review of Economic Studies*, 69(1): 209–244, 2002.

Liu, X. and Prucha, I. R. A robust test for network generated dependence. *Journal of econometrics*, 207(1):92–113, 2018.

Mace, B. J. Full insurance in the presence of aggregate uncertainty. *Journal of Political Economy*, 99(5):928–956, 1991.

Mazzocco, M. and Saini, S. Testing efficient risk sharing with heterogeneous risk preferences. *American Economic Review*, 102(1):428–68, 2012.

Obstfeld, M. Are industrial-country consumption risks globally diversified? *Capital mobility: The impact on consumption, investment and growth*, pages 13–44, 1994.

Obstfeld, M. and Rogoff, K. The six major puzzles in international macroeconomics: Is there a common cause? *NBER Macroeconomics Annual*, 15:390–403, 2001.

Schulhofer-Wohl, S. Heterogeneity and tests of risk sharing. *Journal of Political Economy*, 119 (5):925–958, 2011.

Su, L., Lu, W., Song, R., and Huang, D. Testing and estimation of social network dependence with time to event data. *Journal of the American Statistical Association*, 2019.

Townsend, R. M. Risk and insurance in village india. *Econometrica*, pages 539–591, 1994.

Townsend, R. M. Consumption insurance: An evaluation of risk-bearing systems in low-income economies. *Journal of Economic perspectives*, 9(3):83–102, 1995.

Figures and Tables

Figure 1: **Effect of Increasing δ on Trade Network**

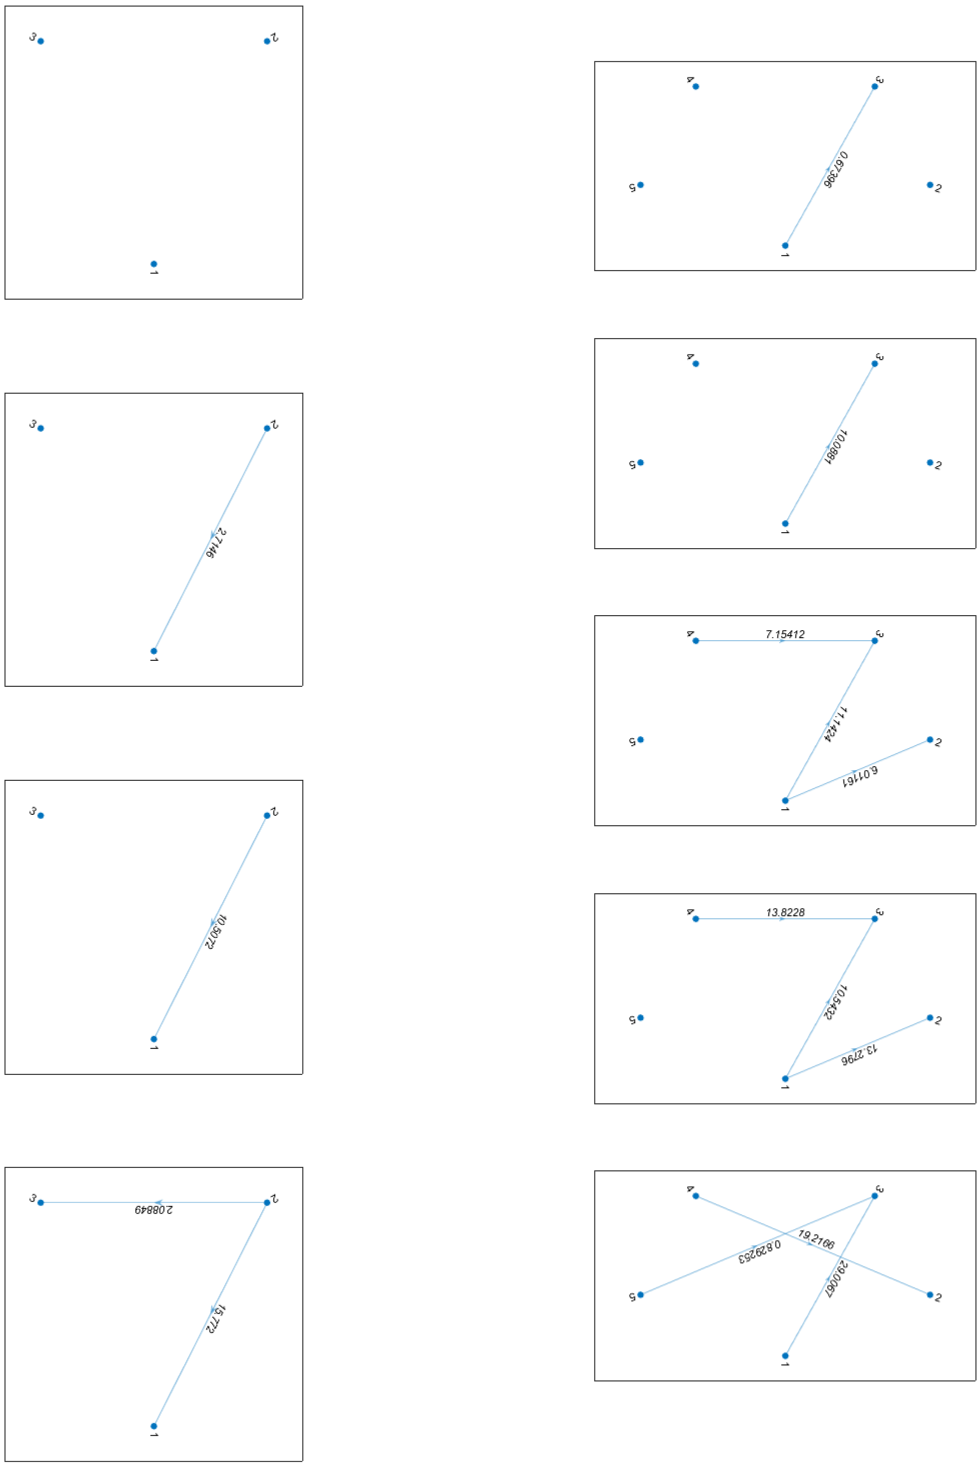

Note: The figure plots the optimal network shapes as δ increase or transaction costs decrease. The left panel is a 3 country example and the right panel is a 5 country example.

Figure 2: **Effect of Increasing δ on Consumption**

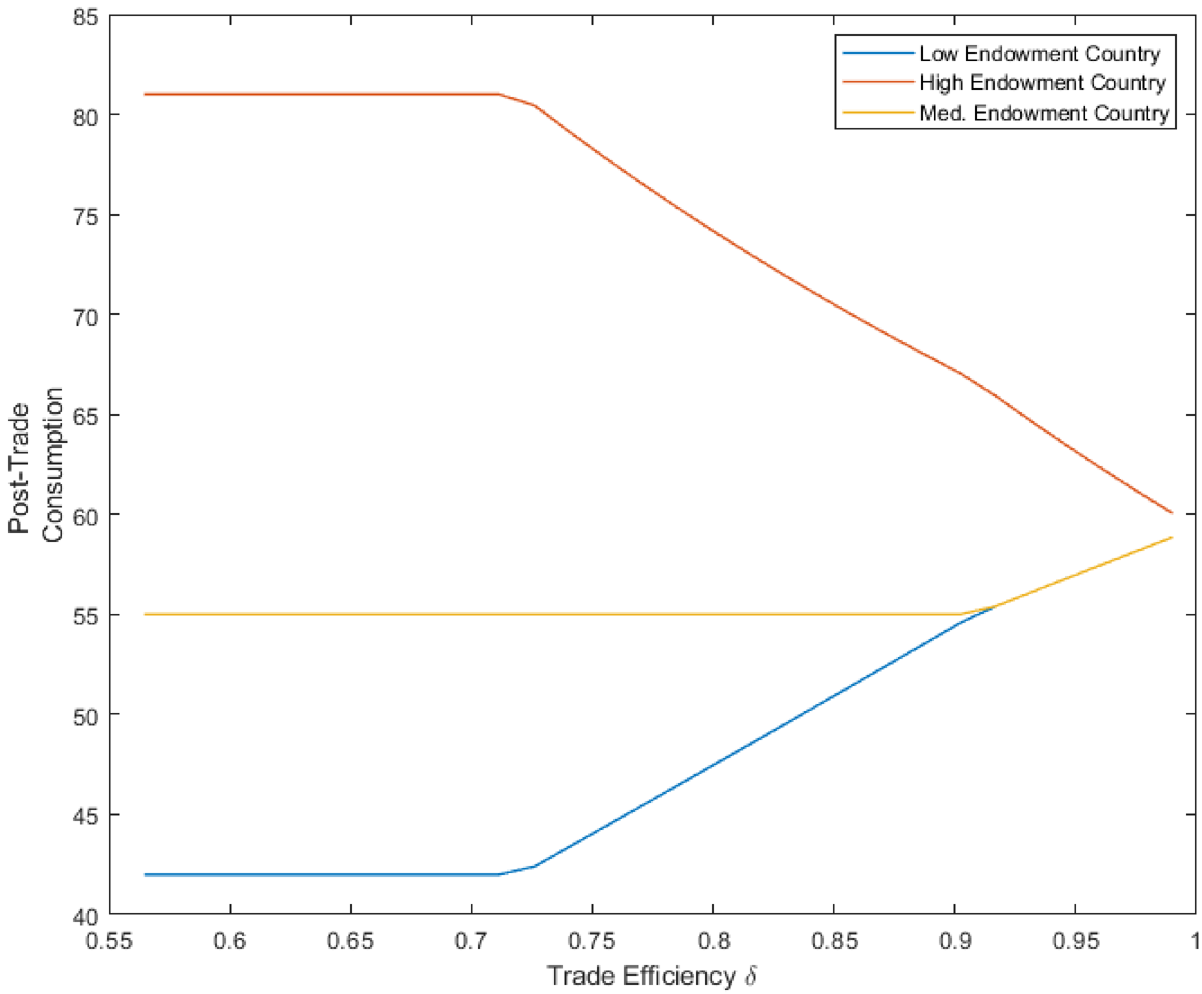

Note: The figure plots the optimal consumption (from social planner's problem) in a 3 country example as δ increase of transaction costs decrease.

Figure 3: **Adding Trade Partners**

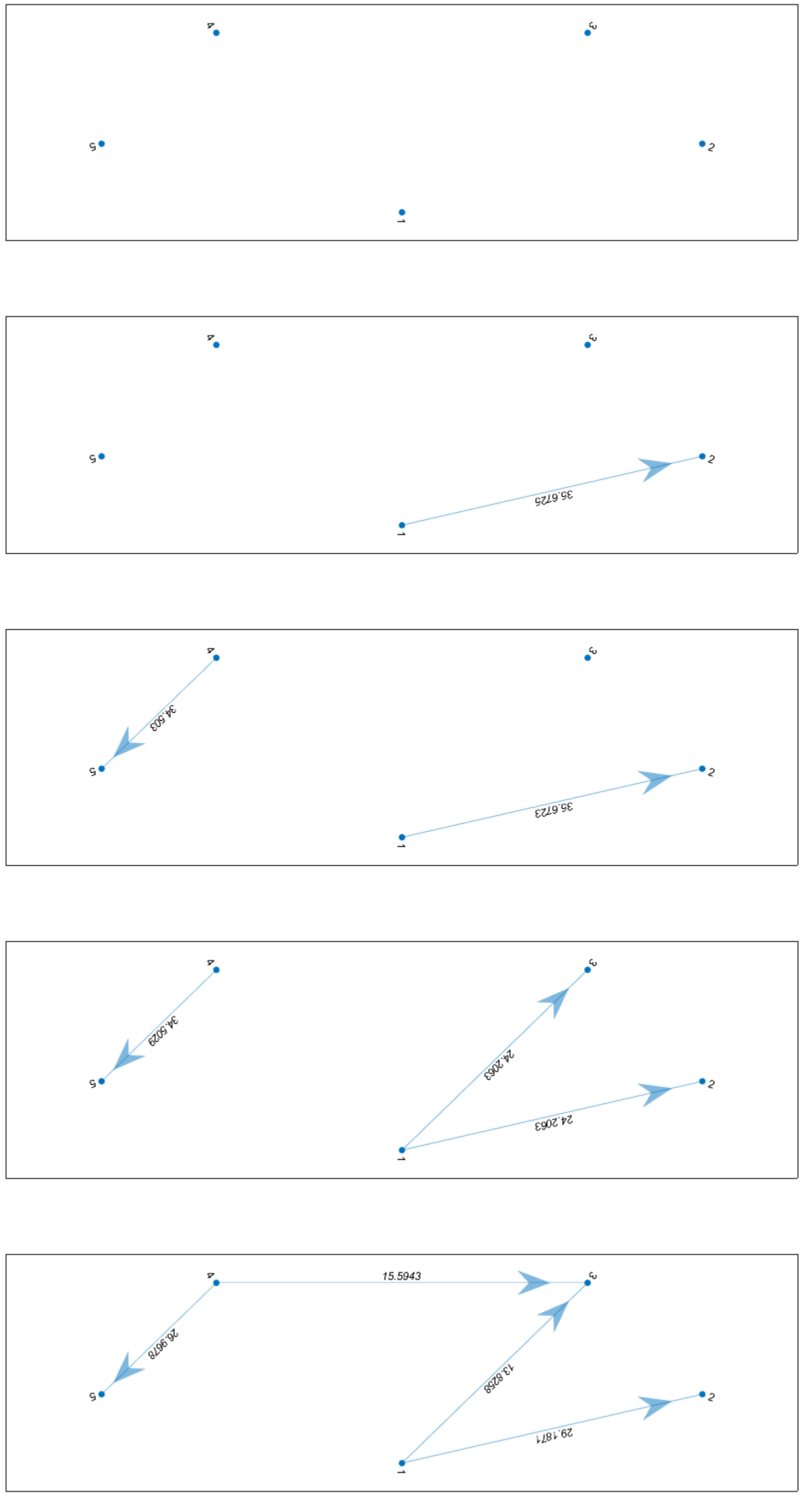

Note: The figure plots the network shapes of a 5 country example under different trading scenarios.

Figure 4: **Efficiency Effects of Adding Partners Consumption Variance and Global Utility**

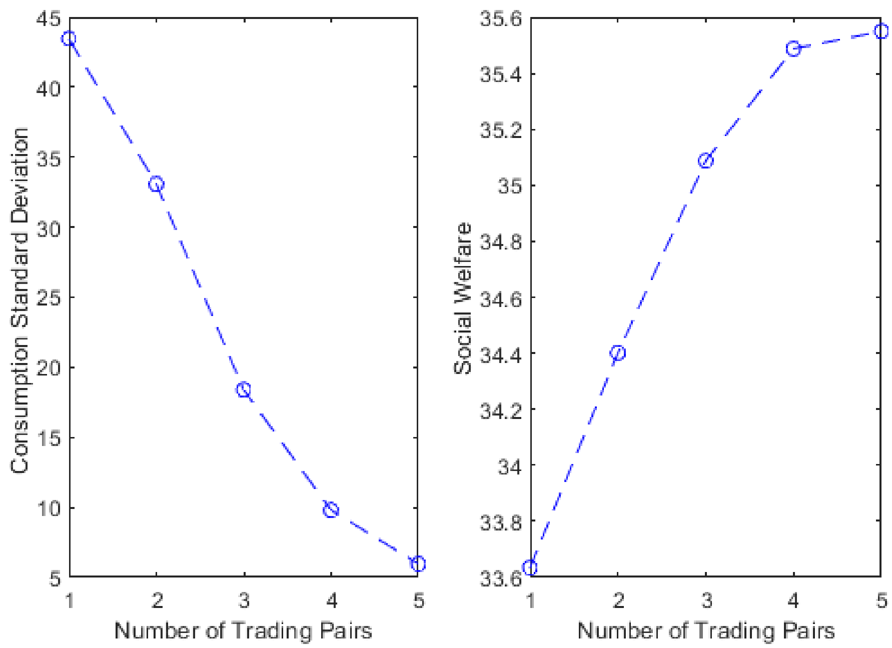

Note: The figure plots the consumption standard deviation and social welfare of a 5 country example when increasing the number of trading pairs.

Figure 5: **Transaction/Trade Costs, Trade Activity and Risk Sharing**

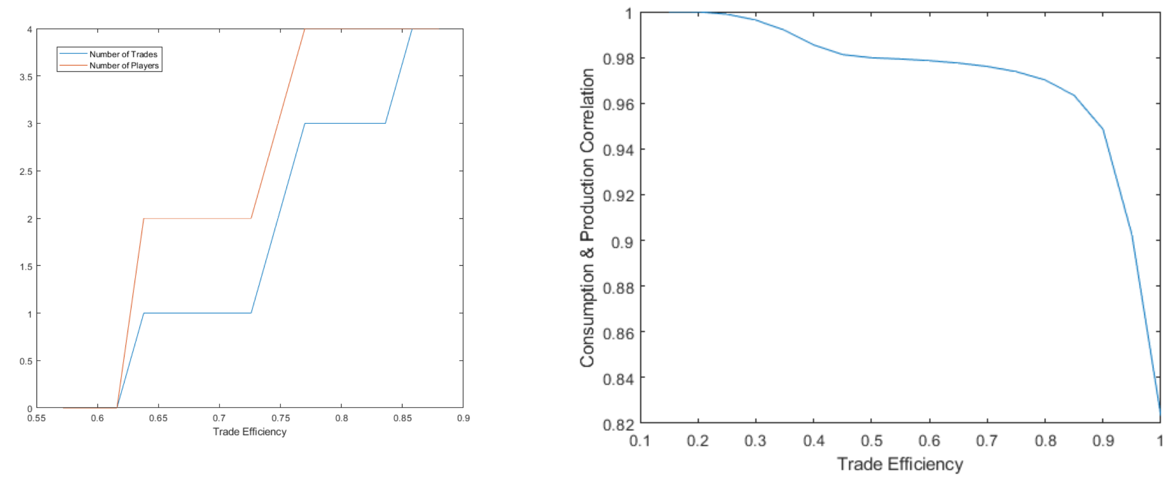

(a) Trade Efficiency and Trade Networks

(b) Trade Efficiency and Risk Sharing

Note: Figure (a) plots the number of trading links and number of members in the network as trade costs decrease. Figure (b) plots the correlation between own consumption and production as trade costs decrease.

Figure 6: **Trade Networks for the Three Commodities**

(a) Rice

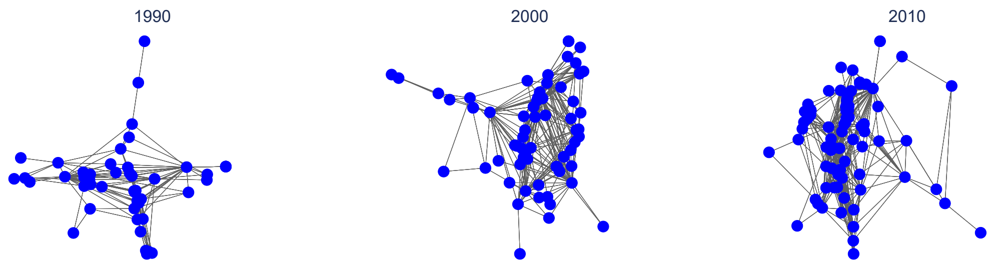

(b) Wheat

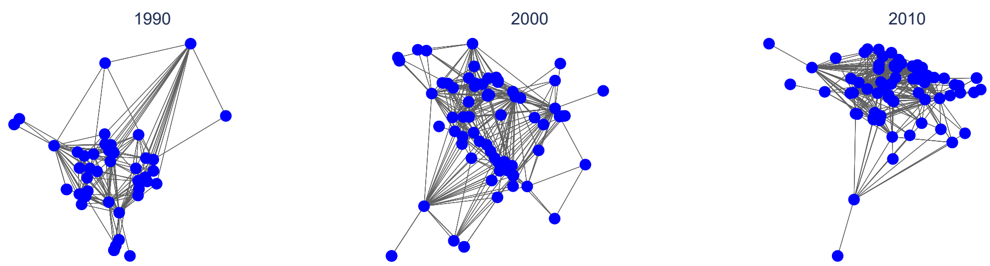

(c) Maize

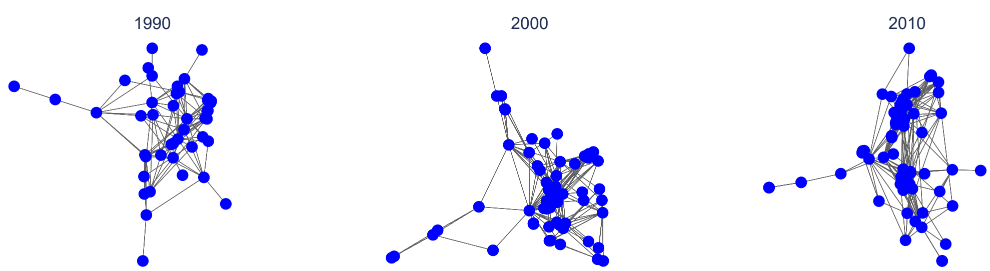

Note: The figure presents the network plots for rice, wheat, and maize for three slices of time. The blue dots are nodes or trading countries and the lines or edges denote trade links between countries.

Figure 7: **Wheat Russia/Ukraine Trade Collapse Counterfactuals**

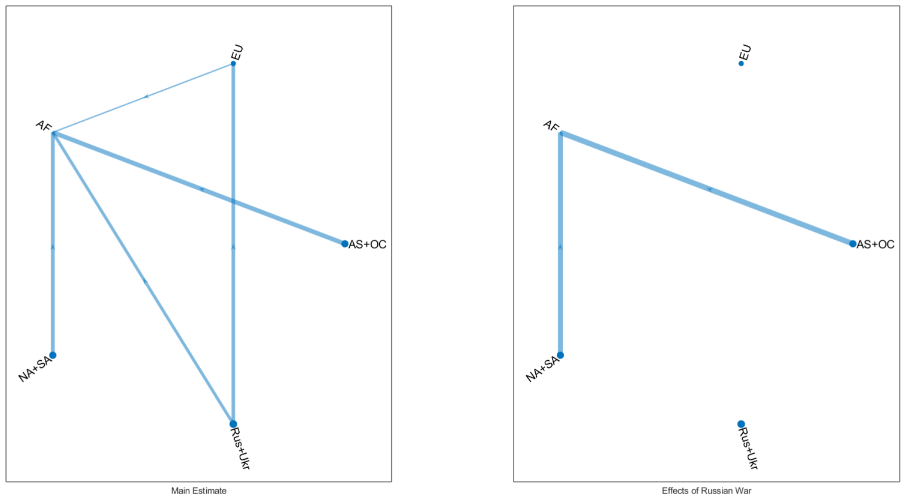

The figure plots equilibrium trade under two scenarios. The node side is the endowment and the line size is the trade flow. The country groups are: NA+SA=North America and South America, EU=European Union, AF=Africa, AS+OC=Asia and Oceania, and RU+UK=Russia and Ukraine.

Table 1: **Simulated Risk-Sharing Regression and Bias**

Dependent Variable: log consumption			
	(1)	(2)	(3)
Log production	0.3913^{***}	0.3377^{**}	0.0647
	(0.0157)	(0.0457)	(0.0325)
Control for $\mathcal{A}$	No	Yes	Yes
Control for Δ	No	No	Yes
Observations	128	128	128
R^2	0.546	0.673	0.925

Note: Table present the estimates of risk sharing regressions from simulated data. The term $\mathcal{A}$ refers to the aggregate risk function. The term Δ refers to the indirect trade weight, which captures transaction costs. Network dependent standard errors in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 2: **Structural Risk-Sharing Test Estimates**

Commodity:	Maize		Wheat		Rice	
Dependent Variable: log consumption	Baseline	Tran. Costs	Baseline	Tran. Costs	Baseline	Tran. Costs
Log production	0.303*	0.159	0.180***	0.149***	0.241***	0.163***
	(0.133)	(0.130)	(0.042)	(0.042)	(0.053)	(0.046)
Preference parameter	1.288***	1.092***	1.252***	1.297***	1.121***	1.281***
	(0.33)	(0.215)	(0.134)	(0.143)	(0.138)	(0.197)
Constant	10.755***	13.249***	12.677***	13.423***	13.771***	15.501***
	(1.931)	(1.925)	(0.502)	(0.652)	(0.924)	(0.812)
Trade-cost constant		4.003		20.225		-1.370
		(34.697)		(33.477)		(26.842)
Common language(ij)		0.809		-0.020		-0.098
		(3.185)		(4.836)		(3.202)
Log weighted distance(ij)		-0.866		-2.309		-1.035
		(2.865)		(4.245)		(2.983)
Country i GATT member(it)		0.274		0.172		0.260
		(6.844)		(6.500)		(5.612)
Country j GATT member(jt)		0.671		0.356		-0.228
		(24.563)		(6.479)		(14.864)
Time difference between i and j (ij)		0.036		0.259		0.071
		(0.555)		(0.694)		(0.651)
Log relative productivity(ij)		-0.266		0.033		-0.002
		(2.179)		(0.967)		(0.706)
Nominal rate of assistance(ijt)		-0.003		-0.009		0.003
		(0.265)		(0.686)		(0.216)
Time trend		0.050		0.038		-0.001
		(0.269)		(0.214)		(0.247)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Transaction Cost Country FE	No	Yes	No	Yes	No	Yes
Observations	1752	1752	1722	1722	1860	1860

Notes: The preference parameter is a coefficient both on the Pareto weights log ratio and the transaction cost. Common language is a dummy variable coded as 1 if a language is spoken by at least 9% of the population in both countries. Population weighted distance between the two countries in kilometers. Time difference in number of hours difference between country i and j. GATT membership is a dummy variable coded as 1 if the country is GATT/WTO member. Relative (historical) productivity is calculate for country-pair and is extracted from the GAEZ database. Nominal rates of assistance (scaled by 10) are from the World Bank, Distortion from Agricultural Incentives Database. Network dependent standard errors in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 3: **Counterfactual Analysis**

	(1)	(2)	(3)	(4)
Corr(Y,C)	Benchmark	Main Estimate	90%	50%
Maize	0.26457	0.78791	0.86175	0.98618
Wheat	0.34561	0.80947	0.93309	0.94149
Rice	-0.41892	0.96605	0.98257	0.9962

Notes: Table presents the correlation between production and consumption simulated for different transaction cost scenarios. Full benchmark is the case of zero transaction costs. Main estimate in column (2) is based on the transaction cost parameters estimated from the data. In columns (3) and (4), we scale down the estimated bilateral δ's (from the main estimate) by the 90% and 50%, respectively.

Appendix

A Model Details

A.1 Symmetric Trade Result

Symmetric trading is not optimal. This motivates excluding the smaller importer in such cases.

Lemma 1. *Given a linear transaction cost of the form in equation (4) and with nonzero fractional bilateral trade parameter $\delta_{ij} \in (0,1]\ \forall i \forall j \neq i$ (assuming both not equal to 1), any nonzero bilateral trade between i and j is one directional, meaning $y_{ij}y_{ji} = 0, y_{ij} \geq 0, y_{ji} \geq 0$.*

Proof. Suppose both i and j trade with each other, meaning $y_{ij} > 0$ and $y_{ji} > 0$. Then, based on the FOC system in equation (6), the two conditions would need to hold simultaneously:

$$\frac{\partial U}{\partial c_i}\delta_{ji} = \frac{\partial U}{\partial c_j} \cap \frac{\partial U}{\partial c_j}\delta_{ij} = \frac{\partial U}{\partial c_i}$$

These can only hold if $\delta_{ji} = 1/\delta_{ij}$. However that can only happen if both are equal to 1 or if one is greater than 1 (either case violating the assumptions). Thus one of these first order conditions does not hold with equality, meaning at least one export is zero. □

A.2 Optimal Consumption Result

Proposition 1. *Optimal c_i^* is a linear combination of endowments.*

Proof. Given a known set of exporters and for a given state, the system of equations that define the optimal export choices can be written as the following:

$$\Delta_{ij}\left(y_j - \sum_{k\neq j}^{N} y_{jk}^E + \sum_{k\neq j}^{N} y_{kj}^E\delta_{kj}\right) = \left(y_i - \sum_{k\neq i}^{N} y_{ik}^E + \sum_{k\neq i}^{N} y_{ki}^E\delta_{ki}\right) \quad \forall \text{ exporter } i$$

Implicit in this equation is that $y_{ij} > 0$ and thus $y_{ji} = 0$ (by Lemma 1 in Appendix A). This can be rearranged into an $\mathbf{Ax} = \mathbf{b}$ matrix format, with $\mathbf{x} = [y_{ij}^E, ...]$. $\mathbf{b} = [y_i - \Delta_{ij}y_j, ...]$. The formula for $\mathbf{A}$ is convoluted and involves the products of transaction costs, relative Pareto weights, and the preference parameter. Then the solution by inversion, due to being square (and assuming full rank), is $\mathbf{x} = \mathbf{A}^{-1}\mathbf{b}$. This is simply a linear combination of the vector $\mathbf{b}$, which in this case is just a linear combination of endowments $y_i \forall i$. Thus the export choices are linear in endowments. Since consumption is also just a linear combination of exports and endowments, optimal consumption will be a linear combination of endowments. □

Specifically, the endowments are y_i and y_j from the same trade network. Optimal consumption has the following pattern, where $\Omega \equiv [\Delta_{ij}, \delta_{ij}\ \forall i \forall j \neq i]$, and f, g^{ik}, h^{ki} are all functions $c_i^* = (y_i \cdot f(\Omega) - \sum_{k\neq i}^{N} g^{ik}(y_k \forall k, \Omega) + \sum_{k\neq i}^{N} h^{ki}(y_k \forall k, \Omega)\delta_{ki})/f(\Omega)$. Both agent pair specific functions (g^{ik} and h^{ki}) are simply weighted sums of all endowments. The endowments of any agent connected via the trade network has non-zero weight. Any agent in a separate network has zero weight.

A.3 Indirect Trade Weight Result

Proposition 2. *The (geodesic path) indirect trade weight from index i to ending partner k is:*

$$\Delta_{i\to k} = \left(\frac{\alpha_k}{\alpha_i}\prod_{(r,s)\in\mathcal{P}_{ik}}\frac{\delta_{rs}^{\mathcal{S}(r,s)}}{\delta_{sr}^{\mathcal{S}'(r,s)}}\right)^{-1/\gamma} \tag{14}$$

Proof. Consider the undirected geodesic path from i to k, which exists if $i,k\in\mathcal{T}$, but may not be unique. There exists a path $i, k_1, ..., k_{n-1}, k_n = k$ with n steps. The first FOC comparing i to k_1 is either $c_i = \Delta_{i,k_1}c_{k_1}$ or $c_{k_1} = \Delta_{k_1,i}c_i$ depending on which is the importer or exporter; either $\mathcal{S}'(i,k_1) = 1$ or $\mathcal{S}(i,k_1) = 1$, which due to Lemma 1 are mutually exclusive. The second is either $c_{k_1} = \Delta_{k_1,k_2}c_{k_2}$ or $c_{k_2} = \Delta_{k_2,k_1}c_{k_1}$. The combined chain is then:

$$c_i = \Delta_{i,k_1}^{\mathcal{S}(i,k_1)}\cdot\Delta_{k_1,i}^{-\mathcal{S}'(i,k_1)}\cdot\Delta_{k_1,k_2}^{\mathcal{S}(k_1,k_2)}\cdot\Delta_{k_2,k_1}^{-\mathcal{S}'(k_1,k_2)}c_{k_2}$$

If $n=2$ then the expression above is the final answer. Now suppose it holds for $n=m$, meaning the path from i to k with m steps yields: $c_i = \prod_{(r,s)\in\mathcal{P}_{ik_m}}\Delta_{rs}^{\mathcal{S}(r,s)}\Delta_{sr}^{-\mathcal{S}'(r,s)}c_{k_m}$.

Now suppose we actually need one more link to reach k from i, meaning $n = m+1$. If m exports to $m+1=k$, then $c_m = \Delta_{m,k}c_k$ and if k exports to m then $c_k = \Delta_{k,m}c_m$. Then we can compare c_i to c_k through c_{k_m}:

$$c_i = \Delta_{i\to k_m}c_{k_m} = \Delta_{i\to k_m}\cdot(\Delta_{m,k}^{\mathcal{S}(m,k)}\Delta_{k,m}^{-\mathcal{S}'(m,k)}c_k)$$

Thus we've shown $\Delta_{i\to k} = \prod_{(r,s)\in\mathcal{P}_{ik}}\Delta_{rs}^{\mathcal{S}(r,s)}\Delta_{sr}^{-\mathcal{S}'(r,s)}$. Substituting out $\Delta_{ik}\equiv(\delta_{ik}\alpha_k/\alpha_i)^{-1/\gamma}$ immediately yields the result. $\square$

Note that the geodesic path is not in general the "lowest cost" path; geodesic paths only consider the number of links and there may be a "longer" path that has a lower indirect cost.[22] Also, in the case of heterogeneous γ, the ratios of γ_i along the path change the formula in scaling each intermediary direct Δ.

B Estimation Details

B.1 Ignoring Transaction Costs and Bias

To see how equation (12) compares with the benchmark of no transactions costs, we derive the test assuming $\delta_{ijt}=1$. The canonical risk-sharing regression is equivalent to the model with $\delta_{ijt}=1$ for all trading pairs. This results in the following regression:

$$\ln(c_{jt}) = (1/\gamma)\ln(\alpha_j) + T_t + \eta\cdot\ln(y_{jt}) + \tilde{\varepsilon}_{jt} \tag{15}$$

Equation (15) is the standard test of risk sharing. Since there are no transaction costs, the trading network is global and aggregate risk, T_t, is only indexed by time. Also, each Δ_{jt} just

[22]Finding the lowest-cost path from all paths is computationally restrictive; restricting the set of paths based on some fixed length above the geodesic length and comparing the trade weights from each path is feasible.

collapses to the Pareto weights. The standard tests, given in (15), are generally estimated using a two-way fixed effects approach where time fixed effects control for aggregate risks and unit fixed effects control for the Pareto weights. The two-way fixed effects approach is not generally appropriate for our case as we allow networks to change every year in an unrestricted manner. This implies that both Δ_{jt} and $\mathcal{A}^*_{\mathcal{T}_t t}$, which depend on the network shape and the position of the agent in the network, will vary across individuals and over time. Aggregate risk and indirect trade weight can only be subsumed into a network-year and a unit fixed effect when the networks are observed and are stable over time.[23]

If in reality we have $\delta_{ijt} < 1$, then the estimate of η from regression (15) is biased. To see how transaction costs can influence the estimate of risk sharing, we decompose the trade weight term using the formula from equation (10) as $-\ln(\Delta_{jt}) = (1/\gamma)\ln(\alpha_j/\alpha_{\bar{i}_t}) + (1/\gamma)\ln(F(\boldsymbol{\delta}_{\mathcal{N}_t}|\mathcal{S}_t))$. The F term is a network specific function of transaction costs for exporters and importers along country j's geodesic path to the index country in the network: $F(\boldsymbol{\delta}_{\mathcal{N}_t}|\mathcal{S}_t) = \prod_{(r,s)\in\mathcal{P}_{\bar{i}_t,j}} (\delta_{rst})^{\mathcal{S}_t(r,s)}/(\delta_{srt})^{\mathcal{S}'_t(r,s)}$. This captures the country specific transaction cost network effects which limit risk sharing. The costs of each country along the path between a country and the central trading country (the index) is an intensive measure of how costly trade is for a country, which affects their consumption level.

The extent of the bias is based on the covariance $\mathrm{Cov}(\log(y_{it}), \ln(F(\boldsymbol{\delta}_{\mathcal{N}_t}|\mathcal{S}_t)))$, which is likely non-zero for two reasons. First, the (optimal) network shape is endogenously determined by the social planner and hence is a function of the endowment distribution. The social planner wants to re-distribute endowment but takes into account the efficiency per transaction; very inefficient linkages are less likely to be optimal. In this case, inefficient countries are more likely to be left out of the network; if a low endowment country is surrounded by inefficient countries, then their trade efficiency is likely low and they may be left out.

Second, the exogenous characteristics that affect transaction costs may correlate with endowments. Low endowment countries may have high barriers to trade, which exacerbates their problem as they can benefit the most from sharing risk, but are least able. Countries with efficient trading cost network effects (those whose linkages are mostly exporters with low costs) may have higher endowments; in this case the covariance is positive and there is upward bias in η.

B.2 Trading Participation Selection

The risk-sharing equation is derived from the FOC of the model; for agents who are not part of any network in a given year, their consumption is simply equal to their production: $\ln(c_{jt}) = \ln(y_{jt}) + \varepsilon_{jt}$, and so there is no "risk-sharing" parameter to test since they are unable to share risk in that period. Note however that the fact that the agent is not trading is endogenous to the model; the realization of shocks in a year combined with transaction costs could lead to an agent not trading. Thus it does raise the question of selection bias present in the sample if we only estimate the risk-sharing regression for trading agents.

If one simply wants to capture how being in a network reduces dependence on own production, then conditioning on the network participants is sufficient. However if one is interested in understanding the extent of global risk sharing induced by variation in endowment shocks, then

[23] If every agent is linked directly or indirectly to every other agent (meaning no isolated networks) then global trading is one network and the aggregate risk function varies only at the yearly level: $\mathcal{A}^*_{\mathcal{T}_t t} = \mathcal{A}^*_t$. In this case, it can be subsumed by a year fixed effect.

one should include excluded countries as their autarkic status in a given year reveals the extent of risk sharing.[24] If the selection is based on observables, then an augmented risk-sharing equation is sufficient to capture membership as follows (expressed in terms of a single network):

$$\ln(c_{it}) = \ln(c_{it}^*) \cdot \mathbb{1}[i_t \in \mathcal{T}_t] + \ln(y_{it}) \cdot \mathbb{1}[i_t \notin \mathcal{T}_t] + \eta \ln(y_{it}) + \varepsilon_{it} \tag{16}$$

B.3 Asymptotic Derivations

The risk sharing test controlling for aggregate risk function and transaction costs is given as:

$$\ln(c_{jt}) = \ln(\mathcal{A}^*_{\mathcal{T}_t t}) + \eta \cdot \ln(y_{jt}) - \ln(\Delta^*_{jt}) + \varepsilon_{jt}, \tag{17}$$

for each $j \in \mathcal{N}_N = \{1, \ldots, N\}$ and $t = 1, 2, \ldots, T$. Let $\mathbf{x}_{it} = (x_{it1}, x_{it2})$, $\mathbf{z}_{it} = \left(\{z_{it}^j\}_1^{|I_{it}|}, \{z_{it}^k\}_1^{|E_{it}|}\right)$ where $|I_{it}|$ and $|E_{it}|$ denote the cardinality of the importer and exporter sets, respectively of country i at time t. Let $\mathbf{f}_t = (f1_t, \ldots, fT_t)$ denote the vector of binary time indicators. Let $\mathbf{w}_{it} = (\mathbf{d}_t, \mathbf{x}_{it}, \mathbf{z}_{it})$ and $\boldsymbol{\theta} = (\boldsymbol{\alpha}, \boldsymbol{\beta})'$. Here $\boldsymbol{\beta} = (\beta_1, \beta_2, \beta_3, \beta_4)'$ and $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_T)'$. Then, one may rewrite the risk sharing regression as:

$$\begin{aligned} y_{it} &= \mathbf{f}_t\boldsymbol{\alpha} + \beta_1 x_{it1} + \beta_2 x_{it2} + \beta_2 \left[\sum_{j \in I_{it}} g(z_{it1}^j, z_{it2}^j; \beta_3, \beta_4) - \sum_{l \in E_{it}} g(z_{it1}^l, z_{it2}^l; \beta_3, \beta_4)\right] + \varepsilon_{it} \\ &\equiv h(\mathbf{w}_{it}, \boldsymbol{\theta}) + \varepsilon_{it} \end{aligned} \tag{18}$$

where $g(\cdot;\cdot) = \log\left(\frac{\exp(\cdot)}{1+\exp(\cdot)}\right)$. The summations over j and l are adding over countries that indirectly connect i with the index country $\bar{i}$ in the importer and exporter sets, respectively. Hence, $\{(y_{it}, \mathbf{w}_{it});\ i \in \mathcal{N}_N; t = 1, 2, \ldots, T\}$ represent the observed sample from the population. Then, stacking equation (18) across time periods, we can write the regression more compactly as[25]

$$\mathbf{y}_i = \mathbf{h}_i(\boldsymbol{\theta}) + \boldsymbol{\varepsilon}_i,\ i \in \mathcal{N}_N \tag{19}$$

where $\mathbf{y}_i = (y_{i1}, \ldots, y_{iT})'$, $\mathbf{h}_i(\boldsymbol{\theta}) = (h_{i1}(\boldsymbol{\theta}), \ldots, h_{iT}(\boldsymbol{\theta}))'$, and $\boldsymbol{\varepsilon}_i = (\varepsilon_{i1}, \ldots, \varepsilon_{iT})'$.

Assumption 1. $\mathbb{E}(\varepsilon_{it}|\mathbf{w}_i) = \mathbb{E}(\varepsilon_{it}|\mathbf{w}_{it}) = 0$

Assumption 1 implies that strict exogeneity holds with respect to risk sharing regression equation. Given the estimating equation in 19, we minimize the following nonlinear least squares objective function,

$$\min_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \frac{1}{NT} \sum_{i \in \mathcal{N}_N} \hat{\boldsymbol{\varepsilon}}_i^2 \equiv \min_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \frac{1}{NT} \sum_{i \in \mathcal{N}_N} (\mathbf{y}_i - \mathbf{h}_i(\hat{\boldsymbol{\theta}}))^2 \tag{20}$$

[24]This would require a different approach: regress the probability of a network forming of a certain type as a function of endowment shocks. This is a selection equation that the post-selection risk-sharing regression ignores.

[25]Note that we do not explicitly introduce a binary indicator s_{it} for missing data given that our application involves an unbalanced panel. We simply assume that the observed data is conditionally independent of the errors, ε_{it}.

with the first order conditions for $\hat{\boldsymbol{\theta}}$ given as

$$\frac{1}{NT}\sum_{i\in\mathcal{N}_N}\nabla_{\boldsymbol{\theta}}\mathbf{h}_i(\hat{\boldsymbol{\theta}})'\hat{\boldsymbol{\varepsilon}}_i = \frac{1}{NT}\sum_{i\in\mathcal{N}_N}\mathbf{h}_i^{(1)}(\hat{\boldsymbol{\theta}})\hat{\boldsymbol{\varepsilon}}_i = \mathbf{0} \tag{21}$$

where

$$\mathbf{h}_i^{(1)}(\hat{\boldsymbol{\theta}})' = \nabla_{\boldsymbol{\theta}}\mathbf{h}_i(\hat{\boldsymbol{\theta}}) = \begin{pmatrix} \frac{\partial h_{i1}(\hat{\boldsymbol{\theta}})}{\partial\theta_1} & \cdots & \frac{\partial h_{i1}(\hat{\boldsymbol{\theta}})}{\partial\theta_M} \\ \frac{\partial h_{i2}(\hat{\boldsymbol{\theta}})}{\partial\theta_1} & \cdots & \frac{\partial h_{i2}(\hat{\boldsymbol{\theta}})}{\partial\theta_M} \\ \vdots & \ddots & \vdots \\ \frac{\partial h_{iT}(\hat{\boldsymbol{\theta}})}{\partial\theta_1} & \cdots & \frac{\partial h_{iT}(\hat{\boldsymbol{\theta}})}{\partial\theta_M} \end{pmatrix}_{(T\times M)}$$

Then, by first-order taylor expansion of equation (21) around the true $\boldsymbol{\theta}$, we have the following influence function representation for $\hat{\boldsymbol{\theta}}$

$$\sqrt{N}(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}) = -\mathbf{Q}_T^{-1}\left(\frac{1}{\sqrt{N}}\sum_{i\in\mathcal{N}_N}\mathbf{v}_i\right) + o_p(1) \tag{22}$$

where, $\mathbf{v}_i = \mathbf{h}_i^{(1)}(\boldsymbol{\theta})\boldsymbol{\varepsilon}_i = \frac{1}{T}\sum_{t=1}^{T}\nabla_{\boldsymbol{\theta}}h_{it}(\boldsymbol{\theta})\varepsilon_{it}$ and $\mathbf{Q}_T = \operatorname*{plim}_{N\to\infty}\frac{1}{NT}\sum_{i\in\mathcal{N}_N}\nabla_{\boldsymbol{\theta}}\{\mathbf{h}_i^{(1)}(\tilde{\boldsymbol{\theta}})\tilde{\boldsymbol{\varepsilon}}_i\}$ $= -\frac{1}{T}\mathbb{E}[\mathbf{h}_i^{(1)}(\boldsymbol{\theta})\mathbf{h}_i^{(1)}(\boldsymbol{\theta})']$.

Now, in order to obtain the asymptotic distribution of $\hat{\boldsymbol{\theta}}$, we need to apply the LLN and CLT for network dependent data to the sequence, $N^{-1/2}\sum_{i\in\mathcal{N}_N}\mathbf{v}_i$, where $\mathbb{E}(\mathbf{v}_i|\mathbf{w}_i) = \mathbf{0}$ due to assumption 1a).[26] Provided that certain regularity conditions hold (mentioned below) which are informed by KMS, we will obtain the following asymptotic normality result.

$$\sqrt{N}(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}) \sim \mathrm{N}\left(\mathbf{0}, \mathbf{Q}_T^{-1}\boldsymbol{\Omega}_T\mathbf{Q}_T^{-1}\right) \tag{23}$$

with the network-HAC variance is given by

$$\boldsymbol{\Omega}_T = \text{Avar}(N^{-1/2}\sum_{i\in\mathcal{N}_N}\mathbf{v}_i) = \lim_{N\to\infty}\frac{1}{NT^2}\sum_{d\geq 0}\sum_{i\in\mathcal{N}_N}\sum_{j\in\mathcal{N}_N^{\partial}(i,d)}\mathbb{E}(\mathbf{v}_i\mathbf{v}_j'|\mathbf{w}_i,\mathbf{w}_j) \tag{24}$$

where $\mathcal{N}_N^{\partial}(i,d) = \{j\in\mathcal{N}_N : D_N(i,j) = d\}$ is the set of neighbors of node i who are *exactly* d-links away. In our context, these cross-correlations result from network dependence among the errors $\boldsymbol{\varepsilon}_i$'s. This is because

$$\mathbb{E}(\mathbf{v}_i\mathbf{v}_j'|\mathbf{w}_i,\mathbf{w}_j) = \mathbf{h}_i^{(1)}(\boldsymbol{\theta})\mathbb{E}(\boldsymbol{\varepsilon}_i\boldsymbol{\varepsilon}_j'|\mathbf{w}_i,\mathbf{w}_j)\mathbf{h}_j^{(1)}(\boldsymbol{\theta})'$$

In particular, we assume weak dependence which means that these covariances diminish as the network-distance between any two pairs of nodes, i and j, grows large. To limit the strength of dependence, we impose regularity conditions that assume the network data to be conditionally

[26]Such an aggregation helps to express the original panel problem solely in terms of the cross sectional dimension. This matters for the asymptotic argument in our setting where we assume T to be fixed and N to approach infinity.

ψ-dependent given the σ-algebra, $\mathcal{C}_N$.

Note that we have assumed conditional ψ-dependence with respect to $\boldsymbol{\varepsilon}_i$. This is easily translated into conditional ψ-dependence of the linear transformation, $\mathbf{v}_i = \mathbf{h}_i^{(1)}(\boldsymbol{\theta})\boldsymbol{\varepsilon}_i$ given $\mathcal{C}_N$ with dependence coefficients, λ_N, from using Lemma 2.1 of KMS. The other conditions required for the LLN and CLT include moments conditions for $\mathbf{v}_i$, as follows

Assumption 2 (Moment conditions)**.** *i) For some $\epsilon > 0$, $\sup_{N\geq 1}\max_{i\in\mathcal{N}_N}||\mathbf{v}_i||_{\mathcal{C}_N,1+\epsilon} < \infty$ a.s; ii) For some $p > 4$, $\sup_{N\geq 1}\max_{i\in\mathcal{N}_N}||\mathbf{v}_i||_{\mathcal{C}_N,p} < \infty$ a.s. where $||\mathbf{v}_i||_{\mathcal{C}_N,p} = \left(\mathbb{E}[|\mathbf{v}_i|^p|\mathcal{C}_N]^{1/p}\right)$.*

Assumption 3 (Denseness and strength of network dependence)**.** *i) $\frac{1}{N}\sum_{d\geq 0}\zeta_N^{\partial}(d)\lambda_{N,d} \to_{a.s.} 0$ where $\zeta_N^{\partial}(d) = N^{-1}\sum_{i\in\mathcal{N}_N}|\mathcal{N}_N^{\partial}(i,d)|$; ii) There exists a positive sequence, $m_N \to \infty$ such that for $k = 1, 2$*

$$N^{1-k}\cdot\boldsymbol{\Omega}_N^{-k}\sum_{d\geq 0}a_N(d, m_N; k)\lambda_{N,d}^{1-\frac{2+k}{p}} \to_{a.s.} 0$$
$$N^{3/2}\cdot\boldsymbol{\Omega}_N^{-1/2}\lambda_{N,m_N}^{1-1/p} \to_{a.s.} 0$$

as $N\to\infty$, where $p > 4$ and $a_N(\cdot,\cdot;)$ is the measure of network denseness used in KMS.

where $\zeta_N^{\partial}(d)$ gives us the average neighborhood size and $a_N(\cdot,\cdot;)$ is a combination of average neighborhood size and average neighborhood shell size. These are being used as measures of network denseness. As KMS put it, this condition appears in the form of a tradeoff between network denseness and strength of dependence needed to ensure that covariance decays as a function of network-distance, d.

A consistent estimator for the network HAC variance can then be obtained from constructing a sample analogue of the expression in (23) where the middle term is taken to be a weighted average of the sample covariances of $\mathbf{h}_i^{(1)}(\hat{\boldsymbol{\theta}})\hat{\boldsymbol{\varepsilon}}_i$ and $\mathbf{h}_j^{(1)}(\hat{\boldsymbol{\theta}})\hat{\boldsymbol{\varepsilon}}_j$ as follows:

$$\hat{\boldsymbol{\Omega}}_T = \frac{1}{NT^2}\sum_{d\geq 0}\mathcal{K}_N(d)\left[\sum_{i\in\mathcal{N}_N}\sum_{j\in\mathcal{N}_N^{\partial}(i,d)}\mathbf{h}_i^{(1)}(\hat{\boldsymbol{\theta}})\hat{\boldsymbol{\varepsilon}}_i\hat{\boldsymbol{\varepsilon}}_j'\mathbf{h}_j^{(1)}(\hat{\boldsymbol{\theta}})'\right]$$

with $\mathcal{K}_N(d) = \mathcal{K}(d/b_N)$ as the weight given to sample covariances as a function of the distance.

Consistency of $\hat{\boldsymbol{\Omega}}_T$ can be established using the results in KMS if we assume the following conditions hold along with Assumption 2) ii).

Assumption 4 (Consistency of the HAC estimator)**.** *There exists a $p > 4$ such that*

i) $\lim_{N\to\infty}\sum_{d\geq 1}|\mathcal{K}_N(d) - 1|\zeta_N^{\partial}(d)\lambda_{N,d}^{1-2/p} = 0$ a.s. and

ii) $\lim_{N\to\infty}N^{-1}\sum_{d\geq 0}a_N(d, b_N; 2)\lambda_{N,d}^{1-4/p} = 0$ a.s.

C Results Appendix

Figure A1: **Degree Distribution of Nodes in Networks**

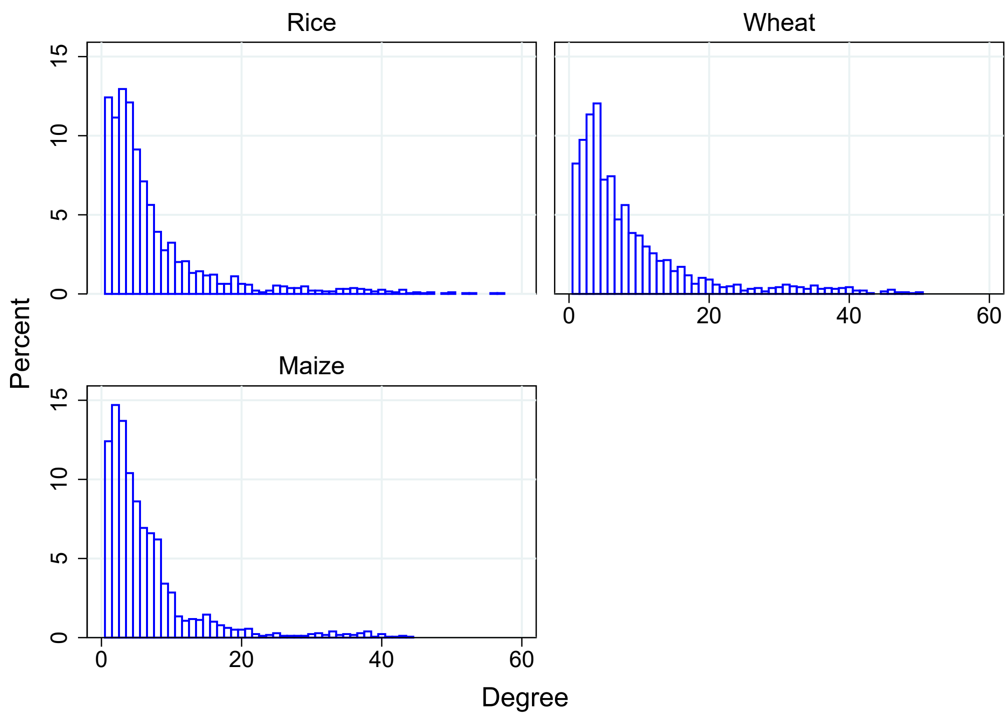

Note: The figure plots the histogram of number of trade links per country for rice, wheat, and maize.

Figure A2: **Average Number of Trade Partners: 1986-2013**

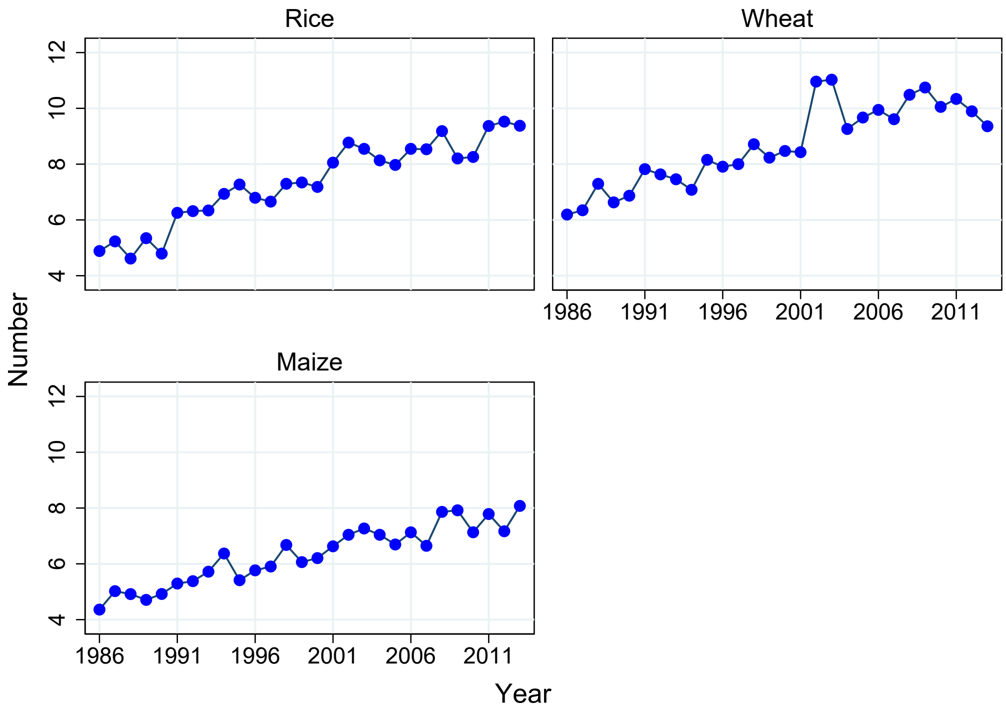

Note: The figure plots trends in the average number of trade partners per country for rice, wheat, and maize.

C.1 Heterogeneous Preferences

As demonstrated in Schulhofer-Wohl (2011), it is more realistic to have agent specific preferences γ_i instead of a constant γ. In this case, the FOC still has an exponential proportional form $c_i^{\gamma_i} = c_j^{\gamma_j}(\alpha_j \delta_{ij}/\alpha_i)$. Just as before, we can, without loss, define $\mathcal{A}^*_{\mathcal{T}}$ to be agent $\bar{i}$'s consumption. Then by the FOC, $c_j = \tilde{\mathcal{A}}^{\gamma_{\bar{i}}/\gamma_j}/(\tilde{\Delta}_{\bar{i}\to j})^{-1/\gamma_j}$, where $\tilde{\Delta} \equiv \Delta^{-\gamma}$. Then the risk-sharing regression has scaled model terms per country:

$$\ln(c_{jt}) = \frac{\ln(\tilde{\Delta}_{jt})}{\gamma_j} + \frac{\gamma_{\bar{i}_t} \ln(\mathcal{A}_{\mathcal{T}_t t})}{\gamma_j} + \eta \cdot \ln(y_{jt}) + \varepsilon_{jt} \tag{25}$$

Note that $\gamma_{\bar{i}_t}$ can change every year depending on the network structure; a previous member chosen for the index could be absent in the network the following year, and thus the identity of $\bar{i}_t = \bar{i}(\mathcal{N}_t)$ changes. Note the presence of the index country's γ term as a coefficient on $\mathcal{A}_{\mathcal{T}_t t}$. This is a network-specific term and thus even if one omits trading costs from the model, the correct model specification with heterogeneous preferences uses trade network information. As the heterogeneous specification has country specific γ_j, it can soak up time-invariant country specific factors that could affect the relationship between production and consumption. Fully addressing all possible alternative explanations for nonzero η beyond transaction costs is beyond the scope of this analysis, but this specification aids in this.

Table A1 reports estimates from the heterogeneous risk preference model in equation (25) for baseline and transaction cost specifications for all three commodities.[27] These regressions include country fixed effects in all specifications. We compare the baseline to the trade-cost-constant specification as the full set of cost controls were not well estimated separately from the country-level slopes. Overall the results are noisier for production, consistent with Schulhofer-Wohl (2011).

For maize, the baseline model yields a production coefficient η of 0.255 (0.137), with the standard error in parentheses. The full model with transaction cost specification yields η equal to 0.1711 (0.08845). This is a 33.1% decrease in η and a 35.3% decrease in its standard error. For wheat, the baseline η is 0.11481 (0.038932), and the model yields 0.11695 (0.066531). This is a 1.86% increase in η and a 71.9% increase in its standard error. Thus, while the coefficient increases a trivial amount, the predictive power of production decreases significantly. Finally, for rice, baseline η is 0.21078 (0.21541), and the model yields 0.18163 (0.20765). This is a 13.8% decrease in the coefficient and a 3.60% decrease in its standard error.

The trade-cost-constant specification under heterogeneous preferences yields larger changes to risk-sharing than in the homogeneous case. This is intuitive as the effects of the network shape are likely dependent on a country's location in the network, which the latter specification averages out.

Table A1: **Heterogeneous Effects**

Commodity:	Maize		Wheat		Rice	
DV: log consumption	Baseline	Trans. Cost	Baseline	Trans. Cost	Baseline	Trans. Cost
Log production	0.255	0.171	0.115**	0.117	0.211	0.182
	(0.137)	(0.088)	(0.039)	(0.067)	(0.215)	(0.208)
Constant	11.548***	12.509***	13.343***	13.493***	13.899***	13.415***
	(0.089)	(0.056)	(0.650)	(3.022)	(0.361)	(0.576)
Trade-cost constant		19.987		19.989***		19.957***
		(19.822)		(1.917)		(0.535)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country γ-FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1752	1752	1722	1722	1860	1860

Notes: Country $\gamma-$FE refers to country specific slopes for preferences. Other parameter estimates omitted. Network dependent standard errors in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

C.2 Additional Results

To see how centrality of the network affects risk sharing, we consider specifications in which we only include countries that are within a certain distance (network-distance) from the index. When only including agents within a 1-link radius from the index, the risk sharing parameter should be smaller than the whole sample. As one increases the number of links included (meaning going away from the center of the network center), risk-sharing would naturally go down; the less connected countries are not as integrated into the trading network. The effects of controlling for transaction costs should also matter at all levels of the radius. For maize, at $\leq$ 3-links, the coefficient decreases by 9.9% from baseline, at $\leq$ 2-links it decreases by 11%, and at 1-link it decreases

[27] For countries without variation over time, we normalize their γ to 1.

by 23%. In addition, at 1-link distance the statistical significance decreases. This indicates that near the center of the network, there is stronger evidence of risk-sharing in Maize after controlling for transaction costs. For the distance from center results, for wheat, $\leq$ 3-links yields a -5.7% change, $\leq$ 2-links has -6.6%, and 1-link has -10.4%. For rice, $\leq$ 3-links yields a -9.6% change, $\leq$ 2-links has a similar decrease, and 1-link has a -5.6% change.

As an alternative sample, we consider a 1 million population cutoff and keep all nonzero trade levels. Results are weaker (smaller or no difference between baseline and full model) and there is less aggregate risk sharing; this is intuitive as many more countries which are less connected (and trade significantly less) are included in this sample. For functional form robustness, first we consider a probit instead of logit for the costs, δ. Second, we consider a more substantial functional form change that still captures the indirect trade effects. We specify the log indirect trade weight as a single logit of the linear difference in covariates across export/import pairs. Across all commodities, we find similar risk-sharing estimates and some small changes to statistical significance of the controls.