Empirics of Air Services Agreements: A Structural Model of Network Formation[*]

Phillip McCalman
Department of Economics
University of Melbourne

Marc Rysman
Department of Economics
Boston University

July 13, 2026

Abstract

Air services agreements are necessary for direct flights between countries and, consequently, are central to the operation of the international commercial airline market. These agreements are bilateral in nature but their coverage is far from universal. To gain insight into why some agreements are formed but not others, we study a new data set on air services agreements from the perspective of strategic network formation. Since an agreement can have implications for countries other than the two directly involved, this is an environment in which externalities are likely to be important. These externalities also suggest that there are multiple equilibria, creating an econometric identification problem. We develop a structural model based on moment inequalities that uses the concepts of pairwise stability and its refinements to generate robust estimating equations. We find that the network structure is important in determining the choices of countries to form agreements and that the jointly optimal network of agreements is substantially different than the observed outcome.

[*] We thank Elie Tamer, Robin Lee, Tom Wollman, Angelo Mele, Richard Davis, Panle Jia Barwick, Dan Ackerberg, Nancy Rose, Krishna Dasaratha, Phil Haile, Steve Berry and many more, as well as several seminar audiences. Excellent research assistance was provided by a surprisingly large number of previous graduate students but Haoran Pan and Cheng Ding deserve special mention. Errors remain our own.

1 Introduction

Air services agreements govern the international airline market. In order to have a commercial flight between two countries, there must be some form of agreement between them, almost always an air services agreement (ASA). Thus, the existing network of ASAs provides structure for the pattern of international flights. Given the importance of direct links for international trade and global commerce more generally, it is somewhat surprising that the coverage of ASAs is far from complete. For the 58 countries we focus on, among the largest countries in the world, more than a third of the country pairs do not have an agreement in our time period. In contrast, we observe that almost all of these country pairs are engaged in bilateral trade. What factors explain the pattern of agreements we observe? Is the realized set of agreements desirable in some broad sense?

In this paper, we address these questions by studying the incentives to form ASAs. We take the perspective of the economic literature on strategic network formation, which models the utility-maximizing choices of agents to form links between each other. In this framework, we interpret countries as nodes (or agents) and ASAs as links that allow communication and trade between nodes. We evaluate whether countries focus on just bilateral characteristics such as the GDP and distance of their partner, or do they also account for the agreements that a partner has signed. To the extent that network structure matters to individual countries, then the choices of pairs of countries to form an agreement has an externality on other countries. We ask what elements of the link structure are important. For instance, are countries particularly interested in connecting to countries that are well-connected, or are countries primarily interested in serving as hubs between unconnected countries (or both)? Finally, we compute the globally optimal network of agreements and compare that to the observed network.

In order to estimate the payoff function of countries and perform counterfactual experiments, we estimate a structural model of the network formation process. Our model follows closely the theoretical literature on strategic network formation, such as discussed in Goyal (2007) and Jackson (2008). In our model, there are no transfers between countries, so a link is formed only if both agents derive positive payoff from the link, which we argue is a realistic description for how ASAs are negotiated. Nash equilibrium in simultaneous choices is often of limited use in network formation contexts, so we follow the theoretical literature and instead study *pairwise stable* allocations.

Even with pairwise stability, predicting outcomes is problematic because for any given set of parameters, there will often be multiple stable allocations. This problem is common in estimating games with strategic interactions, which often generate multiple equilibria. Rather than computing pairwise stable allocations to match to the data, we exploit the fact that pairwise stability requires that observed choices generate higher payoffs than their alternative choices. This revealed preference logic leads to a set of inequality conditions. We map these inequalities into their empirical counterparts, and estimate based on the recent literature on estimating moment inequalities, such as Pakes, Porter, Ho and Ishii (2015). In practice, we follow a version of the method of Cox and Shi (2023) for constructing confidence intervals in partially identified models.

Inspired by refinements to pairwise stability, such as Nash Pairwise Stability and strong stability, we develop additional moment inequalities, and we study how much information these concepts provide in the sense of generating more precise parameter estimates.

In our model, the country-level payoff to an agreement is a reduced-form function of bilateral characteristics and network structure. The bilateral characteristics are similar to those found in the literature on estimating gravity equations (see Head and Meyer, 2014), such as the distance between countries. We use predictions from a gravity equation model directly as an important explanatory variable. For network structure, we focus on measures from the vast empirical literature on social networks (see Prell, 2011; Kolaczyk, 2009), particularly *betweenness centrality* and *eigenvector centrality*. We argue that these capture externalities that we expect to be important in the air services context, such as the desire to be a hub or to be in a dense network. These are related to externalities highlighted in the theoretical literature, such as Jackson and Wolinsky (1996).

Our central data set is called the World Air Services Agreements (WASA) database, and is published by the International Civil Aviation Organization (ICAO), an agency of the United Nations, and was expanded upon by the WTO for 2005. In order to reduce heterogeneity, we select a list of 58 countries that are big, wealthy, or both. An advantage of air services agreements is that they are almost all negotiated bilaterally rather than multilaterally, which makes them similar to existing network theoretical models. A major exception is the European Union, which functions as a large multilateral agreement. We impose various adjustments to address this issue.

A potential concern may be the presence of omitted variables in the determination of the network structure. A country may have many links for some unobserved reason, and failing to account for this may lead a researcher to falsely conclude that countries want to link to countries that have many links. We develop an approach based on countries that make different linking choices across observably similar countries to address unobservable heterogeneity, following Pakes (2010), Pakes et al. (2015), and Wollman (2018), to show that our main result is robust to unobserved heterogeneity. Unobserved heterogeneity is addressed in relatively few of the empirical papers on strategic network formation that we discuss below, so we view this as a contribution of our paper.

We find that network structure plays a statistically significant role in determining choices. In particular, countries form agreements that increase their role as a hub between otherwise unconnected countries. Our moments exploiting refinements of pairwise stability contribute to inference in our main specification.

We use our parameter results to compute the globally optimal network. Our computation addresses two externalities. The first is pairwise, and results from the fact that some agreements do not form even if the sum of payoffs from the agreement is positive because we do not allow side payments between countries. The second is global and arises from the preference to be a hub. When the social planner takes account only of the first issue, the number of agreements grows by more than 16 percentage points. Further addressing the network issue causes the social planner to further reduce the number of agreements by about 1.5 percentage points globally based on a conservative lower bound on the size of the network coefficient. Reductions in the number of agreements allows the planner to create more hubs. Within countries, changes are much larger, with the social planner raising the number of agreements for several countries by more than 50% and reducing for others by as much as a third. Even some countries with little change in the number of agreements see significant reallocation among countries. In general, the model reallocates agreements from small countries to large ones, especially those with strong trading relationships. Overall, we find that the observed set of ASAs differs substantially from the efficient outcome.

From a policy perspective, addressing this issue requires global planning. Currently, the WTO explicitly excludes ASAs from its negotiation, but there have been proposals to bring ASAs under its mandate or otherwise create a more global system of negotiation, and our results support this initiative.[1] An important caveat to our research is that we derive the payoffs to countries from revealed preference about how the countries form agreements, and use this payoff function in calculating counterfactual payoffs. If countries themselves are not maximizing country benefits, perhaps because of political economy issues, we do not address that. While including explanatory variables about political economy in our approach is conceptually straightforward, moment inequality estimation typically can handle relatively few parameters, so we take a parsimonious approach.

2 Literature Review

We study *strategic network formation*, where agents endogenously choose links based on payoffs that depend on the network structure. Early theoretical contributions include Jackson and Wolinsky (1996), with overviews in Jackson (2008) and Goyal (2007). A central empirical challenge in this literature is identification in the presence of unobserved heterogeneity and multiple equilibria.

Our project fits into a growing number of papers that structurally estimate a model of strategic network formation. Recent overviews appear in Graham (2015), de Paula (2020), and Graham and de Paula (2020). Loosely, most previous papers make simplifying assumptions in order to computationally solve outcomes and then use either Bayesian techniques (e.g., Goldsmith-Pinkham and Imbens, 2013; Christakis, Fowler, Imbens and Kalyanaraman, 2020; Mele, 2017; Hsieh and Lee, 2016) or bound the probabilities of outcomes (e.g., Miyauchi, 2016; Sheng, 2020; Gualdani, 2021). Other approaches specify games of imperfect information (Leung, 2015) or search and matching (Currarini, Jackson and Pin, 2009; Boucher and Mourifié, 2017; Badev, 2021). Many of these papers study friendship patterns in surveys of school students.

Relative to this literature, we make several contributions. We provide a new approach based on bounds derived from revealed preference. We exploit refinements to stability, including Nash pairwise stability, which are difficult to incorporate in alternative approaches. Also, we allow for general network externalities rather than restricting interactions to local neighborhoods. In

[1] Havel (2009) is an example advocating for a global system. He discusses support for using the WTO (or GATS) on page 540. GATS itself has a provision to eventually address air traffic rights: Paragraph 5 of the Air Transport Annex states that "the Council for Trade in Services shall review periodically, and at least every five years, developments in the air transport sector and the operation of the annex with a view to considering the further application of the Agreement to the sector." Note that while traffic rights and services directly related to the exercise of traffic rights are currently outside of the scope of the WTO, GATS does cover aspects of the air services sector, in particular: computer reservation system (CRS) services, selling and marketing of air transport services and aircraft repair and maintenance.

addition, our method tractably accommodates large networks and connects empirical measures of network position, such as centrality, to underlying economic primitives. In addition, our focus on international industrial policy is substantially different than other empirical papers on strategic network formation.

Our approach also relates to a broader literature using moment inequalities to study network and industrial organization problems (e.g., Ho, 2009; Holmes, 2011; Houde, Newberry and Seim, 2023), and recent applications in transportation and infrastructure networks (e.g., Bontemps, Gualdani and Remmy, 2023; Yuan and Barwick, 2024; Jeon and Rysman, 2025).

Finally, our application contributes to the literature on international agreements. Existing work on trade agreements typically relies on reduced-form binary choice models that abstract from equilibrium multiplicity (e.g., Baier and Bergstrand (2004), Egger, Larch, Staub and Winkelmann (2011), Chen and Joshi (2010); see Limão (2016) for an overview).[2] In contrast, our framework explicitly accommodates multiple equilibria in network formation. A related empirical literature studies the effects of air services agreements on trade and transportation outcomes (e.g. Cristea, Hummels and Roberson (2017), Micco and Serebrisky (2006)). Rather than treating agreements as exogenous, we model their formation as an equilibrium outcome.

3 Industry and Data

The direct benefit of an agreement is that it allows for direct flights between countries. In addition, the effects of an ASA can easily expand beyond air travel. For instance, if air travel supports communication between executives at trade conferences, the impact of an ASA may be much larger in shipping than air. Recognizing the importance of the nascent air services industry, a large group of countries met in Chicago in 1944 to work out an international agreement on how the industry should be managed.[3] The so-called Chicago Convention failed to reach a comprehensive international agreement but instead led to a treaty that established a framework under which subsequent bilateral agreements would follow, establishing a template for subsequent ASAs. ASAs determine what rights partner-country airlines have, such as whether to pick up passengers or cargo, and also determine issues such as what routes are allowed, how many airlines are allowed on the routes, and whether price changes require government approval.[4]

The Chicago Convention also established ICAO (named as such later, when it became part of the United Nations), which coordinates international standards on flight and airport regulations. ICAO maintains a data set of all ASAs, the World Air Services Agreements (WASA) Database. However, the ICAO process for collecting data led to many missing agreements. For the 2005 data, the WTO undertook a project to capture missing agreements. The WTO data set is the centerpiece of our project. Thus, our project focuses on 2005. The WTO data is a cross-section and

[2] In Appendix A, we discuss a simple example of why ignoring the multiplicity of equilibria in these approaches causes an estimation problem.

[3] Odoni (2009) provides an excellent discussion of the institutional background for air services agreements.

[4] Note that these deals are termed international *agreements* rather than *treaties*. Agreements are typically easier to negotiate. For instance, in the United States, agreements do not require congressional approval, whereas treaties do. From the perspective of airlines, both have the force of law.

contains data only on whether an agreement exists, not on the content of an agreement. For more description of the ICAO data, the WTO data collection process, and our attempts similar to the WTO process, see Appendix B.

To keep our data reasonably homogeneous, we do not include every country in the world in our data set. In order to choose countries, we take the top 50 countries by population and the top 50 by GDP and form the union. From this, we drop countries with a population less than 500,000.[5] We add back in Iceland although its population is less than 500,000 because Iceland has a prominent national airline and an important geographical location. Overall, we have a list of 58 countries from 6 continents. This creates 1,653 unordered country pairs, or 3,306 directed country pairs. When we use ”country pairs” on its own, we mean unordered country pairs.

The Chicago Convention envisioned only bilateral agreements, which is helpful for our purposes because multilateral agreements are more difficult to model in the context of strategic network formation. However, in practice, the EU has functioned as a multilateral air services agreement between all members since 1992. Up until 2002, EU countries continued to sign bilateral agreements with non-EU countries, rather than having the EU negotiate as a whole.[6] Since 2005, several more multilateral air services agreements have arisen or have been proposed, particularly in Asia, and so the fact that we use 2005 data helpfully avoids this issue.

From the perspective of the social networks literature, we term the countries as *nodes* and the agreements as *links*. We assume that all EU countries are linked to each other, but that EU countries independently form links to countries outside of the EU.[7] This reflects the state of the EU in the early 2000s, which we believe best describes our 2005 data.[8] We calculate network statistics including all EU-to-EU links. However, when forming moments for estimation, we do not include EU-to-EU pairs because the formation of the EU air services agreement reflects a different process. Thus, while we observe 1,653 unordered country pairs, we exclude 253 EU-to-EU unordered pairs from the estimation moments, leaving 1,400 non-EU country pairs, or 2,800 directed country pairs

In addition, we utilize the CEPII data set created for Head, Mayer and Ries (2010), which contains several useful bilateral variables, such as distance and indicators for a common border and a common language. See Conte, Cotterlaz and Mayer (2022). For variables that vary over time, such as GDP and population, we use the average of 2001-2005.

[5]Dropping countries with low populations eliminates Antigua, the Bahamas, Barbados, Brunei, Equatorial Guinea, Iceland, Luxembourg, Macau, Malta, St. Kitts, and the Seychelles. We also drop Puerto Rico and Taiwan, whose freedom to negotiate agreements is unclear, at best.

[6]The use of bilateral agreements was subject to a lawsuit filed by the European Commission against member countries. The court found that member countries had the right to negotiate bilateral agreements under the Treaty of the European Union, but that member countries did not have the right to restrict the benefits only to their own national airlines. While some EU countries have updated their bilateral agreements accordingly, bilateral agreements hold little appeal without this last element, and agreements since then have tended towards multilateral agreements between an outside country and all EU countries simultaneously. See `https://transport.ec.europa.eu/transport-modes/air/international-aviation/external-aviation-policy_en`. von den Steinen and Probst (2013) describes some challenges arising in the transition to multilateral agreements.

[7]In recognition of these issues, the WTO did not study EU-to-EU pairs in their augmentation of ICAO data.

[8]In our definition of “EU,” we include the 25 states that joined by 2005 (Bulgaria and Romania joined in 2007) plus Norway, Lichtenstein, Iceland, and Switzerland, which signed the European Common Aviation Area (ECAA) agreements by 2005 (see European Commission, 2010). Note that Lichtenstein is dropped from our final data set because it has a low population. We do not count the Western Balkan states, which were not covered by the European Common Aviation Area in our 2005 data.

A natural question is why countries do not sign ASAs with every other country. Signing agreements has a cost in terms of administration and negotiator time, and also requires ongoing enforcement, which may require registering complaints with foreign countries. In the United States, the State Department has an Aviation Negotiation division that is dedicated to negotiating and enforcing ASAs.

4 Social Network Variables

In our empirical work, we ask how a country values connectedness. Some notation aids in developing concepts of connectedness. Let the set of countries be $\mathcal{N} = \{1, \ldots, n\}$. Let Λ be an $n \times n$ matrix. We denote element $\{i, j\}$ of matrix Λ as λ_{ij}, with $\lambda_{ij} = 1$ if countries i and j have an ASA, and $\lambda_{ij} = 0$ otherwise. Thus, Λ represents the set of links between nodes. Links are undirected, so $\lambda_{ij} = \lambda_{ji}$. The diagonal of Λ is 0 by assumption. Let the function $s_i(\Lambda)$ return the set of countries that are linked to i in network Λ. That is, $s_i(\Lambda) = \{j : \lambda_{ij} = 1\}$. *Degree* is defined as the cardinality of $s_i(\Lambda)$, and is the number of links a node has.

We would also like an index that reflects not only how many links a country has, but also how many links the partner countries have, and the partners of those partners in turn. Labeling this vector of indices as C^{eig}, we define C^{eig} implicitly as:

$$aC^{eig} = \Lambda C^{eig}$$

where a is a proportionality factor. Thus, $C^{eig}(\Lambda)$ is an eigenvector of Λ, and a is the corresponding eigenvalue. The convention is to use the eigenvector associated with the highest eigenvalue, and term $C_i^{eig}(\Lambda)$ the *eigenvector centrality* for node i.

Next, we consider the sense in which a node sits between other nodes. A *geodesic* is a shortest path between two nodes. There may be multiple geodesics between any two nodes. For instance, in our data, countries that do not have a link can always reach each other in two links, but there may be multiple paths by which to do so. Let $P(j, k)$ be the number of geodesics between nodes j and k. Let $P_i(j, k)$ be the number of geodesics between j and k that pass through i. *Betweenness centrality* is:

$$C_i^{betw}(\Lambda) = \sum_{\{j,k: j \neq i, k \neq i, j \neq k\}} \frac{P_i(j, k)}{P(j, k)}.$$

To see how these two measures differ, consider Figure 1. The figure has five nodes, A through E, and five links. The figure displays the nodes and links, as well as the resulting measures of eigenvector and betweenness centrality.[9] We do not display degree, but it is clear: node C has degree 3, node E has degree 1, and the rest have degree 2. Node C is central in this network and has the highest score in both eigenvector and betweenness centrality. However, note the differences between A, B, and D. Nodes A, B, and C make up a cluster that generates the high eigenvector scores for A and B. However, A and B are not hubs – there are no shortest paths

[9]For clarity, we display $P_i(kj)$ instead of C_i^{betw}.

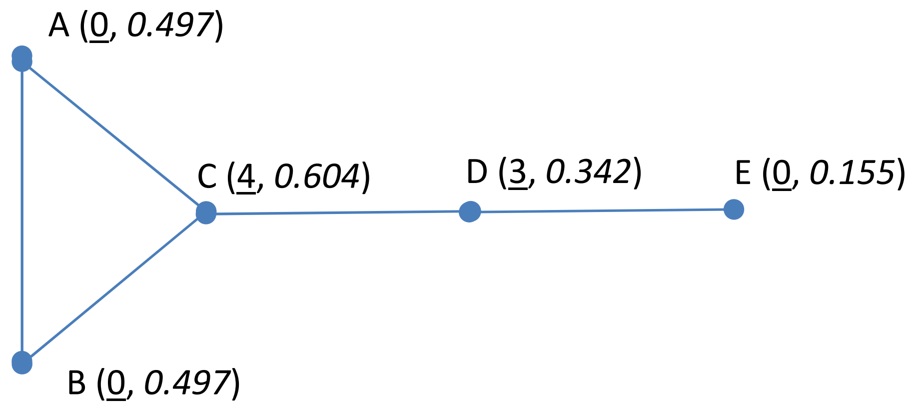

Figure 1
Network Example. Betweenness centrality is underlined, Eigenvector centrality is in italics.

between two nodes that run through A or B. In contrast, D is relatively more isolated than A or B, but is the only route to E from any other node. Thus, D is higher than A and B in betweenness centrality but lower in eigenvector centrality.

Although we do not provide microfoundations (i.e., a game theoretic model in which these measures arise) for these measures, we believe that these statistics can be interpreted to capture some important constructs from the theoretical network formation literature. We focus on two models introduced by Jackson and Wolinsky (1996), the *co-author model*, which focuses on agents that prefer links to nodes that are not linked to others, and the *connections model*, which emphasizes agents that benefit from links to nodes with many links.

In the model of the co-author externality, if player A is connected to B and C, A may be better off if B and C are not connected to each other. In the model, each player has a limited amount of time to devote to each link, and so the connection between B and C takes away time from A's links. In the context of airlines, this is equivalent to saying that A values being a hub between B and C. We believe that the hubbing externality is well captured by betweenness centrality. If we observe a country forming agreements that tend to raise the country's betweenness centrality, we infer that it prefers to be a hub and that the coauthor externality is present. We do not restrict the sign of the coefficient on betweenness centrality. It may be that the link between countries B and C raises growth so much for those countries that traffic to A actually increases, in which case there is still an externality, but with the opposite sign.[10]

The connections model recognizes that if player A is connected to B but not C, then A gains when B connects to C because now A can reach C. This set-up has a natural interpretation in the airlines context: More paths of links between any two countries provide more options for

[10]Note that as Jackson and Wolinsky (1996) define the coauthor externality, the agent is worse off whether B and C form links with each other or anyone else, since any sort of links take time away from their relationship with A, whereas our concept of hubbing would mean that it is particularly the link between B and C that hurts A. In this sense, our concept of hubbing differs from the coauthor externality. We believe that betweenness centrality better captures hubbing than the coauthor externality in this sense.

that country pair. We believe that this measure is well captured by eigenvector centrality, which measures the overall connectedness of any given node. Thus, we look for countries that create agreements to increase their eigenvector or betweenness centrality in order to infer which of these issues are present. In either case, B and C make their decision without accounting for A, so we have an externality.

Now we turn to some simple statistics. The US, Russia, China, and Singapore have the highest degrees (numbers of links) of non-European countries, with 53, 48, 46, and 46 respectively. Counting EU-to-EU links, France and Germany are connected to every country in the world and Switzerland and the UK are connected to all but one each. European countries make up most of the top-10. The mean number of links per country is 37.03, the standard deviation is 11.2, the minimum is 12 and the interquartile range goes from 31 to 46. Thus, there is substantial variation in the number of links.

Our network is highly connected. Connected paths exist between every node and never require more than one step. We find that 65% of country pairs are directly connected, and 35% of country pairs require one intermediate node to reach each other. Excluding EU-to-EU country pairs, 58.6% of pairs have links. This is the outcome variable our model seeks to explain.

Table 1 presents the network statistics that we are interested in. These statistics are calculated including EU-to-EU links. The table is ordered by degree and presents the top 30 countries. It also presents betweenness centrality and eigenvector centrality, along with the rank of each country in these statistics.[11] Note that eigenvector centrality is typically normalized by dividing by the square root of the sum of the squared eigenvector (see Borgatti, Everett and Freeman, 2002). However, to enhance comparability here and for the rest of the paper, we normalize eigenvector centrality to have the same mean as our measure of betweenness centrality.

We see that France and Germany are connected to 57 countries (every other country) and are ranked highest in both types of centrality. As we move down the list, we see larger differences between relative eigenvector and betweenness centrality, which is what allows us to determine which is more important for link formation. For instance, Spain is ranked 11th in eigenvector centrality but 7th in betweenness centrality. Similarly, Canada is 24th and 11th. In contrast, Thailand is 12th in eigenvector centrality but 17th in betweenness and Japan is 18th and 23rd. While clearly correlated, the measures of centrality capture different issues.

A potential concern is that our measures count links as if all links are equal. We might prefer a measure of betweenness centrality that reflects not only the number of paths that pass through a node, but also whether those paths are important. For instance, being between two countries that are close together should matter more than being between countries that are far apart, or being between countries that trade a lot should matter more than being between countries that do not. Brandes and Erlebach (2005) (among others) discuss the concept of *weighted network statistics*, and we use those ideas to construct weighted versions of betweenness centrality and eigenvector centrality. For weights, we use (the inverse of) distance. We have also experimented with trade as

[11] Because two countries are connected to every other country, all countries are either directly connected or connected by a path of no more than one step. Thus, betweenness centrality is akin to a normalized count of how many times a country is connected to both of two unconnected countries. We view this as a reasonable way to capture hubbing.

Country	Degree	Eigenvector Centrality		Betweenness Centrality	
		Rank	Value	Rank	Value
France	57	1	14.66	1	44.58
United Kingdom	57	1	14.66	1	44.58
Germany	56	3	14.58	4	38.94
Switzerland	56	4	14.57	3	39.37
Netherlands	54	5	14.33	6	31.10
United States	54	6	14.21	5	33.82
Belgium	52	7	13.93	7	28.00
Austria	49	8	13.59	11	15.89
Denmark	48	9	13.25	9	18.06
Russia	47	10	13.10	12	15.88
Spain	47	13	12.84	8	19.31
Italy	46	11	12.89	14	14.37
Singapore	46	12	12.85	15	13.26
Sweden	46	14	12.77	10	16.17
China	44	17	12.37	17	12.15
India	44	15	12.50	21	9.92
Pakistan	44	19	12.23	13	15.85
Turkey	44	16	12.45	16	12.32
Thailand	43	18	12.35	22	8.95
Czech Republic	42	20	11.97	19	10.69
Canada	40	25	11.26	18	11.22
Egypt	40	22	11.56	23	8.43
Japan	40	23	11.51	25	7.75
Norway	40	21	11.60	26	7.18
Greece	39	24	11.37	33	6.00

Table 1: Comparing centrality statistics.

Table reports top 25 countries by degree (number of links) in order of degree. For eigenvector centrality and betweenness centrality, the table reports the value and the rank of that value among the 58 countries.

a weighting variable. We discuss the exact calculation in Appendix C, but ultimately we find that our results are robust to using weighted statistics.

5 Model and Refinements

In this section, we present our theoretical model of network formation and we derive inequalities that we bring to estimation. We then discuss refinements to the solution concept and the new inequalities that they generate. We delay discussion of omitted variable bias due to unobserved heterogeneity until Section 8.

5.1 Model

Let the vector x_{ik} describe observable characteristics of i and k, such as GDP and population, and bilateral characteristics, such as the product of GDPs, distance, and perhaps bilateral trade flows. In our preferred specification, we use only two variables for x_{ik}, a constant term and what we term the *gravity score*, the prediction from a gravity equation model that we estimate on trade data, that we discuss below. The function $\psi_i(\Lambda, \alpha)$ captures the payoff to i from a given link structure Λ, parameterized by α. In our approach, it is a statistic such as betweenness centrality, i.e., $\psi_i(\Lambda, \alpha) = \alpha C_i^{betw}(\Lambda)$. The payoff to country i from the network of links Λ is:

$$\Pi_i(\Lambda) = \psi_i(\Lambda, \alpha) + \sum_{k \in s_i(\Lambda)} (x_{ik}\beta + \varepsilon_{ik}).$$

In this expression, the first term represents the overall network benefit, whereas the second term captures the sum of bilateral benefits. Note that we proceed as if $\Pi_i(\Lambda)$ represents true welfare for the country and the country makes decisions to optimize this variable. In practice, decision-making may be manipulated by issues of political economy. Naturally, we can include variables representing political economy concerns in x_{ik}.

The payoff to i from linking to j is:

$$\Pi_{ij} = \psi_i(\Lambda \cup \{i,j\}, \alpha) + \sum_{k \in s_i(\Lambda \cup \{i,j\})} (x_{ik}\beta + \varepsilon_{ik}).$$

Here, $\Lambda \cup \{i,j\}$ represents the network Λ including a link between countries i and j. If i and j are already linked in Λ (i.e., $\lambda_{ij} = 1$), then: $\Lambda \cup \{i,j\} = \Lambda$.

The payoff to not linking to j is:

$$\Pi_{i-j} = \psi_i(\Lambda - \{i,j\}, \alpha) + \sum_{k \in s_i(\Lambda - \{i,j\})} (x_{ik}\beta + \varepsilon_{ik}).$$

Naturally, $\Lambda - \{i,j\}$ indicates network Λ with $\lambda_{ij} = 0$.

Thus, country i benefits from a link with j if:

$$\begin{aligned} \pi_{ij} &= \Pi_{ij} - \Pi_{i-j} \\ &= \psi_i\left(\Lambda \cup \{i,j\}, \alpha\right) - \psi_i\left(\Lambda - \{i,j\}, \alpha\right) + x_{ij}\beta + \varepsilon_{ij} \geq 0. \end{aligned} \tag{1}$$

We assume that the agent measures one of the elements of x_{ik} with error, which is captured by ε_{ik}. That is, we break up x_{ik} into $\{x_{ik1}, x_{ik2}\}$ where x_{ik1} is the first variable in the vector x_{ik}, and x_{ik2} is the remaining vector of variables. Suppose the measurement error is over the first variable, x_{ik1}. The country observes $\tilde{x}_{ik1}$, where $\tilde{x}_{ik1} = x_{ik1} + \varepsilon_{ik}$, and $E[\varepsilon_{ik}] = 0$. The country makes its decisions based on $E[x_{ik1}|\tilde{x}_{ik1}]$, which equals $\tilde{x}_{ik1}$. Because the country's decision depends only on $\tilde{x}_{ik1}$ and x_{ik2} and not on ε_{ik}, we have that $E[\varepsilon_{ik}|\tilde{x}_{ik1}, x_{ik2}, \Lambda] = 0$. In fact, conditioning on any subset of the agent's information leads to the same result. Pakes et al. (2015) and Dickstein and Morales (2018) further discuss how ε_{ik} can be interpreted either as measurement error or expectational error on the part of the agent. Countries make decisions about agreements based on whether they expect the payoff to be positive, i.e., $E\left[\pi_{ij}|\mathfrak{I}_i\right] \geq 0$ where $\mathfrak{I}_i$ is the information set for country i, such as $\mathfrak{I} = \{\Lambda, \tilde{x}_{i1}, x_{i2}\}$.

We believe that it is reasonable to assume that a country cannot predict its payoff from an agreement exactly. However, it differs from the usual treatment of discrete choice variables, which assumes there is a structural error term that explains why observationally identical agents make different choices, and would not be mean independent of the outcome Λ. In Section 8, we discuss adding a structural error term, that is, an error term that is known to the country.

One approach to solving the game might be to specify strategies and then solve for a Nash equilibrium. However, this approach tends to generate unrealistic equilibria, such as when no players link with anyone (Myerson, 1977, considers such a game). Instead, the literature has focused on pairwise stability:

Definition A network Λ is *pairwise stable* if:

1. $E[\pi_{ij}|\mathfrak{I}_i] \geq 0$ and $E[\pi_{ji}|\mathfrak{I}_j] \geq 0 \quad \forall \quad \{i,j : \lambda_{ij} = 1\}$.
2. $\min\{E[\pi_{ij}|\mathfrak{I}_i], E[\pi_{ji}|\mathfrak{I}_j]\} < 0 \quad \forall \quad \{i,j : \lambda_{ij} = 0\}$.

The expectation is taken over ε_{ij}, which is assumed to be orthogonal to decision-making. Condition 1 rules out unilateral deletion of an existing link: if $\lambda_{ij} = 1$, both countries must weakly prefer the link to no link. Condition 2 rules out mutually profitable link formation: if $\lambda_{ij} = 0$, at least one of the two countries must prefer not to form the link. This distinguishes pairwise stability from Nash equilibrium analysis, which generally does not allow two unlinked players to form a link through a bilateral deviation.

This formulation implicitly assumes that utility is non-transferable between agents. If one agent does not benefit from a link, the agent deletes the link and there is no opportunity for the partner to pay the agent to maintain the link. We find a model with non-transferable utility more natural in this environment, and more consistent with our reading of the institutional literature and discussions with knowledgeable experts, which suggests there was little in the way of "side payments" available.[12] However, estimating a model with transferable utility is feasible, and we do so as a robustness check. With transferable utility between pairs of agents, links are formed based on the sum of payoffs. The definition would be 1) $E[\pi_{ij}+\pi_{ji}|\mathfrak{I}_i,\mathfrak{I}_j] \geq 0$ for all $\{i,j:\lambda_{ij}=1\}$ and $E[\pi_{ij}+\pi_{ji}|\mathfrak{I}_i,\mathfrak{I}_j] < 0$ for all $\{i,j:\lambda_{ij}=0\}$.[13]

We now turn to estimation. To ease notation, we let $\Psi_i(\Lambda,j,\alpha)=\psi_i(\Lambda\cup\{i,j\},\alpha)-\psi_i(\Lambda-\{i,j\},\alpha)$. For any pair $\{i,j\}$, let $-i=j$ and $-j=i$. Also, let z_{ij} be a vector of variables observed by the econometrician that country i uses to construct $\tilde{x}_{ij1}$, its prediction of x_{ij1}. Then, $E[\varepsilon_{ij}|\tilde{x}_{ij1},x_{ij2},\Lambda]=0$ implies $E[\varepsilon_{ij}|z_{ij},x_{ij2},\Lambda]=0$ (Dickstein and Morales, 2018). This is true even if z_{ij} is a subset of variables that i uses to predict x_{ij1} rather than the complete set. Intuitively, z_{ij} are akin to instruments for the endogenous variable x_{ij1}. With this notation, we exploit the definitions of pairwise stability to generate moment inequalities, following Pakes et al. (2015):

$$\begin{aligned} E\left[\Psi_i(\Lambda,j,\alpha)+x_{ij}\beta\middle|\lambda_{ij}=1,\Lambda,z_{ij},x_{ij2}\right] &\geq 0 \qquad (2)\\ E\left[\min_{k\in\{i,j\}}\left\{\Psi_k(\Lambda,-k,\alpha)+x_{k,-k}\beta\right\}\middle|\lambda_{ij}=0,\Lambda,z_{ij},x_{ij2},x_{ji2}\right] &< 0. \qquad (3)\end{aligned}$$

These equations capture that for any pair with an agreement, we know that both countries prefer an agreement to no agreement, but for each pair without an agreement, we know only that one of the two countries prefers no agreement.

An alternative moment to Equation 2, equally well motivated by the theoretical model, is that for each country, its lowest payoff among links rather than its average payoff is positive. Formally:

$$E\left[\min_{j\in s_i(\Lambda)}\left\{\Psi_i(\Lambda,j,\alpha)+x_{ij}\beta\right\}\middle|s_i(\Lambda)\neq\emptyset,\Lambda,Z,X_2\right]\geq 0, \qquad (4)$$

where Z and X_2 are the matrices of z_{ij} and x_{ij2}. Equation 4 should be more informative than Equation 2 because requiring the minimum of a set to be positive is more binding than requiring the average to be positive. However, we treat this moment as if it has only one observation per country (further discussed in Section 6.3) and so it is estimated with substantially less precision. We apply both moments, and we explore which moment is more empirically informative below. We refer to the moment inequality in Equation 4 as the *min moment*.

5.2 Refinements of pairwise stability

One criticism of pairwise stability is that it allows for over-connected networks. A set of links can satisfy pairwise stability if agents do not want to sever any single link, even if an agent would wish to simultaneously sever multiple links. This idea leads to a refinement of pairwise stability called *Nash pairwise stability*, and is discussed in Jackson and Wolinsky (1996). Similarly, an agent

[12]For example, (Havel, 2009, pg. 10) writes "The virtually universal approach among members of ICAO (the 190 parties to the Chicago Convention) is that air services treaty negotiations are conducted separately from general trade diplomacy." and continues "Thus, U.S. 'payments' for more liberal aviation regimes must be made in the context of the aviation agreement itself."

[13]If we instead allowed transfers between any agents, so for instance, agent i could pay j and k based on whether j and k form a link, then the network would be at its social optimum. We could still form inequalities based on whether the sum of all payoffs went up or down with each link that we observe or do not, so estimation is also feasible in this context.

may benefit from forming links with multiple agents simultaneously, even if each individual link was not valuable. The possibility of deviating from an allocation by adding multiple links is also not captured by pairwise stability. This concept is addressed by the notion of *strong stability* (Dutta and Mutuswami, 1997; Jackson and van den Nouweland, 2005). Such a deviation requires coordinated action between multiple agents, whereas the deletion of multiple links requires only unilateral actions.

An advantage of our approach is that it is relatively easy to develop moment inequalities capturing refinements such as Nash pairwise stability and strong stability, and to study how these moments affect outcomes. Imposing additional moments is particularly attractive because moment inequalities frameworks lead to partial identification and thus wide ranges for parameter estimates. The extra information embedded in stability refinements could be a valuable source of precision over parameter estimates. However, we observe many players with many potential sets of links that they could add and delete, so it is difficult to consider every possible such deviation from an allocation. Instead, we consider deviations where players are allowed to alter a limited number of links.

Some notation is helpful. Let k be a scalar and let $\mathcal{S}_{ik}^{+}$ be the set of all combinations of k elements of $\mathcal{N} - s_i(\Lambda)$. That is, if $k = 1$, then $\mathcal{S}_{i1}^{+}$ is the set of individual countries to which i is not linked, the set of potential additional single links. If $k = 2$, then $\mathcal{S}_{i2}^{+}$ is the set of all pairs of countries that i could add. Similarly, let $\mathcal{S}_{ik}^{-}$ be the set of all combinations of k elements of $s_i(\Lambda)$. Thus, $\mathcal{S}_{i2}^{-}$ would be all pairs of nodes that i is linked to, and thus could potentially delete. For example, in Figure 1, $\mathcal{S}_{C2}^{-} = \{AB, AD, DB\}$ and $\mathcal{S}_{E2}^{-} = \emptyset$. Denote each element of $\mathcal{S}_{ik}^{+}$ as σ_{mik}^{+}, $m = 1, \ldots, \#\mathcal{S}_{ik}^{+}$, and similarly for $\mathcal{S}_{ik}^{-}$. Our first new stability concept is motivated by Nash pairwise stability and our computational concerns:

Definition A network Λ is *pairwise stable with deletion degree K* if it is pairwise stable and:

$$E[\Pi_i(\Lambda) - \Pi_i(\Lambda - \sigma_{mik}^{-})|\Im_i] \geq 0 \quad \forall \quad i = 1, \ldots, n, \; k = 1, \ldots, K, \sigma_{mik}^{-} \in \mathcal{S}_{ik}^{-}.$$

The definition rules out networks in which agents wish to delete sets of links up to size K. This stability concept rules out more networks than pairwise stability but less than Nash pairwise stability. Our stability concept is equivalent to Nash pairwise stability if $K \geq n - 1$.[14]

We can generate additional moment inequalities by exploiting conditions from these stability conditions. For example, if $K = 2$, we have an observation for each pair of countries that any country is connected to. Three countries that are each connected to the other two generates three observations.

$$E\left[\Psi_i\left(\Lambda, \sigma_{mik}^{-}, \alpha\right) + \sum_{j \in \sigma_{mik}^{-}} x_{ij}\beta \,\middle|\, \sigma_{mik}^{-} \in \mathcal{S}_{ik}^{-}, \Lambda, Z, X_2\right] \geq 0 \tag{5}$$

Here, $\Psi_i\left(\Lambda, \sigma_{mik}^{-}, \alpha\right) = \psi_i\left(\Lambda, \alpha\right) - \psi_i\left(\Lambda - \sigma_{mik}^{-}, \alpha\right)$. This approach generates a separate moment for each value of $k \leq K$. So we add $K - 1$ moments to those derived for pairwise stability (plus

[14]There is no guarantee of existence of a Nash pairwise stable allocations, or our stability concept. A sequence of bilateral joins could lead to a set of links that the country would like to delete.

possible interactions with instrumental variables such as x). In practice, we consider $K = 2$. That is, we consider whether countries would delete pairs of links.[15]

We can also define a *min moment* analogous to Equation 4 for this case, which says that each country does not want to delete its worst pair of links.

$$E\left[\min_{\sigma^-_{mik}\in\mathcal{S}^-_{ik}}\left\{\Psi_i\left(\Lambda,\sigma^-_{mik},\alpha\right)+\sum_{j\in\sigma^-_{mik}}x_{ij}\beta\right\}\middle|\mathcal{S}^-_{ik},\Lambda,Z,X_2\right]\geq 0 \tag{6}$$

Similar to Equation 5, we can define a stability concept motivated by strong stability:

Definition A network Λ is *pairwise stable with addition degree K* if Pairwise stability holds, and:

$$\min\left\{E\left[\Pi_i\left(\Lambda+\sigma^+_{mik}\right)-\Pi_i\left(\Lambda\right)\middle|\Im_i,\right],E\left[\pi_{ji}\middle|\Im_j\right]\forall j\in\sigma^+_{mik}\right\}<0$$
$$\forall\quad i=1,\ldots,n,\ k\leq K,\ \sigma^+_{mik}\in\mathcal{S}^+_{ik}. \tag{7}$$

Equation 7 says that adding sets of partners reduces the payoff to i, or i cannot find a set of willing partners. We have written this equation so that the set of links is not added either because country i does not want to add the set or because one of the potential partners does not want to link to i.[16]

In practice, we expect there to be little benefit to setting K higher than 2 or 3. To take an extreme example, what would be the value of going from $K = 56$ to $K = 57$? We observe only two countries with 57 links, so we would be forming a moment with two observations, which will have little power over our parameters. Furthermore, we expect a very wide set of parameters to satisfy the condition that a country is better off not deleting 57 of its links, so even if there were a few more countries in this set, the moment would still be of little value.

6 Estimation

This section provides sample analogs of our moments and then discusses empirical issues with implementation. A short discussion of asymptotic properties then follows.

6.1 Sample analogs

In order to define the sample analog of pairwise stability, let I^a be the set of directed country pairs (i, j), with $i \neq j$, for which $\lambda_{ij} = 1$, and let I^{na} be the set of directed country pairs (i, j), with $i \neq j$, for which $\lambda_{ij} = 0$. We exclude directed pairs (i, j) in which both i and j are members of the EU.

[15] Computationally, we implement Equation 5 by generating a data set that stacks all of the elements in $\mathcal{S}^-_{ik}$ for each country i. Even when a few countries have more than 30 links, this data set is not particularly large for $k = 2$, and this data set is created in advance of estimation. During estimation, checking linear profit conditions for each observation is computationally cheap. Thus, exploring $K > 2$ appears feasible, although we have not done so yet, for reasons described in the results.

[16] The concept of strong stability allows for deviations that include simultaneous deletion and addition of multiple links. We have not explored this, but developing additional moments expanding on what we have is straightforward.

Let $n^a = |I^a|$ and $n^{na} = |I^{na}|$. Furthermore, let z_{ij} be a vector of instruments for directed pair (i, j). Recall that for any pair $\{i, j\}$, we write $-i = j$ and $-j = i$. Then, our sample moments are:

$$\frac{1}{n^a} \sum_{(i,j)\in I^a} \left(\Psi_i(\Lambda, j, \alpha) + x_{ij}\beta\right) z_{ij} \geq 0 \tag{8}$$

$$\frac{1}{n^{na}} \sum_{(i,j)\in I^{na}} \min_{k\in\{i,j\}} \left\{\Psi_k(\Lambda, -k, \alpha) + x_{k,-k}\beta\right\} z_{ij} < 0. \tag{9}$$

Moment inequalities lead to partial identification. In order to determine whether a point is in the confidence interval, we follow the method in Cox and Shi (2023). Their method generates a closed form for critical values, whereas alternatives typically rely on simulation, so Cox and Shi (2023) is faster and easier to implement. Most of our parameter results are found by searching for the maximum and minimum of each parameter separately, subject to the constraint that the set of parameters must fall within the confidence set. We then report only these maximum and minimum values, rather than the shape of the confidence intervals. Many of our specifications use only two parameters so we use grid searches and graphs of the confidence set to gain a deeper understanding of our estimator. A fuller discussion of our estimation approach appears in Appendix D.

In addition to the moments in Equations 8 and 9, we utilize two more sets of moments. The first is derived from the stability refinements discussed in Section 5.2. Just as Equation 8 is an empirical counterpart to Equation 2, we develop a sample analog to Equation 5. We implement only the case for $K = 2$. Let $\tilde{n}^-$ be the number of observations for this moment, so $\tilde{n}^- = \sum_{i=1}^{n} \#\mathcal{S}_{i2}^-$, the total number of elements of all sets $\mathcal{S}_{ik}^-$ for all $i = 1, \ldots, n$.

$$\frac{1}{\tilde{n}^-} \sum_{i=1}^{n} \sum_{\sigma_{mi2}^- \in \mathcal{S}_{i2}^-} \left(\Psi_i\left(\Lambda, \sigma_{mi2}^-, \alpha\right) + (x_{ij} + x_{il})\beta\right) z_{ijl} \geq 0, \quad \{j, l\} = \sigma_{mi2}^-. \tag{10}$$

In this case, z_{ijl} refers to the instrument for the observation in which i is linked to j and l. There is an analogous moment derived from the definition of pairwise stability with addition degree 2.

A second set of moments that we utilize are the *min moments*, as presented in Equation 4. Let z_i^m be the instrument vector for the min moment applied to observation i. The min moment is:

$$\frac{1}{n} \sum_{i=1}^{n} \left(\min_{j\in s_i(\Lambda)} \left\{\Psi_i\left(\Lambda, j, \alpha\right) + x_{ij}\beta\right\} \right) z_i^m \geq 0 \tag{11}$$

We also provide an analog to Equation 6, the min moment from pairwise stability with deletion degree 2.

$$\frac{1}{n} \sum_{i=1}^{n} \left(\min_{\sigma_{mi2}^- \in \mathcal{S}_{i2}^-} \left\{\Psi_i\left(\Lambda, \sigma_{mi2}^-, \alpha\right) + \left(x_{ij} + x_{il}\right)\beta\right\} \right) z_i^m \geq 0 \quad \{j, l\} = \sigma_{mi2}^-. \tag{12}$$

Not surprisingly, we find that Equations 11 and 12 provide tighter but less precise parameters bounds relative to the corresponding mean moments.

6.2 Empirical issues

There are several more issues to be dealt with before turning to results. We define $\psi_i(\Lambda, \alpha)$ as a scalar function of centrality. For instance, when we focus on betweenness centrality, we have:

$$\Psi(\Lambda, j, \alpha) = \psi_i(\Lambda \cup \{i,j\}, \alpha) - \psi_i(\Lambda - \{i,j\}, \alpha) = \alpha \left(C_i^{betw}(\Lambda \cup \{i,j\}) - C_i^{betw}(\Lambda - \{i,j\}) \right).$$

Inspection of the estimation equations shows that if one vector of parameters satisfies the moments, any multiple will as well. Similarly, setting all parameters to zero will automatically satisfy these moments. That is, the scale of the parameters is not identified. This is standard in discrete choice models. In logit and probit models, we typically address this issue by normalizing the variance of the error term to 1. However, we do not model the distribution of the error term in this paper, so that normalization is not available. Instead, we normalize the coefficient of x_{ij1} to 1. Thus, the remaining parameters should be interpreted as the importance of that variable relative to x_{ij1}. Normalizing the coefficient on x_{ij1} is natural since our assumptions about measurement error exclude it from the instrument vector. That is, we do not estimate a coefficient for the variable that we do not include in the instrument vector.[17]

It is a challenge to implement estimation via moment inequalities in the context of many regressors. In part, this is a computational issue. Because the computational algorithm often involves something like grid search, it can be difficult to implement with many parameters. In addition, the relationship between short and long regression that we are familiar with from linear models does not necessarily hold in inequality estimation. For these reasons, it is important to choose explanatory variables carefully and with an eye towards parsimony.

In thinking of explanatory variables, we are motivated by the literature on gravity equations used to explain trade data, because trade and ASAs are both forms of "links" between countries (see Head and Meyer, 2014). Thus, ideally, we would include all of the explanatory variables that one finds in standard estimation of gravity equations. However, in practice, this entails a great many variables, such as distance, and indicators for shared language, colonial history, and free trade agreements. Much of the gravity literature emphasizes the importance of exporter and importer fixed effects, which greatly increases the number of parameters to be estimated.

Rather than include all of these variables in our moment inequalities algorithm, we instead first estimate a gravity equation model on the list of countries we study. That is, we specify the log of unilateral trade as a linear function of exporter and importer fixed effects, and a series of explanatory variables found to be important in the gravity equation literature. More discussion and details on this estimation appear in Appendix E. We refer to the predicted value from this regression (assuming the shock in the estimation equation is equal to zero) as the *gravity score*. We use the gravity score as an important explanatory variable in our moment inequalities estimation routine. In this way, variables that are known to be important from the gravity literature are used

[17]Some care must be taken in choosing x_{ij1} to make sure that we believe that the true coefficient is positive. If we normalize a negative coefficient to 1, we effectively change the sign of all the other coefficients. In our case, we strongly believe that the gravity score has a positive effect on the likelihood of an agreement.

to predict outcomes in our model of international agreement formation.

Certainly, it is possible that variables such as distance have a different effect on agreement formation than they do on trade. We can address this by including both the gravity score and distance separately as explanatory variables in our agreement formation model. That is, they can be separate values of x_{ij}. We experiment with different specifications and find that these variables have little explanatory value beyond the gravity score. In much of our analysis, we use the gravity score and a constant term as the only explanatory variables, letting the gravity score take the position of x_{ij1}, the variable over which the countries have measurement error and with the coefficient set to 1.

In setting instrumental variables z_{ij}, we include x_{ij2} and $(\psi_i(\Lambda \cup \{i,j\}) - \psi_i(\Lambda - \{i,j\}))$. In our baseline, x_{ij2} is just the constant term. We also use our specification of the gravity score equation to generate extra instruments. In particular, the model implies that predictors of x_{ij1} are good instruments. For instance, suppose the gravity score model relies upon exporter and importer fixed effects, distance between the two countries, and an indicator for sharing a language. The model holds if we assume countries know distance and the language indicator, but do not know the exporter and importer fixed effects, so that the country still has uncertainty about the gravity score. This setup is realistic because countries surely know distance and language, but perhaps do not know future realizations of trade. Because distance and the language indicator are part of the country's information set, we can include the variables in the instrument vector and thus generate extra moments.

We add a constant term to each variable in the instrument vector in order to ensure that each element of z_{ij} is positive because negative instruments can change the sign of the moment that they interact with.

6.3 Asymptotics

Asymptotic properties are challenging in a network context such as ours. We do not provide formal results, although there are a number of related formal treatments. The asymptotic properties in network models depends on whether one imagines observing many networks of a given number of agents, or a single network with the number of agents going towards infinity. In our context with a cross-section of data, it is more natural to think of observing a single network with an increasing number of agents. Menzel (2021, 2024) and Leung and Moon (2024) derive central limit theorems for network formation models in the case of a single network with setups similar to the ones we use. See also Leung (2019).[18]

An important variable is the measure of network structure such as betweenness centrality or eigenvector centrality. One issue is that although this variable differs across observations, it is always computed from the entire network, and we observe only one network. Menzel (2015) takes up the estimation of this type of estimator in a more general strategic framework. A second issue is that as the number of countries goes to infinity, the network structure variables go to zero.

[18] A related issue with asymptotics arises in the estimation gravity equation models, as in Egger and Staub (2016), Charbonneau (2017), Jochmans (2017) and Cameron and Miller (2014).

However, it is important to recognize that eigenvector centrality is normalized by a scalar, and this scalar could be adjusted to asymptote in a way to address this as the number of nodes goes to infinity. Bonacich centrality, which admits eigenvector centrality as a special case, is one way to think about this. Another way to address this issue is to assume that the network variables depend only on countries that are relatively close by in some sense, so the variables stay constant after the number of countries reaches some critical value.[19] An approach along these lines is Menzel (2021). However, to be clear, we compute our network statistics using the entire network.

In our min moments, we take the minimum of a vector of values for each country, and then take a mean across countries. Although the min moment makes use of the same number of observations as the "mean" moments, the minimum statistic may converge to its asymptotic value more slowly than the mean, so the min moment has worse efficiency properties. To be conservative, we treat the min moment as a moment with n observations, as if there was one observation per country. Above, we write Equation 8 with n^a observations but the min moment (Equation 11) with n observations.

7 Results

For our baseline results, we impose five sets of moments: one for linked directed pairs, one for unlinked directed pairs, one for the country-level min moment over linked pairs, one for linked deletion pairs, and one for the country-level min moment over linked deletion pairs. These correspond to Equations 8, 9, 11, 10, and 12. The last two sets of moments capture the refinement concepts in Section 5.2. In our experience, the moments based on strong stability (Equation 7) did not impact our results, and we ignore them going forward.[20] We interact each moment with six instruments: the constant term, the change in the network variable from changing the link, and four variables that are statistically significant from the gravity score regressions: the log of distance, the sum of the log GDP, an indicator for the two countries having a common official language, and an indicator for being in a regional free trade agreement.[21] The number of observations varies across sets of moments, as described in Table 2.

Before presenting results in a table, we graph our moments in order to gain intuition for how estimation works. Graphing is possible because, in our base specification, we estimate only two parameters. We present a color-coded contour map of the Cox and Shi (2023) objective function. In Figure 2, the X axis is the constant term and Y axis is the parameter on betweenness centrality. As noted above, we adopt the normalization that the coefficient on the log gravity score is unity. Each point in the grid is color-coded by its test-inversion rejection confidence level: the lowest confidence level at which that parameter vector would be excluded from the confidence set. Thus,

[19] Dasaratha (2020) studies the asymptotics of eigenvector centrality in a single network using a stochastic block model, although the application is somewhat different.

[20] It is probably not surprising that moments based on deleting multiple links provide more information than moments based on adding multiple links. Deleting multiple links requires only one agent to want to delete, whereas adding requires decisions by multiple agents. Thus, the fact that we do not observe multiple additions is consistent with a wider set of parameters.

[21] The sum of the log GDP is not in the gravity score equation as a result of the presence of exporter and importer fixed effects but seemed natural to include as an instrument.

Table 2: Number of observations in each moment

	Linked directed pairs	Unlinked directed pairs	Countries	Linked deletion pairs
Observations	1,606	1,194	58	4,599
Relevant Equation:	8	9	11,12	10

Note: Columns 1 and 2 count directed country pairs used in estimation. Together with the 506 directed EU-to-EU pairs excluded from estimation, these sum to the total number of directed country pairs in the data ($1,606 + 1,194 + 506 = 3,306 = 58 \times 57$). Equivalently, the data contain 1,653 unordered country pairs.

Figure 2: Graphing the 95% confidence interval

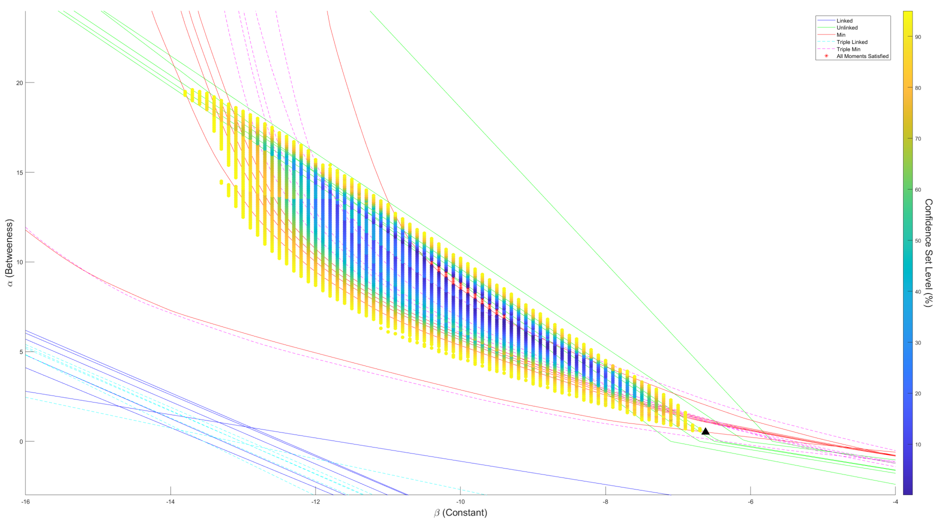

Lines represent different moments from estimation. Points are color-coded by their test-inversion rejection confidence level. The black triangle is the lower bound on the parameter on betweenness centrality in the 95% confidence set.

a value of 90% means that the parameter vector is excluded from the 90% confidence set, while a value of zero means that the parameter vector satisfies all of the sample moment inequalities. If the moments were the population moments rather than sample moments, the set of parameter vectors that satisfy all of the moments would be the identified set. The color coding therefore shows how the confidence set changes as the confidence level varies.

We also graph the moments. That is, we graph the set of parameters that would set each sample moment to zero. Kinks in the lines result from the presence of `min` functions in those moment definitions. There are multiple lines for each moment because we interact the moment with several instruments. The red sliver in the middle of the confidence set is the set of points that satisfy all of the moments. This gives a sense of which inequalities point up and down on the graph. We call the set of parameters that satisfy all of the moments the "sample identified set" for lack of a better term. For limiting the size of the confidence interval, Equation 9 (resulting from the unlinked directed pairs) plays a large role from above and Equations 11 and 12 (the min moments) play a large role from below. Notice that the confidence set extends much farther below the sample identified set than above. That is because there are many fewer observations in the min moments (as described in Table 2) so the min moments are estimated with less precision.

The black triangle represents the lower extremum of the parameter on betweenness centrality. It is just above zero, so we can conclude that the parameter is positive. The straight lines in the bottom left corner result from the linked directed pairs and linked deletion-pair moments. These have no impact on the sample identified set because, consistent with our intuition, the min moments are so much more restrictive. The figure shows us that even accounting for the lack of precision in the min moments, the min moments are still more binding. That is, they are more binding not only on the sample identified set but also on the confidence interval. It is difficult to see from this figure whether the refinements, particularly Equation 12 (the min moment from deletion degree 2) makes a contribution beyond Equation 4, the min moment. In the figure, do the dashed lines add to inference relative to the solid lines? We explore that below.

We now turn to parameter results. Table 3 provides the confidence sets for three specifications. We report only the maximum and minimum of the confidence set for each parameter. The baseline results are presented in the first panel, numbered I. These results correspond to Figure 2. The first row reports the parameter on the network variable betweenness centrality, α in $\alpha\left(C_i^{betw}(\Lambda \cup \{i,j\}) - C_i^{betw}(\Lambda - \{i,j\})\right)$. We can see that the parameter on betweenness centrality is bounded from below by a positive number, 0.51. Thus, we conclude that the parameter is positive and statistically significantly different from zero. That is consistent with countries seeking to become a hub in the agreements network.

In panel II, we replace betweenness centrality with eigenvector centrality. Surprisingly, it is negative, as if countries avoid dense parts of the network. One explanation for this result may be that countries care about betweenness centrality and are indifferent to eigenvector centrality, but eigenvector centrality captures something about betweenness centrality. Indeed, while eigenvector centrality and betweenness centrality are positively correlated (as is clear from Table 1), we find that changes in centrality are negatively correlated. In particular, we find that $\rho(C_i^{betw}(\lambda) - C_i^{betw}(\lambda - \{i,j\}), C_i^{eig}(\lambda) - C_i^{eig}(\lambda - \{i,j\}) = -0.78$ (where $\rho(.,.)$ is the correlation coefficient) when conditioning on linked directed pairs, and $\rho(C_i^{betw}(\lambda) - C_i^{betw}(\lambda \cup \{i,j\}), C_i^{eig}(\lambda) - C_i^{eig}(\lambda \cup \{i,j\}) = -0.84$ for unlinked directed pairs.

To explore this further, we estimate with both betweenness centrality and eigenvector centrality in panel III. We include both network variables as separate instrumental variables. We see that betweenness is still statistically significant and positive, whereas the parameter range on eigenvector centrality includes zero. Thus, the result that eigenvector centrality has a negative effect does not appear robust. Going forward, we focus on specification I.

Table 4 considers several additional specifications. First, we consider adding additional explanatory variables. In panel I, we add coefficients on distance and GDP. Both are insignificant. Recall that these variables also enter through the gravity score. Thus, finding that they are insignificant does not imply they are unimportant. It is consistent with them being important but well controlled for by our specification of the gravity score. In this specification, the parameter on betweenness centrality remains significantly different from zero.

Table 3: Main estimation results

	I		II		III	
	Lo	Hi	Lo	Hi	Lo	Hi
Betweenness	0.51	19.66			0.29	27.51
Eigenvector			-80.87	-9.79	-43.75	8.97
Constant	-13.81	-6.81	-4.05	14.08	-84.00	39.14

Note: The table reports the extremum of the 95% confidence set for each parameter in the specification. Each specification uses multiple moments with different numbers of observations as described in Table 2. Each specification includes gravity score as an explanatory variable with a coefficient set to 1.

Table 4: Robustness results

	I		II		III	
Betweenness	0.41	37.62	0.43	21.63	0.50	22.88
Distance	-1.85	19.44				
GDP	-100	34.24				
Constant	-84.78	0.63	-15.31	-6.52	-14.01	-6.73

Note: Panel I includes additional explanatory variables. Panel II drops the refinements of pairwise stability (Moments 5 and 6). Panel III uses weighted betweenness. The table reports the extremum of the confidence set for each parameter in the specification. Each specification uses multiple moments with different numbers of observations as described in Table 2. We do not search past 100 or -100.

In Panel II, we drop the moments associated with the refinements of pairwise stability, Equations 10 and 12. The parameter interval for betweenness centrality expands. For instance, the lower bound drops from 0.51 to 0.43. This result highlights the value of our approach, which admits refinements of pairwise stability, for narrowing the confidence interval. The role of these moments is subtle. Under betweenness centrality, there are no cases in which a country would want to delete two links but not want to delete either of the links independently. That might lead a researcher to believe that these moments provide no additional information, and indeed, we see in Figure 2 that they are not binding on the sample identified set. But although the refinement moments do not bind on the sample identified set, they still contribute to inference.

In unreported results, we also experiment with dropping other moments in order to explore their contribution to estimation. The results are foreshadowed by Figure 2. Dropping the min moments (Equation 11) leads the confidence set to expand substantially. We do not find bounds for parameters in this case (note that we never search past parameter values of 100/-100). In contrast, dropping the moment for linked pairs (Equation 8) has no impact on our parameters and returns the same results as panel I in Table 3. As we expected, restricting the minimum of a country's links to be positive is substantially more informative than restricting the average to be positive.

Recall that we have normalized eigenvector and betweenness centrality to have the same mean. However, it is evident from Table 1 that eigenvector centrality has a much lower variance. As a robustness check, we normalize eigenvector centrality to have the same mean and variance as betweenness centrality. Note that under such a transformation, the new measure no longer necessarily corresponds to eigenvector centrality, but in any case, the (unreported) results are qualitatively the same, but with estimates scaled as one would expect.

In panel III, we used weighted betweenness rather than betweenness. Results are similar to Table 3. We find that the parameter on betweenness is positive in both specifications and eigenvector centrality is insignificant. Finally, in unreported results, we experiment with our baseline model (Panel I of Table 3) but we use the additive model of link formation. This is briefly described above just after the definition of pairwise stability. This model is consistent with partners being able to make side payments in order to form agreements, so agreements are formed when the sum of two countries' payoffs is positive and not formed when the sum is negative. We fail to find any parameters that satisfy our confidence set restrictions. Thus, we find that the confidence set is empty, which can be understood as a rejection of the side-payment model.

In Appendix F, we repeat our estimation but with our alternative data set based on our expansion of the ICAO data set rather than the WTOs (see Appendix B for more discussion). Results are quite similar, with betweenness centrality being statistically significant and eigenvector centrality being negative when by itself and insignificant when included with betweenness. When including control variables, those variables are insignificant, although betweenness is insignificant in this case, unlike in Table 10. When we compare the parameter interval on betweenness centrality with and without the refinement moments (panel II in Table 11), we find that without the refinements, the parameter interval on betweenness centrality includes zero and thus becomes insignificant. Thus, in the alternative data set, the value of the refinements is particularly striking.

8 Omitted heterogeneity

An important concern in any network formation process is the possibility of omitted variables. We may see many countries linking to a particular country because it is well-connected, or because it is attractive to connect to for some exogenous but unobserved reason. That is, connectivity may be driven by omitted heterogeneity rather than network economics. Naively assuming that there is no omitted heterogeneity may lead to biases in other estimated parameters, for instance, concluding that countries value partnering with highly connected countries when, in fact, those countries are attractive for exogenous but unobservable reasons.[22]

In this section, we address this problem by focusing on different connectivity choices across similarly situated partners, following Pakes et al. (2015). For example in our data, Mexico is connected to Denmark but not Sweden. If we assume that Mexico's unobserved values from connecting to these countries are the same, which country provides a bigger increase in betweenness centrality tells us about how Mexico values this type of connectivity rather than about unobserved heterogeneity. While not a perfect natural experiment, comparing results under this approach provides perspective on the importance of unobserved heterogeneity.

[22]Graham (2017) is another example of a paper discussing unobserved heterogeneity in a network formation setting, although with an approach that is quite different.

8.1 Model

Formally, consider three countries, i, j, and k. Suppose that i is linked to j but not k. That is, $\lambda_{ij} = 1$ and $\lambda_{ik} = 0$. We then have that:

$$\begin{aligned} E\left[\Psi_i(\Lambda, j, \alpha) + x_{ij}\beta + \nu_{ij} \middle| \lambda_{ij} = 1, \Lambda, z_{ij}, x_{ij2}, \nu_{ij}\right] &\geq 0. \\ E\left[\min_{c\in\{i,k\}} \left\{\Psi_c(\Lambda, -c, \alpha) + x_{c,-c}\beta + \nu_{c,-c}\right\} \middle| \lambda_{ik} = 0, \Lambda, z_{ij}, x_{ik2}, x_{ki2}, \nu_{ik}, \nu_{ki}\right] &< 0. \end{aligned} \tag{13}$$

Relative to the Equations 2 and 3, these equations include the term ν_{ij}, which represents omitted heterogeneity; factors observed by the countries but not by the econometrician that lead countries to form agreements.

We cannot use the equations in 13 because of the presence of the unobserved terms. However, we can make progress by assuming $\nu_{ij} = \nu_{ik} = \nu_{ki}$. That is, the unobserved value that i obtains from linking with j and k is the same. Because of the way the *min* function operates in the second line of Equation 13, we must also make the assumption in the second equality, $\nu_{ik} = \nu_{ki}$, that i obtains the same unobservable value from k as what k obtains from i. We relax this below. But with this assumption, the unobserved term drops out of the difference between the two equations in 13:

$$E\left[\Psi_i(\Lambda, j, \alpha) + x_{ij}\beta - \min_{c\in\{i,k\}} \left\{\Psi_c(\Lambda, -c, \alpha) + x_{c,-c}\beta\right\} \middle| \lambda_{ij} = 1, \lambda_{ik} = 0, \Lambda, z_{ij}, z_{ik}, x_{ij2}, x_{ik2}, x_{ki2}\right] \geq 0. \tag{14}$$

This difference is more informative if we use comparison countries that are relatively similar. For instance, in our data, Mexico links to Thailand but not to Sweden. While comparing Mexico's choice to Thailand but not Sweden may be useful, we expect that it will be more useful to constrain ourselves to countries that are more similar, such as Sweden and Denmark. One possibility is to require countries j and k to be on the same continent. In practice, we require both countries to be members of the EU. We choose the EU because it comprises a relatively homogenous set of countries in terms of economic development and governance.

To further tighten the comparison, we take only *nearest neighbors* in terms of our main observable characteristic, the gravity score. Denote x^g_{ij} as the gravity score between i and j. We rank pairs of countries j and k by $|x^g_{ij} - x^g_{ik}|$ and use only the pairs with the lowest 10 values to form observations. For instance, suppose the United States were linked to all but one EU country, Estonia. We would form 10 observations for the US as country i, Estonia as country k, and j being one of the countries with a gravity score with the US among the ten closest to the Estonia-US gravity score. If the US were also not connected to Slovenia, the ten observations would be a mix of observations with k being Estonia or Slovenia based on whichever had the countries closest in terms of gravity score.[23]

For empirical implementation, let $\mathcal{N}^{\text{EU}}$ be the set of EU countries and let $\mathcal{N}^{+}$ be the set of countries that link to some but not all countries in the EU (i.e., there is some variation in their

[23] In our data, the US is not connected to three EU countries: Estonia, Slovenia, and Cyprus.

linking choices):

$$\mathcal{N}^+ = \left\{ i \mid i \in \mathcal{N}, i \notin \mathcal{N}^{\texttt{EU}};\ \exists j,k \in \mathcal{N}^{\texttt{EU}} \texttt{ such that } \lambda_{ij} = 1, \lambda_{ik} = 0 \right\}.$$

Let $\mathcal{H}_i$ be the set of pairs of EU countries where one links and one does not link with i and that are nearest neighbors in terms of gravity score x^g_{ij}. That is:

$$\mathcal{H}_i = \left\{ j,k \mid j,k \in \mathcal{N}^{\texttt{EU}};\ \lambda_{ij} = 1, \lambda_{ik} = 0;\ |x^g_{ij} - x^g_{ik}| \texttt{ among lowest 10} \right\} \texttt{ for } i \in \mathcal{N}^+.$$

Our identifying assumption is:

Assumption 1 $\nu_{ij} = \nu_{ik} = \nu_{ki}$ *for all* $j,k \in \mathcal{H}_i$, $i \in \mathcal{N}^+$.

Let N^+ and N_i be the number of elements in $\mathcal{N}^+$ and $\mathcal{H}_i$. Then, the empirical analog of our moment is:

$$\frac{1}{N^+} \sum_{i \in \mathcal{N}^+} \frac{1}{N_i} \sum_{j,k \in \mathcal{H}_i} \Psi_i\left(\Lambda, j, \alpha\right) + x_{ij}\beta - \min_{c \in \{i,k\}} \left\{ \Psi_c\left(\Lambda, -c, \alpha\right) + x_{c,-c}\beta \right\} \geq 0$$

We also consider an analogy to the min moment in the paper. In this case, that is:

$$\frac{1}{N^+} \sum_{i \in \mathcal{N}^+} \min_{j,k \in \mathcal{H}_i} \left\{ \Psi_i\left(\Lambda, j, \alpha\right) + x_{ij}\beta - \min_{c \in \{i,k\}} \left(\Psi_c\left(\Lambda, -c, \alpha\right) + x_{c,-c}\beta \right) \right\} \geq 0.$$

In practice, we experiment with different numbers of nearest neighbors. That is, we vary the count of comparisons above and below 10. Lower numbers make the comparisons more precise and thus more informative but also reduce the number of observations.

8.2 An alternative comparison

We can take advantage of the special structure of our case to make a less restrictive assumption. In our case, x_{ij} consists of only two variables, the constant term and the gravity score. Both of these are symmetric across countries, so $x_{ij} = x_{ji}$. Thus, $x_{c,-c}$ does not determine which country is the lower value in the `min` function in Equation 14. The lower element of the `min` function is determined entirely by $\Psi_c(\Lambda, -c, \alpha)$. In our linear specification, the term is $\alpha(\psi_c(\Lambda \cup \{c, -c\}) - \psi_c(\Lambda))$. Thus, assuming $\alpha > 0$, which country has a lower value does not depend on α, only on the term $\psi_c(\Lambda \cup \{c, -c\}) - \psi_c(\Lambda))$. We can determine this prior to estimation. In words, we consider only sets of countries i, j, k for which i links with j and not k *and* for which a link between i and k would lead to a smaller change in betweenness centrality for i than k. In this way, it is possible to construct a moment inequality for which we do not have to make assumptions on ν_{ki}.

Formally, define $\mathcal{H}'_i$ to select for elements of $\mathcal{H}_i$ that satisfy the additional condition:

$$\mathcal{H}'_i = \{j,k \mid j,k \in \mathcal{H}_i;\ \psi_i(\Lambda \cup \{i,k\}) - \psi_i(\Lambda) < \psi_k(\Lambda \cup \{i,k\}) - \psi_k(\Lambda)\}.$$

In theory, $\mathcal{N}^+$ would also need to be defined to cover only countries that additionally satisfy

Table 5: Results from differencing approach on the betweenness centrality parameter

Specification	Lower bound	Obs.	Nearest Neighbors
Under Assumption 1:			
1	3.859	3,756	All obs.
2	0.709	510	15
3	0.449	340	10
4	-0.643	170	5
Under Assumption 2:			
5	2.997	1,932	All obs.
6	0.372	240	10

Note: Results from differencing approach described in Section 8. Row 1 uses all EU countries for comparison. Row 2-4 uses nearest neighbors within the EU by gravity score. Rows 5 and 6 uses the pairs selected by Assumption 2.

the new condition, but all countries in $\mathcal{N}^+$ do so in our data so we do not adjust $\mathcal{N}^+$. Now we assume:

Assumption 2 $\nu_{ij} = \nu_{ik}$ *for all* $j, k \in \mathcal{H}_i'$, $i \in \mathcal{N}^+$.

Assumption 2 is more attractive than Assumption 1 because it drops the second equality. The differencing equation no longer requires a `min` function:

$$E\left[\Psi_i(\Lambda, j, \alpha) + x_{ij}\beta - (\Psi_i(\Lambda, k, \alpha) + x_{ik}\beta) \mid i \in \mathcal{N}^+, j, k \in \mathcal{H}_i', x_{ij2}, x_{ik2}, z_{ij}, z_{ik}, \Lambda\right] \geq 0. \quad (15)$$

We can develop estimating equations in an analogous way to the above.

8.3 Results

Exploiting Equations 14 or 15 has two drawbacks. First, the constant term drops out of the differencing so we do not learn about β, only the network parameter α. Second, as in Pakes (2010) and Wollman (2018), the differencing approach identifies only a lower bound for the parameter, not an upper bound. Those papers take other approaches to identify upper bounds, but because the focus of our counterfactual is on the lower bound, we do not pursue that here. Because we learn only about the lower bound for one parameter, we treat the differencing approach as a robustness exercise rather than our main result.

We consider only betweenness centrality as our network statistic. The way unobserved heterogeneity affects betweenness centrality is complex so it is difficult to say whether it should bias α up or down. Results appear in Table 5. The first four rows of results are based on Assumption 1. There are 34 non-EU countries that do not connect to all EU countries. For each of these 34 countries, we find each pair of EU countries that link and do not link with the non-EU country.[24] That

[24] Russia is the only non-EU country with agreements with every EU country.

leads to 3,756 observations. For this case, we find a lower bound for the confidence interval of α of 3.856, much higher than the lower bound in our main result of 0.43. We then impose the nearest neighbor criterion by using only the EU countries that are close to each other in gravity score with the non-EU country. For row 3, we use 10 nearest neighbors for each of the 34 non-EU countries, and thus have 340 observations. In this case, the lower bound is 0.45, very close to our main result. We take this as evidence that our main result is not driven by unobserved heterogeneity.

As we increase the number of nearest neighbors, the lower bound increases. For instance, row 2 considers 15 nearest neighbors, and we find a lower bound of 0.71. Our result that the lower bound is positive is not completely robust. For instance, row 4 shows that when we use only 5 nearest neighbors, the lower bound is -0.64. But this estimation relies on only 170 observations, so perhaps it is not surprising that we find a statistically insignificant result. We then switch to the approach in Subsection 8.2, i.e., the data set for which the non-EU country provides the minimum value in the `min` function with the EU country that does not link. Results are similar, although there are fewer observations that satisfy this criterion. Overall, we conclude that our main result is robust to our method for addressing unobserved heterogeneity.

9 Efficiency of the ASA Network

The results above imply that the structure of the network of agreements influences payoffs over and above bilateral factors alone. The existence of such externalities implies that both the observed number of agreements and the composition of agreements may not be consistent with an efficient outcome. To evaluate this issue, we must first define a measure of social welfare.

We define social welfare globally as the sum of welfare across all countries.

$$S(\Lambda) = \sum_{i=1}^{n} \Pi_i(\Lambda) = \sum_{i=1}^{n} \psi_i\left(\Lambda, \alpha\right) + \sum_{j \in s_i(\Lambda)} x_{ij}\beta + \varepsilon_{ij} \tag{16}$$

This welfare function weighs each country equally, which is a natural assumption given that we do not have a way of turning utility into dollars. Because welfare is defined as the sum of country-level payoffs, the bilateral component of a link appears once for each endpoint: for a link between i and j, social welfare includes country i's payoff from linking to j and country j's payoff from linking to i. We assume all EU pairs of countries are always linked, but otherwise count them in both network statistic computations and social welfare. Thus, we consider 1,400 potential links.

With this welfare function, we can determine what set of links maximize total joint surplus. The globally optimal network differs from the observed network because of two externalities. The first arises within bilateral link formation: countries do not account for their partner's benefit from forming a link. Some agreements are beneficial to two countries jointly but nevertheless do not get signed, which leads to too few agreements. The second externality is that countries do not account for the effect on others through network structure. Based on the results above, countries form agreements that increase their betweenness centrality, and thus make them hubs. This force can lead the network to be either over- or under-connected. Countries form links without taking into account whether doing so contributes to or takes away from another country's status as a hub.

Finding the global maximum of a function of 1,400 binary variables can be very challenging.[25] We proceed by ordering the country pairs and sequentially turning on or off links if doing so increases welfare. We iterate through the ordering until we reach convergence. Note that if $\alpha = 0$, there would be no externality across pairs of countries, and a single iteration would reach the global maximum. Thus, this approach should be more reliable for low values of α. In addition, we repeat the exercise starting from many different network structures and with many different orderings and we always converge to the same outcome.

In order to isolate the two externalities, we perform the exercise again but holding betweenness centrality constant at the level observed in the data. That is, we adjust links to maximize welfare, but as we adjust links, we do not adjust the betweenness centrality values that go into each country's payoff. Thus, we focus only on the bilateral payoff. Note that in this case, sequentially iterating through country pairs finds the global optimum.

In selecting the parameter vector to be used to evaluate equation (16), we are guided by a desire to provide a lower bound on the role of the network externality. Consequently, we set α equal to the lower bound of the confidence interval, the value reported in column 1, Panel I of Table 3. We use the parameter vector from this specification. That is, when we searched for the lowest value of α in the confidence set, we found $\alpha = 0.51$, as reported in Table 3, and a coefficient on the constant term of $\beta = -6.5$. We use this value for β in our calculations.

Results appear in Table 6. If we consider only the bilateral effect of adding links, Column 2 shows that we would add 302 links to the observed number of 821 links but drop 75 other links, for a net change of increasing by 227 links. Thus, the number of links in the network increases from 58.6% to 74.9% of potential links. The global optimum adds 12 fewer links and cuts 10 more for a net increase from the observed network of 205, 22 fewer links than considering only the bilateral effect. Thus, the social optimum calls for 73.3% of links to be formed, about 1.6 percentage points less than without the betweenness centrality effect at its lower bound.

Examining compositional change gives a better sense of the results. In comparing the observed network to the global optimum, 46 countries increase their number of links while 11 countries decrease their number of links. Country-level changes are reported in Table 7. The big increases occur for large countries with relatively few links – Brazil, Indonesia, Mexico, and Saudi Arabia all go from the 27-31 range of observed links to the 46-48 range, with Brazil being the biggest increaser going to 51 links. In contrast, small countries with relatively many links have the biggest decreases: Bahrain, Oman, and Trinidad drop by 10-16 links in the optimal network. The biggest increase in Europe is Ireland going from 10 to 28 links and the biggest decrease is Cyprus going from 12 to 6. Even countries with small changes in the number of links may have large changes in who they link with. For instance, Egypt has no change in its number of links between the observed and socially optimal outcome, but the optimal outcome calls for Egypt to drop 7 links

[25]One approach is the *branch and bound* approach utilized in Seim and Waldfogel (2013). However, it does not guarantee a global maximum and our understanding is that existing implementations do not handle the number of variables in our problem. Leung (2020) proposes an attractive approach but it does not apply when payoffs depend on choices anywhere in the network, such as eigenvector and betweenness centrality.

Table 6: Counterfactual

	Observed	Bilateral Effect Only	Global Optimum
Added		302	290
Subtracted		75	85
Net Increase		227	205
Links	821	1048	1026
Share of Total	58.6%	74.9%	73.3%

Note: Column 1 shows the number of links in the observed data of 1,400 potential links. We do not consider EU-to-EU links. Column 2 solves for the global optimum holding betweenness centrality fixed. Column 3 solves for the global optimum. Column 3 reports added and subtracted links relative to Column 1, not Column 2.

and add 7 others, leading to 14 changes. Primarily, Egypt drops links with distant, relatively small countries such as Ireland and Indonesia, and adds links to nearby countries, such as Algeria and Cyprus, including relatively small ones such as Bahrain and Kuwait.

10 Conclusion

This paper studies the formation of air services agreements as a problem of strategic network formation. We develop a structural model in which countries choose bilateral links, where the payoff from any link depends not only on bilateral characteristics but also on the position of the countries in the wider network. Pairwise stability provides a tractable way to discipline the observed network without requiring us to select among multiple stable networks. The resulting moment inequalities allow us to estimate the importance of network structure and to incorporate refinements of pairwise stability.

The empirical results indicate that network position matters. In particular, countries value agreements that increase their role as hubs in the global ASA network. This effect is robust across several specifications and is reinforced by moments based on stability refinements. We also show that the main result is not simply an artifact of unobserved country attractiveness by using comparisons between linked and unlinked partners that are similar on observable gravity characteristics.

The counterfactual analysis suggests that the observed network is far from globally efficient. A social planner would increase the overall number of agreements, but would also substantially reallocate links across countries. This reflects two distinct externalities: some jointly beneficial agreements fail to form because payoffs are non-transferable, and individual countries do not account for how their link choices affect the network position of others. The resulting efficient network is therefore not merely a denser version of the observed network, but a meaningfully different allocation of international aviation links.

More broadly, the paper shows that air services agreements are not only bilateral policy instruments but also components of a global network. This has implications for how liberalization should be assessed. Evaluating agreements pair by pair misses the externalities generated by network structure. A more coordinated approach to international aviation policy may therefore

Table 7: Counterfactual links by country

	Observed	Social planner	Increase	Change
Bahrain	31	15	-16	22
Oman	29	17	-12	16
Qatar	35	24	-11	17
Trinidad	12	2	-10	10
Cyprus	12	6	-6	16
Kuwait	32	27	-5	17
Switzerland	34	31	-3	5
France	35	34	-1	1
Germany	35	34	-1	1
Iceland	4	3	-1	3
Austria	29	28	-1	9
Egypt	40	40	0	14
United Kingdom	34	35	1	1
Netherlands	31	34	3	5
Estonia	2	6	4	4
United States	53	57	4	4
Belgium	28	32	4	8
Denmark	24	28	4	8
Singapore	46	50	4	8
Thailand	46	51	5	11
Greece	19	24	5	13
Algeria	15	20	5	19
Russia	48	54	6	10
Pakistan	37	43	6	16
Philippines	35	41	6	18
Italy	27	34	7	7
Spain	27	34	7	7
Czech Republic	19	26	7	11
Hungary	15	23	8	10
India	44	52	8	10
Sweden	23	31	8	10
United Arab Em	40	48	8	12
Norway	18	27	9	9
Canada	42	51	9	11
Slovenia	4	13	9	11
South Africa	39	48	9	13
Turkey	42	51	9	15
China	46	56	10	10
Malaysia	39	49	10	16
Israel	31	42	11	13
Finland	15	27	12	14
Iran	31	43	12	18
Hong Kong	38	51	13	15
Poland	13	26	13	15
Venezuela	14	27	13	17
Japan	42	56	14	14
New Zealand	23	38	15	15
Colombia	13	29	16	18
Portugal	11	27	16	20
Argentina	26	43	17	17
Australia	31	48	17	19
Saudi Arabia	31	48	17	21
South Korea	37	55	18	18
Mexico	28	46	18	18
Ireland	10	28	18	20
Chile	22	41	19	21
Indonesia	28	47	19	25
Brazil	27	51	24	24

Note: Column 1 shows the observed number of links by country, dropping EU-to-EU links. Column 2 shows the amount the social planner would choose. Column 3 is the difference between Columns 2 and 1. The rows are sorted on Column 3. Column 4 is the number of changes the social planner makes to that country, so if the planner adds one link and subtracts another, the number of changes is 2.

improve welfare, especially when the value of an agreement depends on how it reshapes access, connectivity, and hub positions across the system.

References

Badev, Anton (2021) "Nash Equilibria on (Un)Stable Networks," *Econometrica*, 89, 1179–1206.

Baier, Scott L. and Jeffrey H. Bergstrand (2004) "Economic Determinants of Free Trade Agreements," *Journal of International Economics*, 64, 29–63.

Bontemps, Christian, Cristina Gualdani, and Kevin Remmy (2023) "Price Competition and Endogenous Product Choice in Networks: Evidence from the US Airline Industry," Toulouse School of Economics Working Paper, n. 23-1415.

Borgatti, S.P., M.G. Everett, and L.C Freeman (2002) *UCINET for Windows: Software for Social Network Analysis*, Harvard, MA: Analytic Technologies.

Boucher, Vincent and Ismael Mourifié (2017) "My Friend Far Far Away: A Random Field Approach to Exponential Random Graph Models," *Econometrics Journal*, 20, S14–S46.

Brandes, Ulrik and Thomas Erlebach eds. (2005) *Network Analysis: Methodological Foundations*, Lecture Notes in Computer Science, Berlin, Heidelberg: Springer.

Cameron, A. Colin and Douglas L. Miller (2014) "Robust Inference for Dyadic Data," Unpublished manuscript, University of California - Davis.

Charbonneau, Karyne B. (2017) "Multiple Fixed Effects in Binary Response Panel Data Models," *The Econometrics Journal*, 20, S1–S13.

Chen, Maggie Xiaoyang and Sumit Joshi (2010) "Third-country effects on the formation of free trade agreements," *Journal of International Economics*, 82 (2), 238–248.

Christakis, Nicholas, James Fowler, Guido W. Imbens, and Karthik Kalyanaraman (2020) "An empirical model for strategic network formation," in Graham, Bryan and Áureo de Paula eds. *The Econometric Analysis of Network Data*, Chap. 6, 123–148: Academic Press, https://doi.org/10.1016/B978-0-12-811771-2.00012-2.

Conte, M., P. Cotterlaz, and T. Mayer (2022) "The CEPII Gravity database," CEPII Working Paper N°2022-05, July.

Cox, Gregory and Xiaoxia Shi (2023) "Simple Adaptive Size-Exact Testing for Full-Vector and Subvector Inference in Moment Inequality Models," *Review of Economic Studies*, 90, 201–228.

Cristea, Anca, David Hummels, and Brian Roberson (2017) "Estimating the Gains from Liberalizing Services Trade: The Case of Passenger Aviation," Unpublished Manuscript, Univeristy of Oregon.

Currarini, Sergio, Matthew O. Jackson, and Paolo Pin (2009) "An Economic Model of Friendship: Homophily, Minorities, and Segregation," *Econometrica*, 77, 1003–1045.

Dasaratha, Krishna (2020) "Distributions of Centrality on Networks," *Games and Economic Behavior*, 122, 1–27.

Dickstein, Michael J. and Eduardo Morales (2018) "What do Exporters Know?" *Quarterly Journal of Economics*, 133, 1753–1801.

Dutta, P and S. Mutuswami (1997) "Stable Networks," *Journal of Economic Theory*, 76, 322–344.

Egger, P., M. Larch, K Staub, and R. Winkelmann (2011) "The Trade Effects of Endogenous Preferential Trade Agreements," *American Economic Journal: Economic Policy*, 3, 113–143.

Egger, Peter H. and Kevin E. Staub (2016) "GLM estimation of trade gravity models with fixed effects," *Empirical Economics*, 50, 137–175.

European Commission (2010) "Common Aviation Area with the Neighbouring Countries by 2010 - Progress Report," Communication from the Commission of the European Commission COM/2008/0596.

Goldsmith-Pinkham, P and Guido Imbens (2013) "Social Networks and the Identification of Peer Effects," *Journal of Business & Economic Statistics*, 31, 253–264.

Goyal, Sanjeev (2007) *Connections: An Introduction to the Economics of Networks*: Princeton University Press.

Graham, Brian S. (2015) "Methods of identification in social networks," *Annual Review of Economics*, 7, 465 – 485.

Graham, Bryan and Áureo de Paula eds. (2020) *The Econometric Analysis of Network Data*: Academic Press.

Graham, Bryan S. (2017) "An Econometric Model of Network Formation with Degree Heterogeneity," *Econometrica*, 85, 1033–1063.

Gualdani, Cristina (2021) "An Econometric Model of Network Formation with an Application to Board Interlocks between Firms," *Journal of Econometrics*, 224, 345–370.

Havel, Brian F. (2009) *Beyond Open Skies: A New Regime for Inernational Aviation*: Kluwer Law International.

Head, K., T. Mayer, and J. Ries (2010) "The Erosion of Colonial Trade Linkages After Independence," *Journal of International Economics*, 81, 1–14.

Head, Keith and Thierry Meyer (2014) "Gravity Equations: Workhorse, Toolkit and Cookbook," in Gopinath, Gita, Elhanan Helpman, and Kenneth Rogoff eds. *Handbook of International Economics*, 4, 131–195: Elsevier.

Ho, Kate (2009) "Insurer-Provider Networks in the Medical Care Market," *American Economic Review*, 99, 393–430.

Holmes, Thomas J. (2011) "The Diffusion of Wal-Mart and Economies of Density," *Econometrica*, 79, 253–302.

Houde, Jean-Fançois, Peter Newberry, and Katja Seim (2023) "Nexus Tax Laws and Economies of Density in E-Commerce: A Study of Amazon's Fulfillment Center Network," *Econometrica*, 91, 147–190.

Hsieh, Chih-Sheng and Lung-Fei Lee (2016) "A social interactions model with endogenous friendship formation and selectivity," *Journal of Applied Econonometrics*, 31, 301–319.

Jackson, Matthew O. (2008) *Social and Economic Networks*: Princeton University Press.

Jackson, Matthew O. and Anne van den Nouweland (2005) "Strongly Stable Networks," *Games and Economic Behavior*, 51, 420–444.

Jackson, Matthew O. and Asher Wolinsky (1996) "A Strategic Model of Social and Economic Networks.," *Journal of Economic Theory*, 71 (1), 44 – 74, `http://search.ebscohost.com.ezproxy.bu.edu/login.aspx?direct=true&db=ecn&AN=0400572&site=ehost-live&scope=site`.

Jeon, Jihye and Marc Rysman (2025) "Investment and Usage of the Subsea Internet CableNetwork," Unpublished manuscript.

Jochmans, Koen (2017) "Two-way models for gravity," *Review of Economics and Statistics*, 99, 478–485.

Kolaczyk, Eric (2009) *Statistical Analysis of Network Data*, Springer Series in Statistics: Springer.

Leung, Michael (2015) "Two-Step Estimation of Network-Formation Models with Incomplete Information," *Journal of Econometrics*, 188, 182–195.

Leung, Michael P. (2019) "A Weak Law for Moments of Pairwise Stable Networks," *Journal of Econometrics*, 210, 310–326.

——— (2020) "Equilibrium Computation in Discrete Network Games," *Quantitative Economics*, 11, 1325–1347.

Leung, Michael P. and Hyungsik Roger Moon (2024) "Normal Approximation in Large Network Models," arXiv:1904.11060v6 [econ.EM].

Limão, Nuno (2016) "Preferential trade agreements," *Handbook of Commercial Policy*, 1, 279–367.

Mele, Angelo (2017) "A Structural Model of Dense Network Formation," *Econometrica*, 85, 825–850.

Menzel, Konrad (2015) "INFERENCE FOR GAMES WITH MANY PLAYERS," *Review of Economic Studies*, 83, 306–337.

——— (2021) "Central Limit Theory for Models of Strategic Network Formation," Unpublished manuscript, New York University.

——— (2024) "Strategic Network Formation with Many Agents," Unpublished manuscript, New York University.

Micco, Alejandro and Tomas Serebrisky (2006) "Competition Regimes and Air Transport Costs: The Effects of Open Skies Agreements.," *Journal of International Economics*, 70, 25 – 51, `http://search.ebscohost.com.ezproxy.bu.edu/login.aspx?direct=true&db=ecn&AN=0871138&site=ehost-live&scope=site`.

Miyauchi, Yuhei (2016) "Structural Estimation of a Pairwise Stable Network with Nonnegative Externality," *Journal of Econometrics*, 195, 224–235.

Myerson, Roger (1977) "Graphs and Cooperation in Games," *Mathematics of Operations Research*, 2, 225–229.

Odoni, Amedeo (2009) "The International Institutional and Regulatory Environment," in Belobaba, Peter, Amedeo Odoni, and Cynthia Barnhart eds. *The Global Airline Industry*: Wiley.

Pakes, Ariel (2010) "Alternative Models for Moment Inequalities," *Econometrica*, 78, 1783– 1822.

Pakes, Ariel, Jack Porter, Katherine Ho, and Joy Ishii (2015) "Moment Inequalities and Their Application," *Econometrica*, 83, 315–334.

de Paula, Áureo (2020) "Econometric Models of Network Formation," *Annual Review of Economics*, 2020, 775–799.

Prell, Christina (2011) *Social Network Analysis: History, Theory and Methodology*: SAGE Publications.

Santos Silva, J.M.C. and Silvana Tenreyro (2006) "The Log of Gravity," *Review of Economics and Statistics*, 88, 641–658.

Seim, Katja and Joel Waldfogel (2013) "Public Monopoly and Economic Efficiency: Evidence from the Pennsylvania Liquor Control Board's Entry Decisions," *American Economic Review*, 103.

Sheng, Shuyang (2020) "A Structural Econometric Analysis of Network Formation Games Through Subnetworks," *Econometrica*, 88, 1829–1858.

von den Steinen, Erwin and Claude Probst (2013) "An Overview of the Air Services Agreements Concluded by the EU," European Parliament Directorate-General for Internal Policies Policy Department B: Structural and Cohesion Policies: Transport and Tourism PE 495.849.

Wollman, Thomas G. (2018) "Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles," *American Economic Review*, 108, 1364–1406.

World Trade Organization (2006) "Second Review of the Air Transport Annex: Developments in the Air Transport Sector (Part III), Quantitative Air Services Agreements Review (QUASAR) Volumes I and II: Note by the Secretariat," Document S/C/W/270/Add.1.

Yuan, Zhe and Panle Barwick (2024) "Network Competition in the Airline Industry: A Framework for Empirical Policy Analysis," *RAND Journal of Economics*, In Press.

Appendices

Appendix A: An example of estimation with multiple equilibria

In the literature on the formation of free trade agreements and other types of international agreements, we may wish to know the probability of an agreement conditional on the formation of other agreements. In this section, we provide a simple example of why naive estimation of conditional choice problems in games is problematic, which explains why we use an approach based on moment inequalities. Rather than write the model in terms of the formation of international agreements, we write in terms of an entry game, which is the canonical discrete game in industrial organization. The notation and model are not particularly tied to our paper, but we have found this point confusing to some readers, so we hope this section is helpful.

Suppose two firms 1 and 2 draw entry costs $c_i \sim U[0,1]$, $i = 1,2$. Entry costs are observable to both firms. The firms simultaneously choose whether to enter or not, denoted as $y_i \in \{0,1\}$. Monopoly profit is $3/4 - c_i$ and duopoly profit is $1/4 - c_i$. Not entering earns 0.

In this case, firm 1's best response function is:

$$y_1 = \mathbb{1}\{c_i \leq 3/4 - y_2/2\}. \tag{17}$$

Firm 2's best response is analogous. In this setting, there are multiple equilibria if both firms draw cost shocks $c_i \in [1/4, 3/4]$. However, one might think that we can ignore the presence of multiple equilibria and estimate best response functions. That is, we can estimate the probability of one firm entering conditional on the other's choice because we can see that the presence of multiple equilibria does not affect Equation 17. However, that is not correct.

To see this, assume that if there are multiple equilibria, firm 1 enters with probability γ. In this case, the probability of each of the four outcomes is in Table 8.

Table 8: Probability of entry outcomes

y_1, y_2	probability
$0,0$	$1/16$
$1,0$	$5/16 + \gamma/4$
$0,1$	$5/16 + (1-\gamma)/4$
$1,1$	$1/16$

In this case, the probability of observing firm 1 enter conditional on 2 entering is:

$$P(y_1 = 1 | y_2 = 1) = \frac{P(y_1 = 1, y_2 = 1)}{P(y_2 = 1)} = \frac{1/16}{1/16 + 5/16 + (1-\gamma)1/4}.$$

Importantly, γ shows up in this equation even though it did not appear in Equation 17. That is, even though in our underlying model, the best response function that governs whether firm 1 enters does not depend on the presence or outcome of multiple equilibria, the empirical relationship we observe in our data still does. Put differently, the best response function follows a uniform distribution, but the uniform distribution is not the correct distribution once we consider data generated from this best response function. If we were to specify a best response function based on an alternative distributional form, such as logit or probit, we would face the same problem: The data do not follow the distribution in our model in the presence of multiple equilibria.

Even if we were to assume that we knew γ or set $\gamma = 1$ or $\gamma = 0$, our problems would not be over. That is because the model makes clear that c_i is correlated with y_{-i}. When one firm draws a low cost, it is likely to enter, and thus its rival is not likely to enter. This creates an endogeneity problem, which is present whether there are multiple equilibria or not.

For these reasons, we utilize a method with minimal reliance on the functional form or exogeneity of the unobserved term, namely moment inequalities.

Appendix B: Data construction

In this section, we discuss data construction in more detail, as well as some variables that appear in our data that we do not use.

As discussed in Section 3, we obtained the WASA database available from ICAO that covers air services agreements. Although agreements are public documents, they are surprisingly difficult to observe for the purpose of creating a database. ICAO relies on countries to self-report any new agreements, but has recently engaged in a more proactive approach to learn about agreements. Even so, ICAO representatives believe that the WASA database contains only a subset of existing agreements. In private communication, an ICAO representative guessed that WASA has between one-half and two-thirds of agreements, although it is difficult to know. Some missing agreements in the 2005 data are quite prominent, such as the agreement between the US and France. Also, agreements are kept in the database until the participating countries indicate that they have a new agreement. Thus, WASA contains a number of agreements between countries that no longer exist. Furthermore, the database contains a number of agreements between EU countries, which we know are superseded by EU legislation.

As discussed in Section 3, the WTO expanded the 2005 WASA database to cover missing agreements and clean out agreements that did not belong. The WTO uncovered "orphan" agreements (i.e., ones not in the WASA database) if there is direct passenger service between two countries. For this purpose, the WTO utilized estimates of global scheduled traffic based on Billing Settlement Plans for 2005 provided by the International Air Transport Association (IATA). WTO representatives argue that this approach is informative because the existence of commercial flights requires that there is some kind of agreement between the two countries in advance. Using traffic to infer the existence of agreements requires some judgment. For instance, observing a very low volume of traffic may be the result of charter flights, which do not require an ASA. But in general, the WTO assumes agreements exist when direct passenger service is observed. We extracted this information from Tables 2 and 3 of World Trade Organization (2006). To our knowledge, the WTO data project is only for 2005.

Before finding the WTO project, we also augmented the ICAO data set to fill in missing agreements and we independently arrived at a method that is very similar to that of the WTO. We used a database on international flights, the TFS database (also available from ICAO), to infer when agreements existed. Some examples of differences between our approach and the WTO's are that we used the average of the number of flights over five years whereas the WTO used a single year and, to address charter flights, we used a cutoff in flight volume to infer that an agreement existed. Our resulting set of inferred agreements is close to what the WTO found. The WTO finds 18 more links than our method out of a potential 1400 links, an increase of 0.13%. We take the WTO data set as our main data set and use our data set as a robustness check. We prefer the WTO database because of its public nature. In addition to the benefits of drawing data from a public entity rather than private researcher work, we are allowed to post the WTO version of the data, whereas our purchasing agreement with ICAO restricts that.

We model the binary choice of countries to form an agreement or not rather than the content of agreements, such as whether they are Open Sky Agreements. Open Sky Agreements are a subset of air services agreements. Open Sky Agreements are a particularly liberal version of air services agreements. While the meaning of "Open Sky" has evolved over time, it has come to mean agreements that have relatively few restrictions on routes, cities, airlines, and prices. We observe an indicator of whether an agreement is considered Open Sky in the part of the data from ICAO, although we do not make use of it in this paper. While Open Sky Agreements are particularly important between the biggest countries, they appear to be less important outside of that group, which is the source of most of the variation in our data. Furthermore, we observe only the existence of an agreement for the agreements added by the WTO, not the content of the agreement, such as the Open Sky indicator.

The WASA database also lists the date that the agreement was signed and any amendments, but we found this difficult to use. If countries signed an agreement and then signed a new agreement, we see only the most recent agreement, and we have no indication of the existence of the earlier agreement. Whereas some countries that have liberalized air services have amended the agreement they originally signed in the 1940's, others have started over with a new agreement. For example, our data would indicate that the US and Canada have had an agreement only since 1992, but it is well known that there was air travel between these two countries long before then. Furthermore, we do not observe the date for agreements added in the WTO process. Thus, we ignore the time dimension of our data set, and treat our data as a cross-section of existing agreements.

Appendix C: Calculation of Weighted Network Statistics

The standard calculation of betweenness centrality is:

$$C_i^{betw}(\Lambda) = \sum_{\{j,k:j\neq i,k\neq i,j\neq k\}} \frac{P_i(j,k)}{P(j,k)}.$$

This is the same equation as written in Section 4. Let $\mathcal{P}_{ijk}$ be the set of geodesics between j and k that pass through i. We denote the elements of this set as ρ_{pijk}, $p = 1, \ldots, \bar{P}_{ijk}$ where $\bar{P}_{ijk}$ is the number of elements in this set. As an example, in Figure 1, $\mathcal{P}_{CBE} = \{BCDE\}$ and $\bar{P}_{CBE} = 1$. That is, there is one geodesic passing from B to E through C and it is $\{BCDE\}$. We denote w_{ijkp} as the weight placed on element p of $\mathcal{P}_{ijk}$. For instance, w_{pijk} may denote the amount of international trade recorded between this sequence of countries or the inverse of the distance of this sequence in kilometers.

We let the total betweenness weight on node i from j and k be:

$$\bar{w}_{ijk} = \sum_{p=1}^{\bar{P}_{ijk}} w_{ijkp}.$$

We write weighted betweenness for node i as:

$$C_i^{wtbetw}(\Lambda) = \sum_{\{j,k:j\neq i,k\neq i,j\neq k\}} \frac{\bar{w}_{ijk}}{\sum_{z:z\neq j,k} \bar{w}_{zjk}}.$$

Note that under our measure, only links with the shortest unweighted geodesics contribute to a node's betweenness centrality. For instance, suppose for node i, there are multiple unweighted geodesics between j and k of length 2. All of those contribute to $C_i^{wtbetw}(\Lambda)$, and no paths of longer length do so. We might imagine that other paths should contribute as well, and i may be between j and k on some paths that have longer unweighted length but shorter weighted length. For instance, a path of three hops may have a shorter distance in kilometers than some paths with only two hops. This issue does not seem important in our context given the connectedness of our network but could be an issue in some other settings.

Weighted measures of eigenvector centrality have already been developed. Essentially, we multiply each element of Λ by the weight between each pair of nodes (such as trade flows or the inverse of physical distance). We then compute eigenvector centrality as before.

Appendix D: Computational Details of Estimation Algorithm

We follow the full vector test (Section 3.1) of Cox and Shi (2023). Our model generates J moments $m_j(\boldsymbol{\theta})$, $j = 1, \ldots, J$ where $\boldsymbol{\theta} = \{\alpha, \beta\}$. For instance, based on Equation 8, we have a moment j defined as:

$$m_j(\boldsymbol{\theta}) = \frac{1}{n^a} \sum_{(i,j)\in I^a} \left(\Psi_i(\Lambda, j, \alpha) + x_{ij}\beta\right) z_{ij}.$$

We write all moments so they are expected to be positive, so that the moment corresponding to agreements that are not formed has a negative sign in front of it:

$$m_j(\boldsymbol{\theta}) = -\frac{1}{n^{na}} \sum_{(i,j)\in I^{na}} \min_{l\in\{i,j\}} \left\{\Psi_l(\Lambda, -l, \alpha) + x_{l,-l}\beta\right\} z_{ij}.$$

For a given parameter vector $\boldsymbol{\theta}$, we compute the $J \times 1$ vector of moments $\boldsymbol{m}(\boldsymbol{\theta})$ and the covariance matrix of the moments $\boldsymbol{\Sigma}(\boldsymbol{\theta}) = \texttt{Var}(\sqrt{n}\boldsymbol{m}(\boldsymbol{\theta}))$. Let $\boldsymbol{\mu}$ be a $J \times 1$ vector. Our test statistic is:

$$T(\boldsymbol{\theta}) = \min_{\boldsymbol{\mu}:\boldsymbol{\mu}\geq 0} \left(\boldsymbol{m}(\boldsymbol{\theta}) - \boldsymbol{\mu}\right)' \boldsymbol{\Sigma}(\boldsymbol{\theta})^{-1} \left(\boldsymbol{m}(\boldsymbol{\theta}) - \boldsymbol{\mu}\right). \tag{18}$$

Solving for $T(\boldsymbol{\theta})$ requires solving a non-linear optimization problem for each parameter vector under consideration. Let r be the number of elements of $\boldsymbol{\mu}$ that are equal to zero, which is the number of active moments at θ. The critical value is $\chi^2_{r,1-\alpha}$, the chi-squared distribution with r degrees of freedom and a confidence level of α (apologies for overusing the symbol α!). We do not implement the refinement of the critical value described in their Supplemental Appendix A.1, which makes our implementation conservative. We reject the parameter from the confidence interval if $T(\boldsymbol{\theta}) > \chi^2_{r,1-\alpha}$. We also record $\hat{\alpha}$ such that $T(\boldsymbol{\theta}) > \chi^2_{r,1-\hat{\alpha}}$, which we use color code the plotted points in Figure 2.

Let θ_i be the i'th element of $\boldsymbol{\theta}$. In order to find the upper bound of θ_i, we solve:

$$\max_{\boldsymbol{\theta}} \theta_i \texttt{ such that } T(\boldsymbol{\theta}) \leq \chi^2_{r,1-\alpha}. \tag{19}$$

The confidence interval is not guaranteed to be convex. We found that iterating between simplex search and compass search, supplemented by graphical inspection such as in Figure 2, worked well for finding the global extremum. Switching from max to min found us the lower bound of θ_i.

We typically started with simplex or gradient search to minimize $T(\boldsymbol{\theta})$ to find $\boldsymbol{\theta}$ that satisfied all moments. We used these as starting values in Equation 19 to find extremum points in the confidence set.

Appendix E: Gravity Equation Estimation

This section details our approach to gravity estimation. We take our dependent variable to be the log of directional trade for each pair of countries, so each pair enters the data set twice, once with one country as the exporter and once as the importer. We use the average of the log of directional trade, averaged over five years 2001-2005. The literature on gravity equations suggests a wide set of explanatory variables that might be included. We focus on five: the log of distance, an indicator

Table 9: Results from gravity equation estimation

	Parameter	S.E.
ln(Distance)	-1.002	(0.045)
Shares a border	0.178	(0.154)
Shares an official language	0.560	(0.998)
Colonial relationship	0.518	(0.168)
Shares a regional trade agreement	0.714	(0.097)
Israel-Muslim majority country	-12.26	(0.423)
Observations	3,306	

Notes: Dependent variable is log of directional trade flow (+0.01). Regression includes exporter and importer fixed effects.

variable for sharing a border, an indicator for sharing a common official language, an indicator for sharing a colonial history, and an indicator for being in a regional trade agreement. In addition, following the literature, we include a full set of exporter and importer indicator variables.

In practice, we increment trade flows by 0.01 in order to address taking log of zero. Some previous papers address zeros in the dependent variable of a gravity equation in different ways, such as Santos Silva and Tenreyro (2006).[26] Because we analyze relatively large countries, the number of pairs that show zero trade is very small. More than half of them involve Israel and a Muslim-majority country, which is an issue outside of the framework of any gravity model. As such, we include an indicator variable for cases in which Israel trades with a Muslim-majority country.

Results appear in Table 9. All point in the expected direction, although several are statistically insignificant. In order to construct our gravity score, we take the predicted value of log exports and log imports between each pair and sum them. That is, we sum the logged variables, and label the result the *gravity score*.

An alternative to our approach is Santos Silva and Tenreyro (2006), who specify the gravity equation in levels rather than logs. Because we use a logged variable as the gravity score, their approach is less attractive to us. However, as a robustness check, we estimate the PPML model of Santos Silva and Tenreyro (2006), and construct the gravity score (i.e., predict trade in each direction in levels, take logs, and sum). We find that the gravity score constructed in this way has a correlation coefficient of 0.915 with our gravity score. Thus, it appears that the two approaches lead to similar results.

Appendix F: Robustness tables

This section presents Tables 3 and 4 using the alternative construction of the data set described in Appendix B. Results are similar in terms of signs and statistical significance. The main difference is in Table 11, where we see that not using the moments drawn from the refinements of pairwise stability leads to an insignificant coefficient on betweenness centrality.

[26]Formally, we specify our model as $\ln(y_i) = x_i\beta + \varepsilon_i$ and assume $E[\varepsilon_i|x_i] = 0$, which differs from Santos Silva and Tenreyro (2006).

Table 10: Main estimation results with alternative data set

	I		II		III	
	Lo	Hi	Lo	Hi	Lo	Hi
Betweenness	0.28	21.32			0.76	23.93
Eigenvector			-75.85	-7.57	-75.84	52.69
Constant	-13.82	-6.52	-4.62	12.87	-55.08	3.55

Note: The table reports the extremum of the confidence set for each parameter in the specification. Each specification uses multiple moments with different numbers of observations as described in Table 2. Relative to Table 3, this table uses the alternative data set described in Appendix B.

Table 11: Robustness results with alternative data set

	I		II		III		IV	
Betweenness	-1.33	59.42	-1.29	22.97	8.57	100.00	15.68	92.06
Eigenvector							-0.48	100
Distance	-29.21	0.17						
GDP	-0.75	5.21						
Constant	-100.00	1.25	-14.45	-6.44				

Note: Panel I includes additional explanatory variables. Panel II drops the refinements of pairwise stability (Moments 5 and 6). The table reports the extremum of the confidence set for each parameter in the specification. Each specification uses multiple moments with different numbers of observations as described in Table 2. Relative to Table 3, this table uses the alternative data set described in Appendix B.