

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

April 24, 2018

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 signed but not others, we study a new data set on air services agreements from the perspective of strategic network formation. Since the signing of an agreement can have implications for countries other than the two directly involved, this is an environment where externalities are likely to be important. These externalities also suggest that there are multiple equilibria for any candidate set of parameters, creating a fundamental identification problem for any econometric analysis. To address this issue we develop a structural model based on moment inequalities that uses the concepts of pairwise stability to generate estimating equations and also introduce methods to implement refinements of pairwise stability. The network structure is found to be important in determining the choices of countries to form agreements, and that the jointly optimal network of agreements would be substantially different than the observed outcome.

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 provide 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 relatively incomplete. For example, for the largest 58 countries, just over half of potential bilateral agreements are realized by an ASA. In contrast, we observe that almost all of these countries 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 between nodes. Our most basic question is whether countries account for the link structure in choosing whether to form agreements. Do countries account for the agreements that potential partners have signed, or do countries just focus on bilateral characteristics? We further ask what elements of the link structure are important, which then implies what externalities are prevalent in the network of agreements. 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. 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.

Pairwise stability places conditions on how observed choices must generate higher payoffs than unobserved choices. This revealed preference logic leads to a set of inequality conditions. Still, predicting outcomes is problematic since 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 predict multiple equilibria. The typical approach adopted is to restrict the environment in some dimension to enhance the feasibility of identifying which equilibrium is selected. This generates methods that are either only applicable in low dimensional environments (small number of agents) or relatively divorced from equilibrium outcomes. Since we are interested in a setting with a large number of countries and would like to exploit the structure of the relevant concepts of economic stability, previous techniques are not feasible for our application.

Predicting outcomes directly is complicated by the issue of multiplicity of equilibria. Rather, we approach the problem by mapping these inequalities directly into their empirical counterparts, and estimate based on the recent literature that employs moment inequalities, such as Pakes et al. (2015). This approach deals directly with the identification issue and we are able to identify a set of parameters that are consistent with the observed outcomes. In practice, we follow a version of the method of Andrews & Barwick (2012) 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 & Meyer, 2013), such as the distance between countries, and in our preferred specification, we use predictions from a gravity equation model directly as an explanatory variable. An important element of our paper is the utilization of useful measures of network structure for capturing the issues of interest. We are motivated by

the externalities studied in the theoretical literature in network formation, particularly Jackson & Wolinsky (1996). They discuss the *coauthor externality* and the *information flow* externality, which we interpret in our context as the incentive to be a hub between otherwise unconnected countries, and the incentive to connect to well-connected countries. For empirical measures of these features, we turn to the vast literature on social networks (see Prell, 2011; Kolaczyk, 2009). We further discuss this literature below, but it provides a wide variety of measures of network features. We rely on *betweenness centrality* to measure hubbing incentives and *eigenvector centrality* to measure the desire for connected partners. For instance, if we observe countries form links that increase their betweenness centrality, we conclude that the desire to be a hub is important. We view developing a relationship between theoretical concepts of externality and empirical measure of network features to be a contribution of our paper.

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. The WASA data base aims to catalog all ASAs. Although we focus only on the existence or non-existence of ASAs, the WASA database contains data on ASAs features, such as whether the ASA is an Open Sky Agreement (which we discuss further below). The WASA database has significant drawbacks, such as the fact that it is missing data on a number of agreements, and that it is difficult to use the data on how long an agreement has been in effect. We discuss these issues and our response below, but surely, WASA is the best database that we are aware of for studying ASAs. To reduce the limitations of the WASA dataset, we augment this data with another from ICOA which tracks the number of flights between countries (TFS dataset). We match WASA up to trade and characteristic data on country pairs. 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. We use data from 2005. Afterwards, a number of new multilateral agreements appear, which would complicate our analysis.

We find that network structure plays a substantial role in determining choices. In particular, countries form agreements that increase their role as a hub between otherwise unconnected countries. That is, countries form agreements that raise their betweenness centrality. Additional moments generated from an empirical analog of Nash pairwise stability play a limited role in identifying parameters, but are important in some circumstances.

Our computation of the globally optimal network of agreements 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 more global, and arises from the preference to be a hub. If we address only the first issue, we would increase the number of agreements since considerations are only bilateral in nature. Further addressing the network issue causes the social planner to substantially change the composition of ASA since the externality is now global. Based on a conservative parameterization of the network variable, we find that these compositional changes are significant – implying that the observed set of ASA’s differs substantially from the efficient outcome.

A potential concern may be the endogeneity of the link 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. This issue is addressed in relatively few of the papers on strategic network formation that we discuss below. The information structure that we impose that generate moment inequalities also generates natural excluded variables, which we exploit as instruments for the network variables. The results allowing for endogeneity are similar to our main specification. Other approaches to endogeneity, such as differencing across similar observation or developing a control function based historical data, are the part of our current research.

To be clear, our research has several caveats. We derive the payoffs to countries from revealed preference about how the countries form agreements, and use this payoff function in calculating counterfactual payoffs. Thus, we ignore political economy concerns, that might cause countries to maximize something other than the welfare of their citizens.

2 Literature Review

The study of *social networks* is a vast field, stretching across psychology, sociology, anthropology, economics, statistics and even physics. See Prell (2011) for an excellent overview and a history of the field, and Kolaczyk (2009) for an overview of statistical issues. The literature is largely empirical, with substantial effort devoted to collecting data, describing data and creating better measures of network position, such as the measures of centrality described above. There are many examples in economics, such as Alatas et al. (2014). An interesting recent example in the field of industrial organization is Fershtman & Gandal (2011), which studies information spillovers in the context of networks of open-source computer programmers. An overview of research using exogenous matching to provide identification appears in Sacerdote (2014), and a more general overview of the econometric study of networks appears in de Paula (2016).

We distinguish between the literature on social networks, and the literature on *strategic network formation*. The social networks literature typically takes the formation of a network as exogenous, or as a reduced-form function of network variables, whereas the strategic network formation literature arises from economics and studies the incentives of agents to form links. Thus, it provides micro-structural models of the formation of networks, and emphasizes stability concepts, and efficiency issues. An early contribution is Jackson & Wolinsky (1996). Two recent overviews are Jackson (2008) and Goyal (2007). Naturally, there is substantial overlap between the literatures on social networks and strategic network formation, in terms of concepts, notation, and empirical examples, and even authors.

Our project fits into a growing number of papers that attempt to structurally estimate a model of strategic network formation.¹ A seminal contribution is Ho (2009), who estimates a model of hospitals joining insurance networks, which is motivated by the solution concept of stability to generate moment inequalities, in the spirit of Pakes et al. (2015). This case is two-sided, in the sense that hospitals match to insurance companies rather than to other hospitals. Several other

¹Note that it is possible to model network formation without modeling what we refer to as strategic network formation. Exponential Random Graph Models are a popular tool outside of economics for modeling network formation, typically as a reduced-form function of network features. Chandrasekhar & Jackson (2014) do in fact provide micro-foundations for a broad class of related models, although their specific assumptions, based on random meetings between groups of agents, probably does not describe our setting.

papers discuss methods or applications for estimating matching games – see Sorensen (2007), Fox (2010) and Agarwal (2013).

The last few years have seen rapid growth in methods for estimating strategic network formation models. Graham (2015) discusses this literature. Imbens & Kalyanaraman (2010), Goldsmith-Pinkham & Imbens (2013) and Mele (2013) present related methods based on Bayesian statistical estimators that require repeatedly solving for the choice of each agent. These papers solve their model by assuming that players make decisions according to an exogenous or random ordering, with myopic decision-making. Thus, their solution concept does not correspond to a standard solution concept in the theoretical literature, although Mele (2013) repeatedly cycles through the set of players, which can be shown to converge to a stable outcome. Hsieh & Lee (2012) also use a Bayesian method and the concept of a potential game in order to avoid issues of multiple equilibria. All three of these papers are motivated by friendship networks in surveys of elementary schools in the AddHealth data.

Sheng (2012) and Miyauchi (2014) use a technique that bears a similarity to the approach of Ciliberto & Tamer (2009) in the entry literature. For a given set of parameters, Sheng (2012) computes the maximum and minimum of the probability of observing a given link structure, and uses these with the observed probabilities to form moment inequalities. To lower the computational complexity of this problem, the paper holds most decisions constant at their observed outcome, focusing on one “sub-network” at a time. Miyauchi (2014) similarly takes a strategy of bounding moments from the data, in this case focusing on the case of non-negative externalities across nodes (which is not satisfied in our model) to simplify the problem.

Leung (2015) models agents in a game of imperfect information, and recommends a two-step estimator that addresses endogeneity and multiple equilibria in a way similar to Bajari et al. (2007). Boucher & Mourifié (2015) develops a likelihood framework, focusing on the case of no externalities between agents. Currarini et al. (2009) develop an estimator based on an underlying model of search and matching between agents. Badev (2013) does as well, and further studies the choice of smoking and how it interacts with friendship formation. The vast majority of these papers study friendship pattern in surveys of school students.

The paper by de Paula et al. (2014) presents a model for aggregate data of populations that meet and match. Graham (2014) studies network formation in a model with transferable utility, and is particularly concerned with the issue of omitted variables. While the paper focuses on the case in which links are formed just on bilateral characteristics, he introduced network structure into the problem using a conditioning argument similar to the solution of Chamberlain (1980) for the fixed-effects logit.

The paper most similar to ours in terms of methods is Uetake (2014), who also uses the concept of stability to generate moment inequalities drawn directly from the consumer's utility functions. Unlike our method, Uetake (2014) simulates the structural error terms in the agent utility function, which requires specific assumptions in a moment inequality context applied to games. Also, Uetake focuses on the case where the researcher observes many networks. The paper studies the propensity of venture capitalists to work with similar venture capitalists (similar in terms of observable characteristics), and less-so on network structure. Note that all of the papers discussed are written for cross-sectional data. Fong & Lee (2013) present a method for studying matching in a dynamic context.

Relative to all of these papers, our project makes several contributions. Our method relies directly on theoretical stability concepts and so is attractive from the perspective economic theory, and we are the first to exploit refinements to pairwise stability, such as Nash pairwise stability, which would be difficult to do in other approaches. Our method can handle large numbers of agents and our method is relatively fast. Our project emphasizes the use of network variables such as centrality, and the link to theoretical concepts such as the co-author externality that we discuss below. Our goal of characterizing externalities and comparing the optimal network to the observed network is central to the theoretical goals of the literature on strategic network formation. Furthermore, our application to agreements between countries is substantially different from other papers, and the resulting conclusions that we can reach are thus quite different as well. Our empirical approach requires strong assumptions on error terms, as in Pakes et al. (2015) and related papers. But as can be seen, all methods require important restrictions in order to make progress. We believe that our project represents an important contribution to the discussion of

appropriate methods for this area, and is complementary to existing work in this area.

Our project also bears on the field of international trade. Some previous theoretical work studies treaties, typically free-trade agreements, through the prism of network formation games, such as Goyal & Joshi (2006), and Furusawa & Konishi (2007). There is also a related empirical literature on the determinants of free-trade agreements, such as Baier & Bergstrand (2004), Egger et al. (2011) and Chen & Joshi (2010). Employing dichotomous dependent variable models, these papers examine comparative static predictions based on unilateral best responses. While equilibrium outcomes are not characterized, this approach implicitly assumes a unique equilibrium, despite acknowledging a range of externalities associated with preferential trade agreements. For an overview and critique of this literature see Limão (2016). In contrast, we allow for a multiplicity of equilibria that can arise in setting with externalities and use the equilibrium definition as the basis for our empirical methodology.

A small empirical literature studies air services agreements, primarily on the topic of Open Sky Agreements. The existing empirical literature largely takes the air services agreements as an explanatory variable that can be used to predict outcomes, such as trade, whereas our project treats the air services agreements as the endogenous variable to be explained. Naturally, the extent to which air services agreements affect economic outcomes is an important justification for our study of the formation of agreements, so we view this literature as highly complementary. Cristea et al. (2012) study the impact on trade of US Open Sky Agreements, using panel-data techniques. Micco & Serebrisky (2006) also look at this issue. There is a substantial literature on air services agreements that exists outside of mainstream economics journals, primarily in fields such as operations research, engineering, and transportation research. One example is Dresner (2008). A recent report commissioned by the Department of State also finds large benefits to agreements (Intervistas, 2015).

3 Industry and Data

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.² 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, essentially 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.³

Open Sky Agreements are perhaps more widely known than air services agreements, although formally, Open Sky Agreements are a subset of air services agreements. Open Sky Agreements are a particularly liberal version of air services agreements. While the Chicago Convention provides a formal definition of Open Sky Agreements, few agreements today live up to that definition.⁴ Open Sky 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, 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.

The Chicago Convention envisioned only bilateral agreements, which is helpful for our purposes since multi-lateral agreements are more difficult to model, particularly in the context of strategic network formation. In practice, there is one multilateral air services agreement, which is between members of the European Union. Since 1992, EU members have participated in a multi-lateral, very liberal agreement. However, up until 2002, EU countries continued to sign bilateral agreements with non-EU countries, rather than having the EU negotiate as a whole.⁵

²Odoni (2009) provides an excellent discussion of the institutional background for air services agreements.

³Note 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.

⁴For instance, formally, Open Sky would mean allowing domestic cabotage, which is allowing foreign airlines to pick up and drop off passengers flying within a country. While the US would say that it has signed many Open Sky agreements, none of them allow that.

⁵The use of bilateral agreements was subject to a lawsuit filed by the European Commission against member countries.

We address this issue in the construction of our data. Since 2005, several more multi-lateral air services agreements have arisen or have been proposed, particularly in Asia, and so we use 2005 data in order to avoid this issue.

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, that is the centerpiece of our data set. At this stage, we make use of only the existence of an agreement, and we model whether countries form an agreement or not.⁶ 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.⁷ Thus, we ignore the time dimension of our data set, and treat our data as a cross-section of existing agreements.

The benefit of an agreement is that it allows for direct flights between countries. Travel between countries would otherwise require connecting flights, which can exist only if connecting countries have signed ASAs. The effects of an ASA can easily expand beyond air travel. If air travel supports communication between executives at trading companies, the impacts of an ASA may be much larger in shipping than air. One question is why countries do not sign ASAs with every other country. Signing agreements has a cost in terms of administration and negotiator's time. Furthermore, a country may be opposed to an ASA in some cases. If one country believes that its airlines will not be competitive on a particular route with the airlines

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 multi-lateral agreements between an outside country and all EU countries simultaneously. See http://ec.europa.eu/transport/modes/air/international_aviation/external_aviation_policy/horizontal_agreements_en.htm.

⁶WASA also contains a list of indicator variables describing the agreement, so it is possible to utilize more information about agreements in the future. An strand of the strategic network formation literature analyzes the intensity of a link, and these indicator variables could be understood along these lines.

⁷For 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.

of another country, that country may be reluctant to open up that route. In the United States, the State Department has an Aviation Negotiation division that is dedicated to negotiating ASAs. Their website highlights the number of Open Sky Agreements they have signed over the last several years. Naturally, we are not privy to how the division decides with which countries to negotiate.

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.⁸ In order to address this problem, we have obtained another database from ICAO, the TFS database. For our purposes, the TFS database indicates which pairs of countries had direct flights at an annual level, either for people, cargo or mail. ICAO representatives argue that this can be useful since any direct flight requires some kind of agreement between the two countries in advance. Note that this database does not allow us to study the details of the agreement, it only determines the existence of an agreement. The TFS database provides an alternative way to construct the network of agreement. In general, the network drawn from the TFS database has more links than WASA, but one is not perfectly contained in the other. In what follows, we provide results based on a merged TFS and WASA database.

We start with the CEPII data set created for Head et al. (2010), which contains several useful bilateral variables, such as distance and indicators for a common border and a common language. There are several matching issues that we address, further described elsewhere. For variables that vary over time, such as GDP and population, we use the average of 2000-2006. Check this!

To keep our data reasonably homogenous and in order to keep our estimation tractable, 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

⁸In private communication, an ICAO representative guessed that WASA has between one-half and two-thirds of agreements, although it is difficult to know. 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.

countries with population less than 500,000 as well as Puerto Rico and Taiwan, whose freedom to negotiate agreements is unclear.⁹ We add back in Iceland although its population is less than 500,000 because Iceland has a prominent national airline. Overall, we have a list of 58 countries from 6 continents. That creates a list 3,306 country-pairs.

From the perspective of the social networks literature, we term the countries as *nodes* and the agreements as *links*. The *degree* is the number of links to a node. In our computations, we count links from EU countries to outside countries but not EU countries to other EU countries.¹⁰ 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 Common Aviation Area (CAA) agreements by 2005 (see European Commission, 2010). Note that Lichtenstein is dropped from our final data set because it has low population. We do not count Western Baltic states, which signed CAA agreements in 2010. We assume Russia inherits all of the agreements of the Soviet Union. The only other post-Soviet state in our sample is Estonia. We assume Estonia does not inherit these agreements. We assume the Czech Republic and Slovenia inherit all of the agreements of Czechoslovakia. Since we drop all EU pairs, we end up with 1400 pairs of countries, or 2,800 decisions by potential partners.

4 Social Network Variables

In our empirical work, we ask how a country values connectedness. We focus on three measures of connectedness. The first is degree, the number of links a country has formed. While straightforward, the literature on social networks has recognized that this measure is often too simple to capture issues of interest, and has developed a number of other measures, broadly termed *centrality measures*. Centrality measures capture more complete measures of the location of a node in a network. For instance, we might be concerned not just with how many links a node has, but also with whether those links are to nodes that themselves have many links. Similarly, we might

⁹Dropping countries with low populations eliminates Antigua, the Bahamas, Barbados, Brunei, Equatorial Guinea, Iceland, Luxemborg, Macau, Malta, St. Kitts, and the Seychelles.

¹⁰This reflects the state of the EU in the early 2000s, which we believe best describes our 2005 data. Keep in mind that non-EU countries have an opportunity for a higher number of links than EU countries, since EU-to-EU links are deleted.

care whether those links are to nodes that are connected amongst each other, or whether a given node is the sole pathway between sets of nodes.

Some notation aids in developing these concepts. Let the set of countries be $\mathcal{N} = \{1, \dots, n\}$. Let Λ be an $n \times n$ matrix. Element $\{\i, j\}$ of matrix Λ equals 1 if countries i and j have an ASA, and 0 otherwise. Thus, Λ represents the set of links between nodes. The diagonal of Λ is 1 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 $s_i(\Lambda)$.

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 wish to solve:

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

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 highest eigenvector, 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 almost always reach each other in two links, but there may be multiple paths by which to do so. Let $P(k, j)$ be the number of geodesics between nodes k and j . Let $P_i(k, j)$ be the number of geodesics between k and j that pass through i . *Betweenness centrality* is:

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

To see how these two measures differ, consider Figure 1. The figure has 5 nodes, $A - E$, and 5 links. The figure displays the nodes and links, as well as the resulting measures of eigenvector and betweenness centrality.¹¹ 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

¹¹For clarity, we display $P_i(kj)$ instead of C_i^{betw} .




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

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 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.¹²

Note that there are no microfoundations for these measures. That is, there is no game-theoretic model of communication across a network or formation of a network from which these measures arise. However, 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 & 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.

The model of the co-author externality recognizes that 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

¹²There exists useful extensions of these measures that account for nodes that vary in importance. That is, it is more valuable to be between important nodes (say, countries with large GDP) than unimportant nodes, and centrality measures can reflect this. We are further developing this issue.

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 this externality is well captured by betweenness centrality. If we observe countries forming agreements that tend to raise the country's betweenness centrality, we will say that countries prefer to be hubs, and that the coauthor externality is present. Of course, 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. Note that as Jackson & 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.

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, where more connected paths between any two countries provide more options for that country-pair market. We believe that this measure is well captured by eigenvector centrality, which measures the overall connectedness of any given node. In either the co-author or connections case, if B and C make their decision without accounting for A , we have an externality. 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.

Now we turn to some simple statistics. The US, Russia and Singapore have the highest degrees (numbers of links), with 55, 47 and 46 respectively. The highest degrees within the EU are France and the UK, each with 35.¹³ The mean number of links per country is 22.44, the standard deviation is 11.1, the median is 22 and the 25th percentile is 13. Thus, there is substantial variation

¹³If we counted agreements in the data between EU countries, the most linked country overall would be Switzerland, and EU countries would make up most of the top 10.

in the number of links.

Our network is highly connected. Connected paths exist between every node, usually requiring only one or two links. We find that 58% of country-pairs are directly connected, and 41% of country-pairs require only one intermediate node to reach each other. The remaining 1% of country-pairs require two links to reach each other, and no country-pair requires more.

Table 1 presents the network statistics that we are interested in. The table is ordered by degree and presents the top 15 countries. It also presents betweenness centrality and eigenvector centrality, along with the rank of each country in these statistics. For instance, we see in the first row that the United States has 54 agreements and has the highest values for both betweenness and eigenvector centrality. The top 4 of each statistic is held by the US, Singapore, Russia and China almost exclusively. Interestingly, Russia is first in betweenness centrality but 5th in eigenvector centrality. Similarly, India is 3rd in Eigenvector centrality but 9th in betweenness. That is, although India is well connected, it relatively rarely provides an intermediate travel point for countries that are not otherwise connected. Hong Kong is similar, which ranks 12th in eigenvector centrality but only 21st in betweenness. While clearly correlated, the measures of centrality capture different issues, and our project uses these differences to gain insight into the nature of externalities that govern the structure of ASAs.

One concern we might have is that our measures just count links, as if all links were 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. CITE discusses 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 both trade and (the inverse of) distance. We discuss the exact calculation in Appendix A, but ultimately we find that our results are robust to using weighted statistics.

Country	Degree		Eigenvector Centrality		Betweenness Centrality	
	Value	Rank	Value	Rank	Value	Rank
United States	54	1	0.97	1	80.98	2
Russia	47	2	0.85	5	86.21	1
Singapore	46	3	0.90	2	28.99	8
China	46	3	0.86	4	64.12	3
India	44	5	0.87	3	25.12	9
Turkey	44	5	0.82	8	45.72	4
Pakistan	44	5	0.85	6	34.71	6
Thailand	43	8	0.84	7	29.55	7
Egypt	40	9	0.79	11	23.18	13
Canada	40	9	0.73	15	38.10	5
Japan	40	9	0.81	9	17.99	17
Hong Kong	38	12	0.79	12	13.26	21
Malaysia	38	12	0.77	13	24.80	10
South Africa	38	12	0.74	14	20.71	14
South Korea	38	12	0.80	10	11.63	23
France	35	16	0.69	19	23.63	11
United Kingdom	35	16	0.69	19	23.63	11

Table 1: Comparing centrality statistics.

5 Model and Refinements

In this section, we present our theoretical model and we derive inequalities that we bring to estimation. We discuss refinements to the solution concept in our base model. We also present a brief subsection on endogeneity in our model.

5.1 Model

This section presents our model of network formation, from which we derive estimating equations. 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 profit 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.

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} . There is some \tilde{x}_{ik1} such that $\tilde{x}_{ik1} = x_{ik1} + \varepsilon_{ik}$. The country makes its decisions based on $E[x_{ik1} | \tilde{x}_{ik1}]$, which is just \tilde{x}_{ik1} . Conditional on the agent's information set, the expectation of ε_{ik} is 0. That is $E[\varepsilon_{ik} | \tilde{x}_{ik1}, x_{ik2}, \Lambda] = 0$. That is because the country's decision depends only on \tilde{x}_{ik1}, x_{ik2} , not on ε_{ik} . In fact, conditioning on any subset of the agent's information leads to the same result. Pakes et al. (2015) and Dickstein & Morales (2013) further discuss how ε_{ik} can be interpreted

either as expectational on the part of the agent, or measurement error. The implication is that we can assume that $E[\varepsilon_{ij}|x_{ij2}, \Lambda] = 0$. Although we cannot condition on x_{ik1} in this statement, the conditioning on Λ is key to developing moment inequalities.

It is certainly a reasonable assumption that the country cannot predict its payoff from an agreement exactly, because it faces this measurement error. In different specifications, we use the gravity score or the product of GDPs as x_{ik1} , but any country has difficulty measuring these variables accurately, and moreover, their mapping into the benefits from an ASA are also uncertain. 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 assume that ε_{ik} is mean independent of the outcome Λ . Below, we discuss adding an error term that is known to the country.^{14 15}

Thus, 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 $\Lambda_{ij} = 0$, and if Λ does not contain a link between i and j , then $\Lambda - \{i, j\} = \Lambda$.

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

$$\pi_{ij} = \Pi_{ij} - \Pi_{i-j}$$

¹⁴Dickstein & Morales (2013) suggest using an instrument for the mismeasured variable, in a context with a structural error term with a normalized variance. As discussed below, we instead normalize the parameter on the mismeasured variable to 1 in a context with no structural error term.

¹⁵We do not take a position on whether $\Pi_i(\Lambda)$ represents true welfare for the country, or is manipulated by issues of political economy. We regard it as the country's objective function, no matter its source. Naturally, we can include variables representing political economy concerns in x_{ik} .

$$= \psi_i(\Lambda \cup \{i, j\}, \alpha) - \psi_i(\Lambda - \{i, j\}, \alpha) + x_{ij}\beta + \varepsilon_{ij} \geq 0. \quad (1)$$

Since the country cannot observe ε_{ij} , the countries instead expects the payoff from an agreement to be positive if:

$$E[\pi_{ij}] = \psi_i(\Lambda \cup \{i, j\}, \alpha) - \psi_i(\Lambda - \{i, j\}, \alpha) + x_{ij}\beta \geq 0$$

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. $E[\pi_{ij}] \geq 0 \quad \forall \quad \{i, j : \Lambda_{ij} = 1\}$.
2. $\min\{E[\pi_{ij}], E[\pi_{ji}]\} \leq 0 \quad \forall \quad \{i, j : \Lambda_{ij} = 0\}$.

The expectation is taken over ε_{ij} , which is assumed to be orthogonal to decision-making. Note that point 1 is a cooperative concept. A network fails pairwise stability if *two* players would like to form a link that have not done so. Point 2 considers only a unilateral deviation. That is, a link does not survive pairwise stability if either player wishes to sever the link.

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 use his surplus from the link to pay the agent to maintain the link. We find a model with non-transferable utility more natural in this environment. However, estimating under a model with transferable utility is feasible, and we do so as a robustness check. With transferable utility between pairs of agents, the definition would be 1) $E[\pi_{ij} + \pi_{ji}] \geq 0$ for all $\{i, j : \Lambda_{ij} = 1\}$ and $E[\pi_{ij}] + E[\pi_{ji}] \leq 0$ for all $\{i, j : \Lambda_{ij} = 0\}$.¹⁶

We now turn to estimation. We exploit the definitions of pairwise stability to generate moment inequalities, following Pakes et al. (2015). Thus, we can use moment inequalities based on the

¹⁶If 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.

definition of pairwise stability:

$$E [\psi_i (\Lambda, \alpha) - \psi_i (\Lambda - \{i, j\}, \alpha) + x_{ij}\beta | \Lambda_{ij} = 1, \Lambda, x_{ij}] \geq 0 \quad (2)$$

$$E \left[\min_{k \in \{i, j\}} \psi_k (\Lambda \cup \{i, j\}, \alpha) - \psi_k (\Lambda, \alpha) + x_{kj}\beta | \Lambda_{ij} = 0, \Lambda, x_{ij} \right] \leq 0 \quad (3)$$

These equations capture that for any pair with an agreement, we know that both countries prefer an agreement to no agreement, but that for each pair without an agreement, we know only that one of the two countries prefers no agreement. We can further interact elements of these conditions with x_{ij} and network statistics to generate more moments.

As is often the case with estimation based on moments, the researcher has some choice in how to specify moments. For example, consider Equation 2, which states that the average of payoffs to observed links must be positive. An alternative moment, equally well motivated by the theoretical model, is that for each country, its lowest payoff among links is positive. Formally:

$$E \left[\min_{j \in s_i(\Lambda)} \psi_i (\Lambda, \alpha) - \psi_i (\Lambda - \{i, j\}, \alpha) + x_{ij}\beta | \Lambda, X \right] \geq 0 \quad (4)$$

Equation 4 should be more informative than Equation 2 because saying that the minimum of a set is positive is more binding than saying the average is. However, this moment has only one observation per country, and so is estimated with substantially less precision. We apply both moments, and we explore which moment is more informative in practice below. We sometimes refer to the moment inequalities 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 agents would wish to sever multiple links. This idea leads to a refinement of pairwise stability called *Nash pairwise stability*, and is discussed in Jackson & Wolinsky (1996). Similarly, an agent 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 & Mutuswami, 1997). Such a deviation requires coordinated action between multiple agents, whereas 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 identification. 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 limited 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} , $m = 1, \dots, \#\mathcal{S}_{ik}^+$, 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:

1. Pairwise stability holds, and
2. $E[\Pi_i(\Lambda)] \geq E[\Pi_i(\Lambda - \sigma_m)] \quad \forall \quad i = 1, \dots, n, k \leq K, \sigma_{mik} \in \mathcal{S}_{ik}^-$:

The second element of 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

$K \geq n$.¹⁷

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.

$$E \left[\psi_i(\Lambda, \alpha) - \psi_i(\Lambda - \sigma, \alpha) + \sum_{j \in \sigma} x_{ij} \beta \middle| \sigma \in \mathcal{S}_{ik}, \Lambda, x_{ij} \right] \geq 0 \quad (5)$$

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 possible interactions with instrumental variables such as x).¹⁸

Similarly, we can define a stability concept motivated by strong stability:

Definition A network Λ is *pairwise stable with addition degree K* :

1. Pairwise stability holds, and
- 2.

$$\min \{ E[\Pi_i(\Lambda + \sigma_m)] - E[\Pi_i(\Lambda)], E[\pi_{ji}] \forall j \in \sigma_m \} \leq 0 \quad \forall i = 1, \dots, n, \forall k \leq K, \forall \sigma_m \in \mathcal{S}_{ik}^+. \quad (6)$$

The second element of the definition 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 of because one of the potential partners does not want to link to i . This definition leads to additional moments, analogously defined to Equation 5.¹⁹

¹⁷There 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.

¹⁸Computationally, 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 not computationally expensive. Thus, exploring $K > 2$ appears feasible, although we have not done so yet, for reasons described with the results.

¹⁹The concept of strong stability would allow for deviations that included both deletion and addition of multiple links. We have not explored this, but developing additional moments expanding on what we have is straightforward.

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 $K = 44$ to $K = 45$? We observe only one country with more than 44 links, so we would be forming a moment with one observation, which will have infinite variance, and thus no 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 44 of its links, so even if there were a few more countries in this set, the moment would be of little value.

5.3 Endogeneity

A central question in any paper on network formation should be exogeneity. For example, if we see that country A forms a link with country B that has many links, we want to know whether A was attracted to B because it has many links or because there was an unobservable feature of B that made it attractive both to A and to other countries. Technically, we can allow for an extra term v_{ij} in Equation 1 that is unobserved to the econometrician but observed by agents and is thus endogenous to the network Λ . The unobserved term ε_{ij} continues as white noise to both the econometrician and the agents. The value of a link is then:

$$\pi_{ij} = \psi_i(\Lambda \cup \{i, j\}, \alpha) - \psi_i(\Lambda - \{i, j\}, \alpha) + x_{ij}\beta + v_{ij} + \varepsilon_{ij}.$$

We address this problem with an instrumental variables approach. That is, we exclude the network variables from the instrument vector for the moment inequalities. Instead, we include variables that describe the bilateral relationship between the two countries, such as distance and whether the countries have a trade agreement. We further describe these variables and their motivation below.

There are several alternative approaches that we might take to addressing endogeneity. For instance, Pakes et al. (2015) suggest assuming that v_{ij} is constant across some observable features of countries i and j , and then differencing this equation across these countries to eliminate the v_{ij} term. In this approach, we consider pairs of countries with similar such observable characteristics

but different values of network statistics, and possibly different choices of whether to form an agreement. Covariation in these variables identifies α in $\psi(\Lambda, \alpha)$.

An alternative approach to endogeneity relies on a control function. We search for a variable that proxies for unobserved features of each country. For this, we focus on older trade data – we are currently using data from the mid-1950’s. This trade data is available in the data set from Head et al. (2010). Constructing centrality statistics from these data is valuable because they should be highly correlated with the intrinsic value of trading with these countries currently, but since these data largely predate the modern commercial airline industry, the statistics should not be endogenous to events in this industry (NOT IMPLEMENTED YET).²⁰

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 the definition of pairwise stability, define $\mathcal{I}(\lambda_{ij} = 1)$ as the set of all pairs of countries $\{i, j\}$ that form an agreement, and $\mathcal{I}(\lambda_{ij} = 0)$ as the set of all pairs of countries $\{i, j\}$ that do not form an agreement. The number of pairs of countries that form an agreement is n^a and the number that do not is n^{na} . We define $\mathcal{I}(\lambda_{ij} = 1)$ and $\mathcal{I}(\lambda_{ij} = 0)$ to exclude pairs of countries in which both are members of the EU. Furthermore, let z_{ij} be a vector of instruments for pair i, j . Also, we now impose that α affects the network measures as a multiplicative coefficient. Then, our sample moments are:

$$\frac{1}{n^a} \sum_{\{i,j\} \in \mathcal{I}(\lambda_{ij}=1)} \left(\alpha (\psi_i(\Lambda \cup \{i, j\}) - \psi_i(\Lambda - \{i, j\})) + x_{ij}\beta \right) z_{ij} \geq 0 \quad (7)$$

²⁰Naturally, there is a tradeoff between collecting older data that is less affected by the airline industry and more recent data that is more highly correlated with contemporary outcomes. Also, much older data do not exist for many country-pairs.

$$\frac{1}{2n^{na}} \sum_{\{i,j\} \in \mathcal{I}(\lambda_{ij}=0)} \min_{l \in \{i,j\}} \left\{ \alpha (\psi_l(\Lambda \cup \{i,j\}) - \psi_l(\Lambda - \{i,j\})) + x_{ij}\beta \right\} z_{ij} \leq 0$$

Hopefully, it is now clear why we require a model in which Λ is exogenous to ε_{ij} . If not, we could not assume that ε_{ij} is mean zero in each equation, in which case each equation would not hold just in observable variables, and we would not have a basis for estimation.

Moment inequalities lead to partial identification. In order to determine whether a point is in the confidence interval surrounding the identified set, we follow the algorithm presented in Andrews & Barwick (2012). Under their framework, we use the objective function that they ascribe to Pakes et al. (2015), and we use the “normal approximation” that they suggest, which assumes that moments are normal for purposes of inference rather than relying on simulation. The normal approximation substantially decreases the computational time of the estimator. 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 intervals. We then report only these maximum and minimums, rather than the shape of the confidence intervals. Many of our specifications use only two parameters, and so we use grid searches and graphs of the confidence interval to gain a deeper understanding of our estimator. A fuller discussion of our estimation approach appears in Appendix B.

In addition to the moments in Equation 7, we utilize two more sets of moments. The first are derived from the stability refinements discussed in Section 5.2. Just as Equation 7 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 $\tilde{n}^- = \sum_{i=1}^n \#S_{i2}^-$, the total number of elements of all sets S_{ik}^- for all $i = 1, \dots, n$.

$$\frac{1}{\tilde{n}^-} \sum_{i=1}^n \sum_{m \in S_{i2}^-} (\alpha (\psi_i(\Lambda) - \psi_i(\Lambda - \sigma_{mik})) - (x_{ij} + x_{il})\beta) z_{ijl} \geq 0, \quad \{j,l\} = \sigma_{mik}. \quad (8)$$

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 additional degree 2.

A second set of moments that we utilize are the *min moments*, as presented in Equation 4. Averaging over all links as in Equation 7 mixes together many different types of links, both those that are very valuable and those that are only marginally valuable. We may learn more if we focus on those that are marginally valuable, that is that the least valuable link has positive value. Let z_i^m be the instrument vector for the min moment applied to observation i . The min moment is:

$$\frac{1}{n} \sum_{i=1}^n \left(\min_{j \in s_i(\Lambda)} \psi_i(\Lambda, \alpha) - \psi_i(\Lambda - \{i, j\}, \alpha) + x_{ij}\beta \right) z_i^m \geq 0 \quad (9)$$

Not surprisingly, we find that this moment provides a tighter, but imprecise, bound to α and β .

6.2 Empirical issues

There are several more issues to be dealt with before turning to results. 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.²¹

It is a challenge to implement estimation via moment inequalities in the context of many re-

²¹Some 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 or the product of GDPs has a positive effect on the likelihood of an agreement.

gressors. In part, this is a computational issue. Since the computational algorithm often involves something like grid search, it can be difficult to implement with many parameters. Also, this points to an inherent lack of robustness of estimation in the context of partial identification. With typical estimation based on equalities, such as estimation of a linear model via ordinary least squares, adding explanatory variables that are orthogonal to existing explanatory variables has no impact on the coefficients of the existing explanatory variables. However, with moment inequalities, we will identify only a range of parameters for this new orthogonal variable, and that will typically expand the range of parameters that are part of the confidence interval for all of the other variables. Thus, estimation via moment inequalities is not robust to adding orthogonal variables in the same sense as standard 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, since trade and ASAs are both forms of “links” between countries (see Head & Meyer, 2013). 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 C. 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 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 variables such as distance separately as explanatory variables in our agreement formation model. That is, they can be separate values of x_{ijk} . 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 our baseline model, we let the instruments z_{ij} be x_{ij2} and $(\psi_i(\Lambda \cup \{i, j\}) - \psi_i(\Lambda - \{i, j\}))$. However, we can use our specification of the gravity score equation to generate extra instruments. In particular, the model that derives the moment inequalities requires that the country does not observe x_{ij1} , but it is consistent with the model if the country observes predictors of x_{ij1} . 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 still 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 exactly 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. Because we are concerned that the network may be endogenous, we experiment with dropping the term $(\psi_i(\Lambda \cup \{i, j\}) - \psi_i(\Lambda - \{i, j\}))$ from the instrument matrix, and instead relying on this expanded instrument vector to provide identification.

Note that 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. In addition, we scale all of the network variables by the mean of betweenness centrality, so all of the same mean and are comparable in variation.

6.3 Asymptotics

Asymptotic properties are challenging in network context such as ours. We do not provide formal results, although there are a number of related formal treatments. The study of 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. In addition, this is the approach taken of all papers we know of that study the asymptotic properties of gravity equation estimators. See Egger & Staub (2014), Charbanneau (2014), Jochmans (2015) and Cameron & Miller (2014).

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.²² One way to imagine addressing this issue is 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. However, to be clear, we do not compute our network variables to reflect this notion.

7 Results

For the baseline results, we compute the network variables but we weight links by the inverse of log distance, so for instance, being between two countries when your path is very long counts less towards betweenness than if the path is short. A full description of our computation of weighted network statistics appears in the appendix. We also impose 5 sets of moments, one for linked pairs, one for unlinked pairs, one for the min moments, one for linked triplets of countries and one for unlinked triplets, the last two sets of moments capturing the concepts in Section 5.2. Each

²²We thank Antonio Mele for making this point to us.

Observations	Linked pairs	Unlinked pairs	Countries	Linked triplets	Unlinked triplets
	1,624	588	58	4,599	6,678
Relevant Equation:	2	3	4	5	6

Table 2: Number of observation in each moment

Explanatory Var.	Linked pairs		Unlinked pairs		Countries		Linked triplets		Unlinked triplets	
	Low	Hi	Low	Hi	Low	Hi	Low	Hi	Low	Hi
ln(Gravity Score)	1	1	1	1	1	1	1	1	1	1
Constant	-16.612	-6.163	-6.384	-5.954	-100	-7.787	-14.612	-9.264	-8.279	-6.814
Between Cent.	0.191	34.78							1.248	5.506
Eigenvalue Cent.			0.327	1.499			0.127	2.471	-0.206	3.818
Closeness Cent.					12.782	100	8.191	100	-8.28	3.81

Notes: All specifications are based on five moment equations: Linked pairs, unlinked pairs, the min moments, linked triplets and unlinked triplets. Different number of observations apply to each moment. The instrument vector consists of explanatory variables, except for ln(gravity score). The table reports the highest and lowest value within the confidence set for each variable. We do not search past 100, and such entries can be regarded as unbounded.

Table 3
Results for different network variables

moment is interacted with a set of instrumental variables, which in the baseline case consists of all explanatory variables except for gravity score. The number of observations varies across sets of moments, as described in Table 2.

Table 3 provides the confidence sets for a range of specifications. As noted above, we adopt the normalization that the coefficient on the log gravity score is unity. The baseline results are presented in first two columns, where column (1) provides the minimum bound for a parameter associated with the confidence set – that is, the column reports the argmin for each variable rather than the vector associated with the minimum. Column (2) reports analogous results for the maximum.

The first three specifications use only the gravity score, a constant term, and one network variable: betweenness centrality, eigenvector centrality and closeness centrality. We have not previously described closeness centrality, but it is the average of the minimum of the distance of a node to all other nodes. It captures something similar to eigenvector centrality and we use it as a robustness check for that variable. In each case, we see that the 95% confidence interval for the coefficient on the network variable (α in $\alpha(\psi_i(\Lambda \cup \{i, j\}) - \psi_i(\Lambda - \{i, j\}))$) is bounded above zero, suggesting that network position matters to countries when they choose their agreements.

In determining the confidence interval, we do not search above 100 or below -100, so reporting those coefficients is equivalent to saying that the confidence interval is unbounded. Thus, there is no upper bound on the parameter on closeness centrality.

The last two panels try different combinations of the network variables. The fourth panel shows that including eigenvector and closeness centrality simultaneously does not change the results by much. However, the last panel shows that including all three measures simultaneously does affect the results. In particular, betweenness centrality is still bounded above zero, but the other two measures of centrality have confidence intervals that bracket zero. Thus, we conclude that the result that betweenness centrality is important is the most robust of these results.

The next table studies which moment sets provide identifying power. The first panel in Table 4, labeled *baseline* repeats the first panel of Table 3. The next three panels provide results with only three sets of moments rather than five. We always include the moments for linked pairs and unlinked pairs (based on the equations in 7. Each specification includes one additional moment, either the min moment (based on Equation 9) or one of the two refinements (based on Equation 8 and the analog for the addition case). The result when including the min moment is exactly the same as the baseline case, whereas using the other sets of moments (and dropping the min moment) leads to unbounded results. Thus, while we believe the ability to impose refinements of pairwise stability is a general advantage of our method, it does not provide any power in this example. We return to a case where they are useful below. But before doing so, it is useful to consider why refinements of pairwise stability are not helpful. In the case in which the network payoff is measured by betweenness centrality and the coefficient on the network is positive, there is no case in which a country would want to delete a pair of links but not each link individually. While theoretically, it is possible to construct a network in which a country would want to add two links but not each link independently, it appears that does not characterize our example because the analogous moment provides no identifying power.

We next consider several robustness checks. In Table 5, the first panel again presents the baseline result for betweenness centrality using unweighted network variables. The second panel adds four explanatory variables: the log of distance between the country pairs, the log of the

Explanatory Variables	Base Line		Deletion		Addition		Min Moment	
	Low	Hi	Low	Hi	Low	Hi	Low	Hi
In(Gravity Score)	1	1	1	1	1	1	1	1
Constant	-16.158	-6.158	-100	100	-100	100	-16.158	-6.158
Between Cent.	0.185	29.873	-100	100	-100	100	0.185	29.873

Notes: The first panel is based on five moment equations: Linked pairs, unlinked pairs, the min moments, linked triplets and unlinked triplets. The following three panels include one of the latter three moment sets. Different number of observations apply to each moment. The instrument vector consists of explanatory variables, except for In(gravity score). The table reports the highest and lowest value within the confidence set for each variable. We do not search past 100, and such entries can be regarded as unbounded.

Table 4
Results for different combinations of moments.

product of GDPs, a dummy for having a common language, and a dummy for having a free trade agreement. Three of the four are not significantly different from zero. The product of GDPs is bounded below zero. While this result might be surprising on its own, recall that all four of these variables are predictors in the gravity score. Thus, the negative sign on the log of GDPs could imply that the total effect is positive but that the variable has a stronger effect on the gravity score than on value of an agreement. Betweenness centrality is still bounded in the positive range. Similarly, the fact that distance is not significantly different than zero does not imply that distance does not matter, only that it is well accounted for via the gravity score. Overall, we conclude that including these explanatory variables does not strongly affect the results, and that including them in the gravity score is an effective way to include explanatory variables in our model. We find similar results for other combinations of explanatory variables.

In the third panel of Table 5, we use the four variables as instruments but not explanatory variables. Betweenness centrality is again bounded above zero, and the range of the confidence interval is reduced slightly from the baseline result. In the fourth panel, we drop the network variable from the list of instruments. Thus, this specification treats the network variable as an endogenous variable. We include the measures of distance, joint GDP, language and trade agreement in the instrument vector. Thus, we instrument for the network variable with these four variables. The result is similar to the third panel. While the confidence interval is expanded, the

Explanatory Variables	Base Line		Additional Explanatory Vars		Additional Instruments		Without Cent. Instrument	
	Low	Hi	Low	Hi	Low	Hi	Low	Hi
In(Gravity Score)	1	1	1	1	1	1	1	1
Constant	-16.158	-6.158	-10.856	-7.517	-14.508	-6.285	-15.652	-6.258
Between Cent.	0.185	29.873	4.335	13.007	0.341	24.813	0.308	28.296
In(distance)			-0.79	0.128				
In(GDP1*GDP2)			-0.458	-0.185				
common language			-8.911	21.186				
trade agreement			-7.317	9.636				

Notes: Each panel is based on five moment sets: Linked pairs, unlinked pairs, the min moments, linked triplets and unlinked triplets. Different number of observations apply to each moment. In the first two panels, the instrument vector consists of explanatory variables, except for In(gravity score). Panel 2 includes 4 additional variables as explanatory variables and instruments. Panel 3 includes them as instruments. Panel 4 includes them as instruments but drops the network variable as an instrument. The table reports the highest and lowest value within the confidence set for each variable. We do not search past 100, and such entries can be regarded as unbounded.

Table 5
Different variables in instrument vector and explanatory vector

change is not large, and the parameter on betweenness centrality remains significantly greater than zero.

As a further robustness check, we also consider the case of bilateral transfers, or side payments, between countries. In this case, pairs of countries form a link if the sum of their payoffs from the link is positive, and do not form a link if the sum of payoffs is negative. We do not allow third-party payments, in which a country can pay a pair of countries to form a link or not.²³ We do not impose the min moments in this case, because the minimum across a country's payoffs from its links may well be negative in this case (not counting the unobserved side-payment). One might be concerned, because Table 4 shows that the min moment provides important identifying power in the case without bilateral transfers. However, remarkably, the first panel of Table 6 shows that the result that betweenness centrality is positive is robust to the issue of side payments. The second and third panels explore dropping one of the moment sets from stability refinement, and shows that they provide important identifying power in this case. Allowing countries to delete pairs of countries restricts the coefficient on the moment inequality to be positive.

²³Such a model would lead countries to the social optimum. We can consider a model. The inequality would compare whether total payments to all countries with a link were greater than total payments without the link. We have not pursued this model.

Explanatory Variables	Base Line*		Deletion		Addition	
	Low	Hi	Low	Hi	Low	Hi
ln(Gravity Score)	1	1	1	1	1	1
Constant	-20.693	100	-20.701	100	-100	100
Between Cent.	6.135	13.303	6.14	13.308	-100	100

Notes: This model allows for bilateral transfers between countries. The first panel is based on four moment sets: Linked pairs, unlinked pairs, linked triplets, and unlinked triplets. The second panel drops the moment based on unlinked triplets and the third drops the one based on linked triplets. Different number of observations apply to each moment. The instrument vector consists of explanatory variables, except for ln(gravity score). The table reports the highest and lowest value within the confidence set for each variable. We do not search past 100, and such entries can be regarded as unbounded.

Table 6
Results when countries are allowed to make bilateral side-payments

8 Efficiency of the ASA Network

The results above imply that the form 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 address two main issues. First, what is our measure of social welfare? Second, since our estimates are set identified, how do we determine which parameters to use in the evaluation?

To address the first question, we define social welfare globally as the sum across all countries.

$$S(\Lambda) = \sum_{i=1}^n \Pi_i(\Lambda) = \sum_{i=1}^n \psi_i(\Lambda, \alpha) + \sum_{k \in s_i(\Lambda)} x_{ik} \beta + \varepsilon_{ik} \quad (10)$$

This welfare weighs each country equally, which is an assumption given that we do not have a way of turning utility into dollars . etc.

With this welfare function in hand we can ask: What set of links maximizes total joint surplus? In answering this question we must account for two types of externalities. The first arises within potential links – countries do not account for their partner’s benefit from forming a link. So if the distribution of benefits from link formation are sufficiently asymmetric, then it is possible

that only one of the two countries receives a positive pay-off net of the costs of negotiating and maintaining the agreement. This suggests that some agreements that are potentially beneficial in an aggregate bilateral sense, nevertheless do not get signed – leading to too few agreements. Second, countries do not account for the effect on others through network structure. Based on the results above, countries are particularly interested in forming agreements that tend to make them hubs – put them on paths that shorten links between otherwise unlinked countries. Since this is a negative externality there is a tendency for networks to be over-connected in this case (i.e. forming a link in an attempt to make themselves more central can reduce the payoff to some other country that was previously a hub, and this negative impact is not included in country level payoffs.).

To provide insight into the role of each of these externalities we proceed sequentially. First, we evaluate the net benefit within a bilateral link starting from the observed set of agreements. If the total bilateral benefit is positive we turn on that link. While this does not generate an efficient outcome, it does provide a measure of agreements that are not formed due to an asymmetric distribution of benefits. From this point we then allow the network parameter to play a role. For every candidate set of agreements we evaluate equation (10) and iterate through links one at a time, flipping it on or off to see if it improves the outcome. We continue to iterate until there is no improvement from flipping any link. Starting from many different points and many different orderings we have found leads to the same result.

In selecting the parameter vector to be used to evaluate equation (10) 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 arg min for betweenness from column 1 of Table 3. The appropriate choice for the fixed cost of negotiating and maintaining an agreement is determined by following the implications of the parsimonious model represented by columns (1) and (2) of Table 3. In particular, this model implies that within an agreement the distribution of benefits is symmetric except for the network term. Based on this structure we select the fixed cost to be just sufficient to ensure that no additional agreements would be signed in the absence of the network effect. This benchmark reflects a setting where the only motivation to an sign agreement is based on bilateral factors (i.e.

$\alpha = 0$). Ultimately, we use the parameter vector $(1, -10, .81)$ to perform the welfare analysis.

Using these parameters we are able to evaluate the efficiency of the observed network structure. One very simple metric is based on the difference in the total number of agreements between the observed and the efficient outcomes. We calculate that the efficient configuration would require 41 fewer agreements. The reduction is not surprising since the incentive to become a hub is associated with a negative externality. However, focusing on the change in the aggregate number of agreements masks a relatively large change in the composition of agreements – who should be linked to who.

Examining compositional change gives a better sense of the role of network structure since changes in one agreement can change the incentives for a social planner to add or delete links between countries in other parts of the network. We find that the difference in the set of agreements between the observed and efficient outcomes is substantial. In particular, we find that 41 countries increase their number of links while 14 countries decrease their number of links. The big increases occur for large countries with relatively few links – Saudi Arabia, Iran, Algeria, Brazil. In contrast, big decreases are small countries with relatively many links – Singapore, Bahrain, Oman, Switzerland.

9 Conclusion

To be completed.

References

- Agarwal, N. (2013). An empirical model of the medical match. Unpublished Manuscript, MIT.
- Alatas, V., Banerjee, A., Chandrasekhar, A. G., Hanna, R., & Olken, B. A. (2014). Network structure and the aggregation of information: Theory and evidence from indonesia. Unpublished manuscript, MIT.
- Andrews, D. W. K. & Barwick, P. J. (2012). Inference for parameters defined by moment inequalities: A recommended moment selection procedure. *Econometrica*, 80(6), 2805 – 2826.
- Badev, A. (2013). Discrete games in endogenous network: Theory and policy. Unpublished manuscript, Board of Governors of the Federal Reserve.
- Baier, S. L. & Bergstrand, J. H. (2004). Economic determinants of free trade agreements. *Journal of International Economics*, 64, 29–63.
- Bajari, P., Benkard, L., & Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75, 1331–1370.
- Boucher, V. & Mourifié, I. (2015). My friend far far away: Asymptotic properties of pairwise stable networks. Unpublished manuscript, University of Toronto.
- Cameron, A. C. & Miller, D. L. (2014). Robust inference for dyadic data. Unpublished manuscript, University of California - Davis.
- Chamberlain, G. (1980). Analysis of covariance with qualitative data. *Review of Economic Studies*, 47, 225–238.
- Chandrasekhar, A. G. & Jackson, M. O. (2014). Tractable and consistent random graph models. NBER Working Paper #20276.
- Charbanneau, K. B. (2014). Multiple fixed effects in binary response panel data models. Bank of Canada Working Paper 2014-17.
- Chen, M. X. & Joshi, S. (2010). Third-country effects on the formation of free trade agreements. *Journal of International Economics*, 82(2), 238–248.
- Ciliberto, F. & Tamer, E. (2009). Market structure and multiple equilibria in airline markets. *Econometrica*, 77, 1791–1828.
- Cristea, A., Hummels, D., & Roberson, B. (2012). Estimating the gains from liberalizing services trade: The case of passenger aviation. Unpublished Manuscript, Univeristy of Oregon.
- Currarini, S., Jackson, M. O., & Pin, P. (2009). An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, 77, 1003–1045.
- de Paula, A. (2016). Econometrics of network models. CEMMAP Working Paper CWP06/16.
- de Paula, A., Richards-Shubik, S., & Tamer, E. (2014). Identification of preferences in network formation games. Unpublished manuscript, Harvard University.
- Dickstein, M. J. & Morales, E. (2013). Accounting for expectational and structural error in binary choice problems: A moment inequality approach. Unpublished manuscript, Princeton University.
- Dresner, M. (2008). An analysis of U.S. international air transport agreements: From liberal bilaterals to the U.S.-EU multilateral accord. *Journal of Transport Research*, 10, 115–127.

- Dutta, P. & Mutuswami, S. (1997). Stable networks. *Journal of Economic Theory*, 76, 322–344.
- Egger, P., Larch, M., Staub, K., & Winkelmann, R. (2011). The trade effects of endogenous preferential trade agreements. *American Economic Journal: Economic Policy*, 3, 113–143.
- Egger, P. H. & Staub, K. E. (2014). GLM estimation of trade gravity models with fixed effects. forthcoming, *Empirical Economics*.
- 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.
- Fershtman, C. & Gandal, N. (2011). Direct and indirect knowledge spillovers: The ‘social network’ of open-source projects. *RAND Journal of Economics*, 42, 70 – 91.
- Fong, K. & Lee, R. S. (2013). Markov-perfect network formation: An applied framework for bilateral oligopoly and bargaining in buyer-seller networks. Unpublished Manuscript, Stern School of Business, New York University.
- Fox, J. (2010). Estimating matching games with transfers. Unpublished Manuscript, University of Michigan.
- Furusawa, T. & Konishi, H. (2007). Free trade networks. *Journal of International Economics*, 72, 310–335.
- Goldsmith-Pinkham, P. & Imbens, G. (2013). Social networks and the identification of peer effects. *Journal of Business & Economic Statistics*, 31, 253–264.
- Goyal, S. (2007). *Connections: An Introduction to the Economics of Networks*. Princeton University Press.
- Goyal, S. & Joshi, S. (2006). Bilateralism and free trade. *International Economic Review*, 47(3), 749 – 778.
- Graham, B. S. (2014). An empirical model of network formation: detecting homophily when agents are heterogeneous. Unpublished manuscript, University of California - Berkeley.
- Graham, B. S. (2015). Methods of identification in social networks. *Annual Review of Economics*, 7, 465 – 485.
- Head, K., Mayer, T., & Ries, J. (2010). The erosion of colonial trade linkages after independence. *Journal of International Economics*, 81, 1–14.
- Head, K. & Meyer, T. (2013). Gravity equations: Workhorse, toolkit and cookbook. CEPR Discussion Paper, #9322.
- Ho, K. (2009). Insurer-provider networks in the medical care market. *American Economic Review*, 99, 393–430.
- Hsieh, C.-S. & Lee, L.-F. (2012). A structural modeling approach for network formation and social interactions - with applications to students’ friendship choices and selectivity on activities. Unpublished Manuscript, Ohio State University.
- Imbens, G. & Kalyanaraman, K. (2010). An empirical model for strategic network formation. Unpublished Manuscript.
- Intervistas (2015). The economic impacts of air service liberalization. U.S. Department of State, <http://www.state.gov/e/eb/tra/ata/index.htm>.
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.
- Jackson, M. O. & Wolinsky, A. (1996). A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1), 44 – 74.

- Jochmans, K. (2015). Two-way models for gravity. Unpublished manuscript, Sciences Po.
- Kolaczyk, E. (2009). *Statistical Analysis of Network Data*. Springer Series in Statistics. Springer.
- Leung, M. (2015). Two-step estimation of network-formation models with incomplete information. Forthcoming, *Journal of Econometrics*.
- Limão, N. (2016). Preferential trade agreements. *Handbook of Commercial Policy*, 1, 279–367.
- Mele, A. (2013). A structural model of segregation in social networks. Unpublished Manuscript, Johns Hopkins University.
- Menzel, K. (2015). Inference for games with many players. Forthcoming, *Review of Economic Studies*.
- Micco, A. & Serebrisky, T. (2006). Competition regimes and air transport costs: The effects of Open Skies Agreements. *Journal of International Economics*, 70, 25 – 51.
- Miyachi, Y. (2014). Structural estimation of a pairwise stable network with nonnegative externality. Unpublished manuscript, MIT.
- Myerson, R. (1977). Graphs and cooperation in games. *Mathematics of Operations Research*, 2, 225–229.
- Odoni, A. (2009). The international institutional and regulatory environment. In P. Belobaba, A. Odoni, & C. Barnhart (Eds.), *The Global Airline Industry*. Wiley.
- Pakes, A., Porter, J., Ho, K., & Ishii, J. (2015). Moment inequalities and their application. *Econometrica*, 83, 315–334.
- Prell, C. (2011). *Social Network Analysis: History, Theory and Methodology*. SAGE Publications.
- Sacerdote, B. (2014). Experimental and quasi-experimental analysis of peer effects: two steps forward? *Annual Review of Economics*, 6, 253–272.
- Santos Silva, J. & Tenreyro, S. (2006). The log of gravity. *Review of Economics and Statistics*, 88, 641–658.
- Sheng, S. (2012). Identification and estimation of network formation games. Unpublished manuscript, University of Southern California.
- Sorensen, M. (2007). How smart is smart money: A two-sided matching model of venture capital. *Journal of Finance*, 62, 2725–2762.
- Uetake, K. (2014). Estimating a model of strategic network formation and its effects on performance: An application to the U.S. venture capital markets. Unpublished manuscript, Yale School of Management.

Appendices

Appendix A: Calculation of Weighted Network Statistics

To be completed.

Appendix B: Computational Details of Estimation Algorithm

Our model generates J moments $m_j(\theta)$, $j = 1, \dots, J$. For instance, based on Equation 7, we have a moment j defined as:

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

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:

$$m_j(\theta) = -\frac{1}{2n^{na}} \sum_{\{i,j\} \in \mathcal{I}(\lambda_{ij}=0)} \min_{l \in \{i,j\}} \left\{ \alpha (\psi_l(\Lambda \cup \{i,j\}) - \psi_l(\Lambda - \{i,j\})) + x_{ij}\beta \right\} z_{ij}.$$

For a given parameter vector θ , we compute the $J \times 1$ vector of moments $m(\theta)$ and the covariance matrix of the moments Σ . We then compute $\tilde{m}(\theta) = \Sigma^{-1/2}m(\theta)$. That is, we studentize the moments with the Cholesky decomposition of the covariance matrix. We define our objective function as:

$$S(\theta) = \sum_{j=1}^J n_j (\min(\tilde{m}_j(\theta), 0))^2.$$

Andrews & Barwick (2012) term a similar objective function where the moments are studentized only by the diagonal of the covariance matrix to be the objective function from Pakes et al. (2015). In order to determine if a parameter vector is in the confidence interval, we must first determine the degrees of freedom. The degrees of freedom counts the number of binding constraints at a particular parameter value:

$$df(\theta) = \# \left[j \in 1, \dots, J : \sqrt{n_j} \tilde{m}_j(\theta) < \kappa. \right]$$

where $\#$ refers to the cardinality function, and κ is a cutoff value. Finally we accept a vector of parameters into the confidence interval if:

$$CS = \{ \theta : \tilde{\chi}^2(S(\theta), df(\theta)) \leq 1 - \alpha \}.$$

where α is a confidence level. Also, $\tilde{\chi}^2$ is sum of squared half-normals.

In order to find the end points of the confidence set that we report in tables such as Table 3, we follow the following algorithm:

- Step 1: Find a value θ_0 such that $S(\theta_0) = 0$.
 - Gradient search works well here.
- Step 2: For each element of θ , find the maximum and minimum value for which there is a parameter vector in the confidence interval. For instance, to find the maximum value of the i 'th element of θ in the confidence interval, we solve:

$$\max_{\theta} \theta_i \text{ such that } \tilde{\chi}^2(S(\theta), df(\theta)) \leq 1 - \alpha.$$

- We use penalized simplex search here, using θ_0 as a starting value.
- If θ_i hits ± 100 , we call θ_i unbounded in that direction.

In practice, we use $\kappa = 2.35$. That is suggested as in Andrews & Barwick (2012) for the normal approximation, although other parts of the paper suggest cutoffs as low as 1.5. Higher numbers create more conservative confidence intervals, so we use 2.35. We use $\alpha = 0.05$, so we construct 95% confidence intervals.

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

In practice, we increment trade flows by 1 in order to address taking log of zero. Some previous papers address zeros in the dependent variable of a gravity equation in different ways (for example Santos Silva & Tenreyo, 2006). Because we analyze relatively large countries, the number of zero in trade are very small. More than half of them involve Israel and a Muslim-majority country, which is an issue presumably outside of the framework of any gravity model. As such, we include dummy for cases in which Israel trades with a Muslim-majority country.

Results appear in Table 7. 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. It might seem more natural to exponent the predicted values first and then add them and then perhaps take the log, but the exponent of a predicted value is not equal to the prediction of the log variable, and furthermore, the gravity score enters in an essentially reduced-form way, which may address any misspecification on our part.

	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 (+1). Regression includes exporter and importer fixed effects.

Table 7
Results from gravity equation estimation