

ECONOMIC
RESEARCH
FORUM



منتدى
البحوث
الاقتصادية

2018

working paper series

PRICE SYNCHRONICITY, INTER-FIRM NETWORKS,
AND BUSINESS GROUPS IN THE MIDDLE EAST
AND NORTH AFRICA

Michael Siemon

Working Paper No. 1267

PRICE SYNCHRONICITY, INTER-FIRM NETWORKS, AND BUSINESS GROUPS IN THE MIDDLE EAST AND NORTH AFRICA

Michael Siemon¹

Working Paper No. 1267

December 2018

Send correspondence to:
Michael Siemon
Cornell University
mcs296@cornell.edu

¹Department of Sociology, Cornell University.

First published in 2018 by
The Economic Research Forum (ERF)
21 Al-Sad Al-Aaly Street
Dokki, Giza
Egypt
www.erf.org.eg

Copyright © The Economic Research Forum, 2018

All rights reserved. No part of this publication may be reproduced in any form or by any electronic or mechanical means, including information storage and retrieval systems, without permission in writing from the publisher.

The findings, interpretations and conclusions expressed in this publication are entirely those of the author(s) and should not be attributed to the Economic Research Forum, members of its Board of Trustees, or its donors.

Abstract

Business groups are an essential part of the political economy of almost all capitalist countries. Although they have been intensely studied in regions like Latin America and East Asia, the study of business groups in the Middle East and North Africa is relatively less developed. This study presents evidence for the value-relevance, which is measured in terms of over-time correlations of stock returns, of family business groups, government ownership, and other inter-firm relationships among 1185 publicly traded firms in 11 countries in the Middle East and North Africa. Due to the difficulty in obtaining direct observations of business group membership, business groups are inferred with methods from network analysis. More specifically, I use a community-detection algorithm to look for clusters of different types of relationships. Next, I apply a Bayesian multilevel model to estimate the associations between group comembership (as well as other relationships), and pairwise stock returns correlations. Seven exchanges in the sample show evidence in favor of the value-relevance of inferred family business groups while six show additional correlations due to government ownership beyond that associated with co-ownership more generally.

Keywords: Networks, Community Detection, Business Groups, Stock Markets, Synchronicity, Ownership, Corporate Governance, Middle East and North Africa, Bayesian Methods, Multilevel Models

JEL Classifications: C11, D85, G32, L14, N25, O16

ملخص

تشكل مجموعات الأعمال جزءاً أساسياً من الاقتصاد السياسي في جميع الدول الرأسمالية تقريباً. على الرغم من الدراسة المكثفة لمجموعات الأعمال في مناطق مثل أمريكا اللاتينية وشرق آسيا ، إلا أن دراسة هذه المجموعات في الشرق الأوسط وشمال أفريقيا هي أقل تطوراً نسبياً. تقدم هذه الورقة دليلاً على العلاقة بين القيمة والارتباط والتي يتم قياسها من حيث الارتباط المتزايد لعوائد المخزون مع مرور الوقت، ومجموعات الشركات العائلية ، وملكية الحكومة ، والعلاقات بين شركات أخرى يبلغ عددها 1185 شركة متداولة علناً في 11 دولة في الشرق الأوسط. وشمال أفريقيا. نظراً لصعوبة الحصول على الملاحظات المباشرة لعضوية مجموعات الأعمال ، يتم استنتاجها بأساليب تعتمد على تحليل الشبكة. وبشكل أكثر تحديداً ، أنا أستخدم لوغاريتم اكتشاف المجتمع لكي أبحث عن مجموعات من أنواع مختلفة من العلاقات. بعد ذلك ، أقوم بتطبيق نموذج بايزي المتعدد المستويات لتقدير الارتباطات بين العضوية المشتركة في المجموعات (بالإضافة إلى العلاقات الأخرى) ، وإقوم بالتوفيق بين الروابط المرتبطة بالعوائد. تبين سبعة تبادلات في العينة دليلاً لصالح ارتباط القيمة بالقرابة لمجموعات الشركات العائلية المستتبطة ، في حين تظهر ستة تبادلات إضافية بسبب ملكية الحكومة ما وراء ذلك من ارتباط بالملكية المشتركة بشكل عام.

1. Introduction

Business Groups are a fundamental part of the social organization of economic activity throughout the world (Khanna and Yafeh 2007; Granovetter 2010; Colpan et al. 2010; Carney et al. 2011; Colli and Colpan 2016). Despite the growth of research on business groups, few studies have investigated their role in the Middle East and North Africa (MENA). Lack of available data is one reason for this neglect². This paper helps to address this lack and contributes to the literature on business groups more generally by presenting a network-based community detection procedure for inferring group membership and providing evidence that these latent groups influence the comovement of stock prices among public firms in the Middle East and North Africa.

2. Literature Review and Hypotheses

The goal of this paper is to establish that network communities correspond to socially meaningful business groups by showing that comembership in an inferred group is associated with higher pairwise synchronicity in firms' stock prices. This argument relies on two assumptions, that business groups are associated with observable clusters of network ties and that firms within the same business group tend to have highly correlated stock price returns. This section describes prior work justifying these assumptions. It begins by offering a conceptual definition business groups and presenting empirical work that helps to establish a link between the formal notion of network communities and the socially constructed phenomenon of a business group. It then describes a variety of research supporting the expectation that group membership influences investors' perceptions of firm value with special attention to evidence for the value-relevance of such firm-specific information on MENA stock exchanges. It concludes by offering a set of hypotheses regarding the relationship between pairwise price comovement, directly observed connections, and inferred group membership.

2.1. Business Groups as Network Communities

Since the 1970s the business group literatures has developed from a collection of case-studies into a thriving area of research that spans the disciplines of sociology, management, finance, and economics. For example, Strachan's (1976) early study of *grupos* in Nicaragua continues to influence current work due to its rich ethnographic account and can be seen as a precursor to the numerous investigations of Korean and Japanese business groups prompted by the rapid economic growth of East-Asian economies during the Cold War (Orrú et al. 1989; Steers et al. 1989; Biggart 1990; Kim 1991). This research helped foster a growing awareness of business groups as a concept positioned somewhere between the hierarchy of a firm and the decentralization of a market (Granovetter 1995), and subsequent work has tended to follow a common template of purchasing or collecting data on the business groups in a particular country and analyzing their impact on outcomes like firm performance or innovation.

² Many authors of business group studies in other countries note the extensive field work necessary to gather their data. Some countries, however, benefit from well publicized data sets.

Overall, however, the findings of this literature have been mixed, and two major summaries both conclude that rather than having a uniform and measurable impact on firm performance, business groups appear to operate differently in different contexts (Khanna and Yafeh 2007; Carney et al. 2011). While frustrating from a policy perspective, this contingency is an inherent feature of business groups. Unlike the limited liability firm or the nation-state, business groups are not defined with reference to global standards or norms. Available evidence suggests rather that they tend to emerge endogenously as a response to local conditions or as a result of economic elites moving to consolidate their privileged positions (Colpan et al. 2010). As a result, determining which companies do or do not constitute a business group can be a significant challenge, and while single-country studies are free to rely on local criteria, theoretical or comparative treatments must identify areas of common ground.

Shared definitions tend to emphasize three features. First, it is essential that the firms within a group be legally independent. Large conglomerates are thus excluded on the grounds that they too much resemble a single large firm (Colpan et al. 2010). Second, these independent firms are bound by some combination of formal and informal ties. Ownership is thus seen as neither necessary nor sufficient for group membership (Granovetter 1995; Khanna and Yafeh 2007; Colpan et al. 2010). This is despite a narrower conception of business groups in some economic studies that stems from a substantive concern with the ability of pyramidal ownership structures to expropriate funds from minority shareholders (Bertrand et al. 2002; Bae et al. 2002; Morck et al. 2005). Third, these ties should persist over time. This rules out temporary alliances or purely transactional forms affiliation that may easily arise or dissolve with changing circumstances.

Beyond these three points studies tend to emphasize other traits depending on their area of interest. For example, Guillen (2000, p.362) includes the criterion that groups “are active in a wide variety of industries”, reflecting his concern with diversification. As Granovetter (2010) argues in his sociological account of the business group literature, many other potential defining attributes of business groups in fact vary considerably across contexts, and this includes key areas such as the extent of centralized authority, dominance by financial institutions, and relationships with the state.

In contrast with this prior work on business groups, this paper does not begin with direct observations of group membership, and a clear working definition of what can plausibly be defined as a business group is thus all the more essential. The Method section describes the basis for using a community detection algorithm developed by Mucha et al. (2010) to search for sets of firms that satisfy the definition presented above. Empirical support for this approach to detecting business groups is provided by Khanna and Rivkin (2006) in their study of the predictors of group co-membership, but given the lack of direct observations, this part of the analysis should be considered exploratory. The primary challenge of this approach is that community detection methods necessarily assign each firm to at least one cluster, regardless of that community’s underlying significance. To mitigate this concern, I focus on family-dominated business groups since the presence multiple individuals with similar names in a single community indicates a plausible underlying basis of solidarity necessary for the network community to function as a true group.

2.2. Business Groups and Pairwise Price Synchronicity

The idea that firms within the same business group have correlated returns is based on a variety of empirical evidence. As case studies of Korean chaebol and Japanese keiretsu demonstrate, business groups can be more or less centralized, but both cases entail coordination of important firm decisions at the group level (Lincoln et al. 1992; Hamilton and Biggart 1988). This type of shared control is also associated with the pooling of resources (Guillen 2000). Although business groups do not appear to have a universally positive or negative effect on firm performance (Khanna and Yafeh 2007), mechanisms like internal credit markets have been shown to reduce the variance of firm outcomes both cross-sectionally and over time (Lincoln et al. 1996; Almeida et al. 2015). A direct test of the influence of group comembership on pairwise synchronicity is provided by Khanna and Thomas (2009) in their study of inter-firm relationships in Chile. In their sample of 187 listed firms from 1996, they find that group comembership is associated with an increase of between 0.057 and 0.094 in the correlation of returns, depending on the other variables included in the model.

Motivation for linking business groups and price similarity also comes from the literatures on corporate governance and the economic value of political relationships. Interest in business groups was at least partly motivated by the extraordinary success of East Asian economies during the 1980s, but the 1998 Asian financial crisis helped to focus attention on the costs as well as the benefits of these types of economic arrangements (Rajan and Zingales 1998). Strong informal networks can facilitate trust, flexibility, and the sharing of resources, but at the same time can carry connotations of corruption, obligation, and opacity (Granovetter 1985; Uzzi 1997; Dieleman and Sachs 2008). In terms of stock prices, reliance on informal relations is at odds with the logic of a public market for corporate ownership, since only a small subset of potential investors are likely to participate in these networks.

The corporate governance literature's concern with conflicts between the *de facto* and *de jure* control of organizations is directly relevant to this problem. In the US context, Ferreira and Laux (2007) find that firms with fewer protections against takeover bids by outside investors are less synchronous. The threat of a hostile takeover, however, is not a concern for listed companies in most parts of the world, and studies of corporate governance in emerging markets instead focus on potential conflicts between controlling and minority shareholders (La Porta et al. 2008). One source of this friction are so-called ownership "pyramids", which amplify an ultimate owner's control over firms at the bottom of an ownership chain by exploiting the fact that effective control can be achieved with substantially less than a full majority of shares. In their study of French listed firms Boubaker et al. (2014) use this disparity as a measure of the divergence between majority and minority shareholder interests and show that it is associated with less firm-specific stock-price variation. Gul et al. (2010) report a similar finding for China, where they identify a quadratic association between ownership concentration and the synchronicity between a firm's price and market-wide trends. These two studies reinforce the idea that if a top shareholder's effective control is proportional to the share of the firm's profits to which their ownership entitles them, they are less likely to hide earnings and "tunnel" profits toward firms where

they have greater direct ownership using various forms of self-dealing (Bertrand et al. 2002; Bae et al. 2002).

In short, price correlations between firms in the same group should reflect the effects of coordination and resource sharing as well as their common exposure to governance issues. The latter point is especially important because there is reason to doubt the link between firm value and underlying fundamentals for business group members (Bae and Jeong 2007). Thus, even if the quality of the information disclosure for a particular firm or market is weak, the share prices for members of the same group might be correlated due to the influence of a shared distortion.

The prices of firms and business groups with political ties have also been found to move together in response to exogenous political events (Fisman 2001; Chekir and Diwan 2014). Similar studies have demonstrated how connections to government elites lower firms' costs of capital and note that such benefits do not outlast transitions to new parties or regimes (Leuz and Oberholzer-Gee 2006; Johnson and Mitton 2003; Fan et al. 2014). These studies typically measure political ties by obtaining observations of informal ties or by coding the political affiliations of board members and managers.

2.3. Value-Relevance of Firm-Specific Information in the MENA Region

The Middle East and North Africa are an attractive area for studying the associations between networks, business groups, and stock prices for three reasons: the prevalence of informal networks linking the state and the economy, the recent widespread reforms of equity markets, and the largely unexplored role of business groups in both these phenomena. The instability that has long plagued the region cannot be separated from a crisis of high population growth and low job creation, and studies have repeatedly blamed the abuses of a privileged minority of political and economic elites for this poor economic performance (Heydemann 2004; Benhassine et al. 2009). These allusions to endemic corruption in the region have tended to take the form of narratives that link government policy with the interests of their alleged cronies, but in recent work Diwan et al. (2016) has added to this literature by providing clear quantitative evidence that the Egyptian industries that experienced the entrance of politically connected businesses in the 2000s had lower employment growth relative to what would be expected in comparison to other sectors in the Egyptian economy and the same sectors in other countries.

Such evidence again highlights the potential social cost of economic systems that rely on informal relationships. Generalizing this finding to other countries in the region, however, is not easy. While institutions based on informal relationships will tend to exacerbate inequalities between connected and unconnected individuals, they nevertheless have the potential to serve as robust sources of trust and support in otherwise difficult institutional environments (Nee and Oppen 2012). Indeed, the diversity of findings on the social desirability of business groups supports the idea that the balance between these two forces is complex and not well understood (Khanna and Yafeh 2007). Thus, the negative impact of cronyism in Egypt might not apply to the experience of other countries in the region, and better understanding of the causes of economic stagnation in the the Middle East and North Africa will

require more empirical work to establish under what conditions informal relations between states and businesses do more harm than good.

Unfortunately, replicating studies like Diwan et al. (2016) requires information that is difficult to obtain. The authors first interviewed business elites about the incidence of businessmen with perceived political ties and then compared these connections against their knowledge of the effective structures of political power in Egypt, in order to eliminate ties with no real influence on government policy. Access to both types of information can of course be limited, especially in politically volatile countries with poor transparency, but the revitalization of equity markets in the Middle East allows for a less direct approach using only publicly-available data. The growth and reform of these markets has improved the disclosure of information about the ownership, governance, financial performance, and value for a greater number of listed firms, and the availability of these data has encouraged a growing literature on the relationships between firm governance, economic performance, and firm value in the region.

This work can be divided into three categories based on the outcome of interest. The first examines the predictors of specific firm governance policies, and the most relevant of these for the present study focus on the quality of information disclosures, as measured by composite indices created by analyzing the contents of annual reports. In the case of the Jordanian stock market, Haddad et al. (2015) find that in their sample of 57 annual reports from non-financial firms, companies with more concentrated government ownership have higher disclosure scores, while those with greater family ownership have lower scores. A similar study of 2007-2011 annual reports from 667 listed financial firms in the Gulf Cooperation Council countries³, however, finds that firms with board members affiliated from ruling families of the respective country have lower scores across three separate disclosure metrics (Al-Hadi et al. 2016). These contrasting results again highlight the difficulty of assessing implications of government ties.

The second category of studies involving listed firms in the region concerns the impact of firm and country characteristics on the behavior of the markets themselves. Beginning again with studies of the information content of stock prices, Abu-Ghunmi et al. (2015) analyze a sample of 116 nonfinancial firms in Jordan from 2000 to 2010 and find that firms with higher ownership concentration have less firm-specific volatility, which agrees with the results of Boubaker et al. (2014) and Gul et al. (2010). In the case of Tunisia, however, Galanti et al. (2017) find that analyst recommendations are only weakly predictive of firm value relative to similar studies in other countries, suggesting that prices reflect noise or private information, perhaps as a result of low levels of informed trading. Indeed, the informativeness of a firm's stock price is intrinsically related to the frequency of trades since each transaction incorporates new data about market participants' expectations. Hearn (2014) approaches this issue of liquidity using a sample of over 300 firms on the Moroccan, Tunisian, and Egyptian stock markets in order to analyze the association of firm- and country-level measures with the average transaction costs of trading a firm's shares both before and after the regime changes in early 2011. His

³ Kuwait, Bahrain, Qatar, the United Arab Emirates, Saudi Arabia and Oman

results support the use of the percent of days with non-zero returns as a proxy for liquidity and also suggest that firms affiliated with family business groups have marginally lower transaction costs in Morocco, Tunisia, and among larger Egyptian firms.

The final set of studies focuses on the associations between different governance structures and directly observed measures of firm performance and market value. For example, Uddin et al. (2014) find evidence that government-owned firms in the UAE have a higher return on assets but lower valuations relative to their total assets. In Oman, Rajab et al. (2015) investigate how a composite measure of the quality of governance by a firm's board of directors predicts the interest it pays on its debt and find that creditors are less willing to offer cheaper loans in response to improved governance in family-dominated firms. An alternative perspective on family governance is given by Mnasri and Ellouze (2015), who find that family firms in Tunisia are more productive, but only in less-competitive sectors of the economy.

Taken together, these studies highlight the complexities of corporate governance, family business groups, and government connections in the diverse political economies of the Middle East and North Africa. There is evidence that governance issues can arise in a variety of contexts across the region, but studies that rely only on formally disclosed information inevitably suffer from the selectivity of these data in terms of both the firms that disclose and the measures that are available. In order to generate causal evidence about the costs and benefits of informal economic relationships, more information is required that speaks directly to the mechanisms involved. For example, Chekir and Diwan (2014) supplement stock price and other publicly available data for Egyptian listed firms with "common knowledge" from stock brokers about which firms have received special assistance from the state. They use this extra information to show that these politically connected firms suffered disproportionate decreases in value due to political events like the 2011, uprising and furthermore that this value is linked to greater access to state-subsidized credit but not better utilization of the extra capital.

Nevertheless, even formalizing common knowledge can be challenging, and this study instead seeks to make fuller use of the data contained in common disclosures by converting them into a network, which is a formal representation of the connections between individuals and organizations. This approach has a long history in sociology and organizational studies (Mizruchi 1996), and the key innovation of this study is the use of community detection algorithms to search for socially meaningful clusters within these networks (Newman 2006; Mucha et al. 2010). Although the amount of research on business groups in the Arab countries of the Middle East and North Africa is small compared to other areas of the world, they have been identified in published studies of Morocco (Saadi 1989) as well as Tunisia and Egypt (Hearn 2014), and research on the impact of informal connections in economic governance would not be complete without accounting for the role of these ubiquitous and diverse institutions (Granovetter 2010). As such, two of the goals of this study are to use pairwise price synchronicity to provide evidence on how the value-relevance of inter-firm relationships varies across the Middle East and North Africa and to confirm the validity of inferred business group measures.

The basic requirement for this approach to be successful is that stock prices in the region reflect enough firm-specific information for the influence of inter-firm relationships to be observable. Although I am not aware of other studies of pairwise comovement in the MENA region, two of the studies mentioned above provide evidence that firm prices do respond to important firm-specific information. In addition to demonstrating the value of political connections in Egypt, Chekir and Diwan (2014) also provide direct evidence that investors in the Egyptian market are both aware of the connections enjoyed by listed firms and that relationships with the state induce synchronicity through shared vulnerability to political instability. The results of Abu-Ghunmi et al. (2015) are also encouraging in that they show the same decrease in firm-specific information with greater ownership concentration that has been found in countries with large and active markets like France and China (Boubaker et al. 2014; Gul et al. 2010). Less-direct evidence is available from a wider literature on the extent to which accounting-based measures of firm value predict market prices. In general, this literature finds that measures of firm earnings and book-value are moderately predictive of share prices in Kuwait (El Shamy and Kayed 2005; Al-Hares et al. 2012), Saudi Arabia (Al-Sehali and Spear 2004), Jordan (Abuzayed et al. 2009), Egypt (Ragab and Omran 2006; El-Sayed Ebaid 2011), and Morocco and Tunisia (Anandarajan and Hasan 2010). Because they rely on annual reports for their measures of earnings and book-value, they can only give a coarse indication of the relationship between firm-specific information and price. Nevertheless, the fact that earnings and book-value show a consistent relationship to annual variations in price is evidence that the value of companies in the region do depend at least somewhat on firm-specific information.

2.4. Hypotheses

In conclusion, this paper tests both the value-relevance of inter-firm relationship in the Middle East and North Africa as well as the ability of the community detection algorithms to identify meaningful business groups through the following hypotheses. First, given the results of Khanna and Thomas (2009) I expect to observe greater price comovement between firms who share a directly observable relationship.

Hypothesis 1: Firms that have a direct or indirect ownership connection or share a director or owner have greater pairwise synchronicity.

Finding support for this hypothesis would help to confirm the results of Khanna and Thomas (2009) for a wider sample of countries. Second, given the continuing economic influence of the state in the Middle East and North Africa, government ownership should be especially important in shaping the valuation of firms.

Hypothesis 2: Firms that are owned by parts of the same national government have greater pairwise synchronicity beyond that associated with shared owners in general.

Government ownership is of course not equivalent to the types of crony connections found in Chekir and Diwan (2014) and Diwan et al. (2016), but it still reflects direct state involvement. Third, compared

with government ownership, business group membership is more difficult to observe. My approach combines the network communities with information about the surnames of owners, directors, and top management as follows.

Hypothesis 3: Firms that are members of the same inferred network community and both have at least one owner, director, or top manager from that community's dominant family have greater pairwise synchronicity than that associated with their directly observed director and ownership ties.

Finally, the diversity of the political and economic environments in the region suggests that these hypotheses might apply more or less strongly for firms on different exchanges. Given the small sample of exchanges in the data, I do not offer a formal hypothesis concerning the sources of this variation.

3. Methods

3.1. Community Detection in Networks

My analysis is based on four observed networks, which are defined as sets of vertices linked by edges. The two principle networks are a bipartite network of director-firm affiliations and a directed network of ownership relationships. The former is bipartite in that it involves two distinct types of vertices, firms and directors, linked by edges representing an employment relationship, which can only exist between nodes of different types, i.e. a firm cannot serve as a director for another firm. The ownership network is directed in order to reflect the inherent asymmetry of ownership relations and allow for mutual cross-holdings, such that Firm A owning shares in Firm B is distinct from an edge in the opposite direction. These networks provide the following basic measures of the relatedness of two given firms: director interlocks, which is a binary measure of whether two firms share at least one director; direct ownership; and shared ownership, defined as the geometric mean of the total percent of shares held in each firm by all shared owners. The other two networks reflect family connections and government control. The Data section below provides more details.

Despite the variety of relational data encoded by these networks, their substantive implications are not always clear. Director-interlocks can occur for a variety of reasons and do not necessarily indicate a meaningful economic relationship between two firms (Palmer 1983; Mizruchi 1996; Haunschild and Beckman 1998; Chu and Davis 2016). A shared surname, moreover, is neither a necessary nor sufficient indicator of family ties (Khanna and Rivkin 2006). Even government ownership can have an ambiguous relationship with state control because of the widespread involvement of state pension and social security funds in the region's stock markets (Heydemann 2004). In order to better identify meaningful connections between firms, I use a multiplex community detection algorithm to identify clusters of firms that appear across one or more of the four networks (Mucha et al. 2010). This procedure relies on a generalization of network modularity, which is a clustering score indicating the extent to which a specified set of nodes are more interconnected than would be expected under a null model of randomly formed connections (Newman 2006). In the multi-network context, each network can be thought of as a slice of a larger structure in which occurrences of a vertex in different networks

are linked by user-specified weights, and the multiplex community detection algorithm adapts Newman's spectral partitioning approach to search for a high modularity partition.

Prior research on the role of network connections in structuring business groups suggests that the communities found by the algorithm might correspond to these important structures. For example, Fisman (2001) notes that the data he purchased from consultants in Indonesia to evaluate the relevance of regime instability on the valuation of politically connected groups was based on their assessment of the directors and owners of various firms. Furthermore, Khanna and Rivkin (2006) find in their analysis of business group comembership in Chile that director-interlocks and ownership are significant predictors of two firms belonging to the same group.

An important drawback to modularity-based methods, however, is that there is no guarantee that a given network will have a clearly optimal partition, and highly dissimilar partitions can have similar modularity values (Good et al. 2010). I approach this difficulty from three directions. First, I distinguish family-dominated communities from those lacking a clear basis of solidarity. Second, directly including these family as well as government connections in the community detection algorithm provides additional data about where groups might be found and should improve the reliability of the method. In other words, they give greater certainty to the algorithm if they correspond to similarly located concentrations of director-interlock and ownership ties, but are easily ignored if they do not. Finally, testing the results against independently generated but related data provides perhaps the strongest verification, short of comparing them against true group labels. Pairwise price correlations provide a such a data set.

A further complication of modularity-based community detection is their intrinsic resolution limit that inflates the score of weakly connected clusters in large networks (Fortunato and Barthelemy 2007). This occurs because as the number of nodes grows, the expected value for a given edge becomes close to zero. Thus, the modularity increase from a single edge between otherwise unconnected groups of nodes can outweigh the penalty incurred by merging all of the other unlinked nodes. The scale at which this limit occurs, however, can be adjusted using a parameter, γ which increases the penalty of grouping unconnected nodes, and it is possible to set γ to return results consistent with prior expectations of community size. Nonetheless, this approach is limited in that it imposes a uniform increase on the penalty for including unconnected nodes in the same community, which is less appropriate for analyzing the nested but interconnected networks in this analysis of multiple exchanges in the MENA region. Each public firm observation was generated within one or more exchanges, which vary considerably in terms of the practice and regulation of corporate governance (Nagy Eltony and Babiker 2005). This creates variation in the density of firm relationships that are reported in official disclosures, so rather than assume uniformity where little exists, my approach was to specifically incorporate this hierarchical structure by first identifying top-level communities that roughly correspond to exchanges, and then to treat each top-level community as a separate network and rerun the algorithm for each exchange.

3.2. Bayesian Analysis

I employ a Bayesian multilevel robust linear regression model to estimate the association between inter-firm relationships and stock-price similarity. This section first describes the basics of Bayesian statistical modeling as well as the reasons for preferring this approach to more traditional analyses that use maximum-likelihood and p-values. Next, it describes the multilevel structure of the model, which accounts for the clustering of dyads within exchanges and firms, as well as the use of a t-distribution to model prediction errors. Finally, it describes the estimation technique used to fit the model, which relies on sampling parameter values from their posterior distributions given the data and prior expectations about their likely values.

Bayesian methods can be defined as an alternative to the type of statistical analysis that has historically dominated the social sciences. This traditional approach is often referred to as “frequentist” in the Bayesian literature, reflecting its emphasis on sample size in determining the validity of inferences through mechanisms like p-values and standard errors. Hypothesis testing under this framework relies on the thought-experiment of repeatedly obtaining a similar data set according to the same data-generating process, although such replication is often impractical for observational studies. A p-value of 0.05 for a null hypothesis test thus implies that if the null hypothesis were true and we were given 100 such replications, we would expect only 5 of them to generate a sample statistic, such as a regression coefficient, larger than that returned by the observed data. Furthermore, frequentist statistical inferences typically take the form of single point-estimates of unknown population parameters that maximize the likelihood of observed data under a given model. The validity of these estimates is a direct function of sample size in the sense that sample statistics will converge to a population parameter as the size of the sample grows.

In contrast, Bayesian methods avoid the sometimes counter-intuitive idea of replicated datasets, and instead focus on the full distribution of an estimated parameter given observed data and prior beliefs. Bayesian interval representations of estimated parameters thus have the more immediate interpretation that is often often incorrectly given to frequentist confidence intervals, namely that the unknown parameter has a given probability of being within certain range (Gelman et al. 2014, p. 33). Such distributions depend on the model, observed data, and prior beliefs, and the first two are combined to form the likelihood while the latter is referred to as the prior. These two elements are combined to generate a posterior distribution for the unknown quantities of interest. This process is at the core of Bayesian inference and yields several desirable features, such as a reduced dependence on large sample sizes, greater model flexibility, and explicit representations of uncertainty. To realize these benefits, however, the analyst must specify both the structure of the model and the prior distributions of parameters. Furthermore, because they derive parameter distributions by combining the likelihood with prior information, Bayesian models can rarely be estimated by maximum-likelihood algorithms and instead rely on computationally intensive sampling routines to approximate these posterior distributions.

I chose to model the data with Bayesian methods primarily because of their greater flexibility. Because of the difficulty of deriving precise point estimates for important parameters, maximum-likelihood algorithms for multilevel models can encounter problems when dealing with non-nested structures, multiple coefficients that vary by group, and small numbers of groups (Gelman and Hill 2007, p. 345). Each of these issues applies to the current analysis.

3.3. Model

I estimate the association between relational measures and stock-price similarity using a multilevel robust linear regression model. A multilevel structure is useful in this case because it allows important coefficients to vary by exchange. Alternative approaches would be to estimate separate models for each exchange, for example by including interactions between an exchange indicator and all coefficients of interest, or else to estimate a single coefficient for all exchanges. The first approach allows for flexibility at the cost of ignoring information from other exchanges, while the second ignores the structure of the data in favor of including all available information into a single estimate. Multilevel modeling is a compromise between these two approaches, and in a Bayesian framework these types of models allow the data and prior to explicitly determine the extent to which group-specific estimates are drawn toward the posterior estimate of the higher-level parameter (Gelman and Hill 2007, p. 251). This occurs through the estimation of a covariance matrix for all coefficients that vary by exchange, which is used along with the overall mean and data to estimate the particular value for each group.

Multilevel models have a long history of being used to model dyadic data (Van Duijn et al. 1999; Zijlstra et al. 2006). Analyses of data from multiple network structures often take their observations from a set of distinct but comparable units like schools or families, and multilevel models are an effective means of incorporating heterogeneity among clusters of observations. Furthermore, the mixed structure of multilevel models, which allows for combinations of coefficients that are either constant or varying across groups, include popular network models as a special case. A fundamental challenge for regression analyses of dyadic data is that observations are not independent due to the fact that measures from two dyads that include the same node will tend to be correlated as a result of the shared influence of that node's characteristics. One strategy for addressing this issue is to decompose the error terms into node-specific effects and conditionally independent dyad-level residuals (Van Duijn et al. 2004). This approach is similar to the idea of non-nested multilevel models, which allow for each observation to be associated with multiple overlapping group indicators, each having their own intercept.

Because stock-price correlations are a continuous measure of firm similarity, they can be predicted with a linear regression model. However, as described in the next section, they are over-dispersed relative to a normal distribution, and the heavy tails of the country-specific distributions is indicative of the presence of numerous outliers. To mitigate the influence of these outliers, I fit a robust linear regression that replaces the usual assumption of a normal distribution of the model residuals with a t-distribution (Gelman et al. 2014, p. 444). This allows the model to more easily ignore extreme observations by assigning a greater probability to large residual values, at the cost of estimating an

extra degrees of freedom parameter for the residual distribution. Taking the preceding factors into account, the model has the following form:

$$(1) y_i = \beta \mathbf{x}_i^T + \gamma_j \mathbf{z}_i^T + \delta_k + \delta_l + e_i$$

where i is the index of the dyad composed of nodes k and l ; j is the index of the exchange where dyad i was observed; y_i is the Fisher-transformed stock price correlation; $\mathbf{x}_i$ is a vector containing the variables whose associations with price comovement do not vary by exchange in the model; and $\mathbf{z}_i$ contains the other variables whose coefficients are allowed to vary. I did not allow the industry coefficients to vary because it is likely that the greatest variation in their impact occurs among industries rather than among locations. This is primarily a practical decision since each additional coefficient that is allowed to vary by group significantly increases the computation time. Returning to the model, β is a vector of group-invariant coefficients and the other parameters of the model are defined as follows:

$$(2) \gamma_j \sim N(\mu_\gamma, \Sigma_\gamma)$$

$$(3) \delta_k \sim N(0, \sigma_\delta^2)$$

$$(4) \delta_l \sim N(0, \sigma_\delta^2)$$

$$(5) e_i \sim t_v(0, \sigma^2)$$

where γ_j is a vector of varying coefficients corresponding to exchange j , which includes an intercept term; μ_γ is a vector of the mean values of each coefficient in γ_j for the overall sample; Σ_γ is a covariance matrix that includes information on the variance of each coefficient among the different exchanges as well as the covariances between different coefficients; δ_k and δ_l are node-specific intercepts for firms k and l , respectively, with a variance σ_δ^2 ; e_i is the residual for dyad i ; and v and σ are the degrees of freedom and scale for the t -distribution, respectively.

Compared to basic forms of linear regression, the model defined by Equation 1 adopts a more elaborate structure in order to conform to basic assumptions about how the data were generated. This structure is represented not only by partitioning the predictors into sets of constant and varying coefficients, but also by the additional parameters in Equation 2, Equation 3, Equation 4, and Equation 5. These values are referred to as hyperparameters since their role is to define the distributions of the model parameters. For example, μ_γ can be thought of as a vector of the group-invariant estimates for the regression coefficients corresponding to each variable in $\mathbf{z}_i$, and the coefficient for any particular exchange is thus a deviation from this value. The magnitude of this deviation depends on the data as well as the hyperparameter Σ_γ .

3.4. Prior Distributions

As noted above, Bayesian inference involves estimating the posterior distributions of model parameters, which depend on both the likelihood function of the observed data under the model as well as the prior distributions of parameters, which encode a researcher's expectations about parameter values before analyzing the data. Whether due to a lack of prior research or the number of hyperparameters, however, strong prior knowledge is not always available or practical to include, and it might in any case be desirable to make inferences directly from the data at hand and with minimal influence from prior assumptions. In order to fit the above model, I rely on weakly-informative priors,

which are defined as including less information than might be available but nevertheless incorporating knowledge of basic constraints necessary to obtain reasonable parameter values (Gelman et al. 2014, p. 55). For example, correlations are constrained within the interval $[-1, 1]$, and even after performing the Fisher transformation the difference between the maximum and minimum value of the dependent variable is only 4.42. Furthermore, the difference between any two other observations is likely to be much lower, and I thus assign a normal prior distribution with a mean of zero and a standard deviation of 5 for the exchange-invariant coefficients and the means of the varying coefficients⁴. Another example of this type of weak constraint involves variance hyperparameters, which must be non-negative. More formally, the various prior distributions for the model parameters described above are as follows:

$$(6) \beta_n \sim N(0, 5^2)$$

$$(7) \mu_{\gamma m} \sim N(0, 5^2)$$

where n and m index the exchange-invariant and varying coefficients, respectively. Σ_γ is decomposed into a vector of coefficient standard deviations, σ_γ , and off-diagonal correlations Ω_γ such that:

$$(8) \Sigma_\gamma = D(\sigma_\gamma) \Omega_\gamma D(\sigma_\gamma)$$

where D is a function that produces a square matrix, $\mathbf{D}$, such that the diagonal elements $\mathbf{D}_{mm} = \sigma_{\gamma m}$ where m indexes the parameters that vary by exchange and all off-diagonal values are 0. Ω_γ and σ_γ are a matrix and vector, respectively, defined by the following prior distributions:

$$(9) \sigma_{\gamma m} \sim \text{half-}t_3(1)$$

$$(10) \Omega_\gamma \sim \text{LKJ}(1)$$

where $\text{LKJ}(\zeta)$ refers to the LKJ-Correlation prior (Lewandowski et al. 2009) and half- t denotes a centered t -distribution that is “folded” at zero. The remaining parameters are given the following weakly-informative priors:

$$(11) \sigma_\delta \sim \text{half-}t_3(1)$$

$$(12) v \sim \text{gamma}(2, 0.1)$$

$$(13) \sigma \sim \text{half-}t_1(\sigma_y)$$

where σ_y is the standard deviation of the dependent variable. With the exception of the priors for the coefficients in Equation 6 and Equation 7, all of the above priors are the defaults suggested by the software used to fit the model (see Buerkner 2017).

3.5. Model Estimation

Bayesian models are typically fit using Markov Chain Monte Carlo (MCMC) methods which take samples from a target distribution until they converge to a stable representation of the desired values. There are several variations on this approach, and this paper is indebted to recent advances in Hamiltonian Monte Carlo (HMC) that have been made widely available through the Stan statistical platform (Carpenter et al. 2016). Compared to other staple techniques like Metropolis-Hastings and Gibbs samplers, HMC reduces the random-walk behavior of successive samples by using information

⁴ Because the data consist of both binary and percent values, a more informative prior might assign a higher standard deviation to the percent variables. However, having the same prior across all coefficients leads to an important speed increase for the estimation algorithm (Buerkner 2017)

from the gradient of the log-likelihood to produce weakly correlated draws that still converge to the appropriate distribution (Gelman et al. 2014, p. 300). HMC, however, depends on parameters which must be tuned during the analysis in order to ensure efficient and accurate results. A core feature of Stan is that it uses the recently developed no-U-turn sampler (NUTS) and other methods to automate this process (Hoffman and Gelman 2014). The use of Stan, which was developed in C++, has been facilitated by packages in other languages that make use of its tools, and this study uses the Bayesian Regression Models using Stan (brms) package in R (Buerkner 2017).

4. Data

This study investigates the association between inter-firm relationships and price synchronicity using data from 1185 publicly traded firms on 12 stock exchanges in 11 countries in the Middle East and North Africa. The units of analysis are all pairs of firms listed on the same exchange. Including cross-listings, this yields a total of 69,845 observations with sufficient data.

4.1. Price Similarity


Pairwise price correlations are based Datastream's record of each firm's adjusted daily closing price in the local currency from January 1, 2010 through December 31, 2016. I chose this wide period because the low liquidity of many firms in the region often makes it impractical to estimate a correlation parameter for every pair of observations from a single year⁵. I calculated each firm's daily percent return as $r_{it} = (P_{it} - P_{it-1})/P_{it-1}$ where r_{it} is the return for firm i on day t and P_{it} and P_{it-1} are the adjusted closing prices for day t and the previous trading day, respectively. As Morck et al. (2000) note, although Datastream claims these prices are adjusted for events that yield abnormally large shifts in price, numerous changes of over 100% of a firm's value were observed, and I follow their procedure by recoding these days as missing.

A more serious issue with these data is the sparsity of non-zero returns over the seven-year period. In order to minimize the noise from infrequently traded stocks, I followed Khanna and Thomas (2009) and calculated the correlation in returns for all pairs of firms within each exchange using only those days for which both had a non-zero value. I use only observations with at least 25 non-zero returns between the two firms. This accords with the finding that non-zero returns are a convenient proxy for liquidity in the region (Hearn 2014). I chose the cutoff of 25 by inspecting the average correlation of each firm relative to its mean correlation with other firms in the exchange, and then selecting a value that would exclude all observations with abnormally high or low values⁶.

⁵ For each exchange with N publicly traded firms, this involves $N(N-1)/2$ parameters. Many firms in the region have fewer non-zero returns in a year than there are other firms on their exchange, and estimating $N-1$ parameters from fewer than N observations will necessarily produce noisy estimates.

⁶ An alternative approach is to weight each dyadic observation by the number of non-zero returns common to the two firms. Doing so, however, introduces the further problem of specifying the correct weighting scheme and also increases the time needed to sample from the resulting posterior distributions.

Figure 1: Quantile-quantile plots for Jordan and Saudi Arabia demonstrate the high proportion of extreme values in the data relative to a normal distribution



I removed the constraint that correlation values must fall within the interval $[-1, 1]$ using a Fisher transformation. This has little effect on values with an absolute value of less than 0.5, but increases the absolute value of correlations more as they get closer to -1 or 1 , which have infinite values under the transformation⁷. The resulting data have a higher proportion of extreme values than would be expected from a normal distribution. Figure 1 demonstrates this behavior for an extreme and a moderate case. These quantile-quantile plots compare the Jordanian and Saudi price correlation values to a normal distribution. In the Jordanian data, the values at bottom-left and top-right curve away from the fitted line, indicating that the quantile values corresponding to those points are lower and higher, respectively, than would be expected if the data were normally distributed. The Saudi data are closer to the fitted line and show a less pronounced but similar pattern.

4.2. Observed Networks

The data on corporate networks come from profiles for publicly traded firms provided by Mubasher.Info, a investor platform based in Saudi Arabia, Kuwait, and the UAE, that has been used in recent studies of corporate governance in the region (Hearn 2014; Hearn et al. 2017). From the point of view of network analyses, a crucial feature of Mubasher.Info is that they provide the names of owners, managers, directors, and subsidiaries in Arabic as well as Latin characters. This eliminates the difficulty of matching individual and company names that have been transliterated or translated according to different conventions, as will be described in more detail below. The exchanges are: Morocco, Tunisia, Egypt, Palestine, Jordan, Saudi Arabia, Oman, the United Arab Emirates with exchanges in both Abu Dhabi and Dubai, Qatar, Kuwait, and Iraq⁸. I supplemented this information

⁷ There are no instances in these data of firms with perfect correlations of -1 or 1 .

⁸ A previous version also included data for Bahrain. However, the small number of listed firms (21) prevented effective estimation of the exchange- and firm-level intercepts, which invalidates the conditional independence assumption described above.

with SIC codes from Bureau van Dijk's Orbis database, and data on government control from Zawya.com.

4.3. Matching

Access to the original Arabic spelling of individual and company names is essential because corporate networks must often be inferred by matching occurrences of identical or very similar names is association with different firms. Both individual and company names, however, can easily be recorded in slightly varying ways with abbreviations, word order, regional differences in orthography, and simple spelling mistakes all contributing to differences that can mask the underlying equivalence between two names. Spurious matches are another serious problem.

To deal with the first issue, I use adaptations of the name matching algorithm described in Colomer (2012). This procedure relies on a function that generates a match score based on a pairwise comparisons of the words in two names. I employ a different algorithm for individual and firm names, but both are based on identifying specific patterns of similarity, for example abbreviations or single character differences based on Levenshtein distance and weighting the score of a specific form of similarity by the inverse of the frequency of each word. In the case of firms, I compared words regardless of their position in the name, but I maintained word order for individual names. Also, I used a weighted combination of English or French versions of firm names and the Arabic one, at 35% and 65%, respectively. I based person name matches on only the Arabic version.

The key to this process was calibrating the algorithm by repeatedly assessing its ability to distinguish between likely and spurious matches. This included specific adjustments like ruling out matches between Muhammad and Ahmad, which in typical Arabic spelling differ only by a single letter, and manually coding the equivalence between appearances of names where both the given and family name have a frequency above a certain threshold. In this case, if I could not directly establish the connections between occurrences of a name, for example by locating a biography describing employment with two firms, I defaulted to leaving them unmatched. I verified the matches by manually checking the results for all public firms, and all other names that were linked with a score close to the threshold.

4.4. Variable Construction

In order to test the hypotheses described above, I use a measure derived from the community detection analysis, observations of direct ownership and interlock ties, and measures of joint ownership. I operationalize the ownership measures using the natural logarithm of one plus the raw percent value. The raw shared ownership variables are defined as the geometric mean of the percent of each firm's shares owned by the same entities or by entities controlled by the same government. If the majority of a public firm's shares are government-owned, I code that firm as a government entity for the purposes of computing the government ownership of its subsidiaries. The community measure is a proxy for comembership in a family business group and is defined as belonging to the same community and having a member of the most prominent family in that group as a director, top manager, or owner. I

also include assignment to the same two- and three-digit SIC codes in the analysis to mitigate the confounding influence of being active in the same industry.

4.5. Data Description

Table 1 presents a summary of the data at the exchange level as well as for the entire sample. It includes only dyads involving the 1185 firms with at least one dyad with 25 or more days of non-zero returns during the 2010-2016 period. The number of firms on each exchange is thus a product of its size as well as liquidity. The mean price correlations also vary substantially by exchange. The density measures refer to the number of observed connections on an exchange divided by the total number of possible dyads.

Table 1: Mean and count statistics for observations with sufficient price data for analysis are presented by exchange and the entire sample

Exchange Variable	Abu Dhabi	Dubai	Egypt	Iraq	Jordan	Kuwait	Morocco	Oman	Palestine	Qatar	Saudi Arabia	Tunisia	Entire Sample
N Firms	62	54	132	77	227	197	68	97	45	43	167	60	1229
Mean Price Correlation	0.0675	0.1736	0.3013	0.1234	0.0404	0.1240	0.0425	0.1920	0.0208	0.2360	0.4138	0.0739	0.1653
N Dyads	1760	1274	8586	2661	25253	18218	2272	3212	898	903	13821	1770	80628
N Director-Interlocks	144	60	112	59	446	291	117	185	85	249	423	63	2234
Director-Interlock Density	0.0818	0.0471	0.0130	0.0222	0.0177	0.0160	0.0515	0.0576	0.0947	0.2757	0.0306	0.0356	0.0277
N Direct Ownership Ties	7	15	26	63	157	116	32	50	56	11	51	55	639
Direct Ownership Density	0.0040	0.0118	0.0030	0.0237	0.0062	0.0064	0.0141	0.0156	0.0624	0.0122	0.0037	0.0311	0.0079
N Shared Ownership Ties	139	39	2219	1405	1724	704	335	675	185	177	1398	226	9226
Shared Ownership Density	0.0790	0.0306	0.2584	0.5280	0.0683	0.0386	0.1474	0.2101	0.2060	0.1960	0.1012	0.1277	0.1144
Mean Shared Ownership	0.0120	0.0062	0.0103	0.0148	0.0066	0.0065	0.0151	0.0267	0.0229	0.0346	0.0114	0.0175	0.0099
N Family Group Dyads	53	13	37	24	131	121	29	74	46	59	76	19	682
Family Group Density	0.0301	0.0102	0.0043	0.0090	0.0052	0.0066	0.0128	0.0230	0.0512	0.0653	0.0055	0.0107	0.0085
Mean Shared Gov. Own.	0.0804	0.0200	0.0813	0.0323	0.0084	0.0074	0.0277	0.0698	0.0077	0.0966	0.0278	0.0286	0.0262
N Same 2-Digit SIC Code	193	148	540	289	1757	1917	132	174	56	75	1076	186	6543
N Same 3-Digit SIC Code	167	113	215	186	1064	1001	54	82	43	54	779	106	3864

5. Results

This section presents the results of the analysis in three parts. First, it briefly how the inter-slice strength parameters in the community detection procedure were tuned to best fit the data and visualizes the partitions returned by the algorithm. Next, it presents the distributions of the estimated parameters from the multilevel model. Finally, it verifies the validity of the model by comparing predicted values to the observed data.


5.1. Community Detection

Detecting communities from multiple networks requires the analyst to set the strength of the connections between instances of the same node across the different slices, and this value can vary among different pairs of networks (Mucha et al. 2010). For example, in order to analyze a series of observations of the same network across time it might make sense to connect each slice with only the previous and next observation, but the analyst must still decide how strongly to couple the different

cross-sections based on their expectation of the stability of the underlying communities relative to potentially noisy variations in the observed relationships.

Applying the technique to the four cross-sectional networks described above requires two fundamental choices. First, a node will not necessarily belong to the same community across different networks, and this leads to four options for assessing group comembership based on assignment to the same community in one of the following ways: in the ownership network, in the interlock network, in both networks, or at least one of the two. Second, in any of the above cases the results will also depend on the strength of the connections among the different networks. Several scenarios appear plausible. For example, if one network is a significantly better indicator of underlying comembership than the other, then it would make sense to favor relatively weak connections and take measurements based on that network. On the other hand, if the networks are equally informative, then it might be better to enforce a strong connection between them and code communities based on membership in either network. A third possibility is that most of the network ties do not reflect group comembership; but that where groups do exist, there are concentrations of connections in both networks. In this case, the best solution would be to favor weak connections, and to measure communities based on community comembership in both networks.

Figure 2: The four networks used in the multiplex community detection analysis are shown as nodes in a multi-network structure. The number by each edge is the strength of the connection between each network. Networks with no edge are not connected



I attempt to choose among these scenarios using the data themselves. Because the goal is to correctly identify family business groups, I measured various dimensions of family control at the community level, such as overall participation and mean ownership by the dominant family. I then compared a

range of strengths of connections among the different networks across the four measurement schemes according to their ability to identify the greatest number of communities with the highest level of family control. Based on this comparison, I chose to measure family groups based on community comembership in the ownership network using the inter-slice connections shown in Figure 2. Overall, these parameters reflect the greater relevance of ownership ties for indicating group structures, although interlocks still influence the algorithm. I did not include connections between the government network and the family and interlock networks because they have no shared nodes. The government network's tie to the ownership network is stronger than those coming from the family network in order to reflect the possibility that shared surnames do not indicate family ties.

Figure 3 and Figure 4 provide a visualization of the results of the community detection algorithm at the region- and exchange-level, respectively. Both figures depict a composite network created by adding an edge to the network if two nodes are connected by either an ownership or interlock tie. Overall, the figures highlight the community structure as well as the degree of interconnections between publicly traded firms in different countries.

Figure 3: Nodes are colored and shaped by location-based community. Inter-community edges have been lightened





Figure 4: This visualization has the same layout as Figure 3. Nodes are colored and shaped by their lower-level community assignment. Inter-community edges have been lightened



5.2. Multilevel Regression

Table 2: Median and 90% posterior intervals for entire-sample coefficients

Coefficient	5%	Median	95%
Intercept	0.0704	0.1399	0.2096
Shared Director	0.0036	0.0099	0.0169
Direct Ownership	0.0010	0.0085	0.0159
Shared Ownership	0.0043	0.0080	0.0120
Same Family Group	0.0012	0.0154	0.0278
Gov. Ownership	0.0001	0.0027	0.0049
Same SIC (2-Digit)	0.0151	0.0182	0.0213
Same SIC (3-Digit)	0.0320	0.0361	0.0402


The model in Equation 1 was estimated using four simultaneous HMC chains, each with a total of 1000 sampling draws. The first 500 draws in each chain were considered a warm-up phase during which the algorithm tunes the HMC parameters and searches for a high-probability area of the parameter space. The inferences reported below are thus based on a total of 2000 draws from the second half of each chain. The use of multiple chains allows for a metric to assess the convergence of the sampling algorithm by comparing the within and between-chain variance of the sampled values. A measure of over 1.1 indicates further samples are necessary for convergence (Gelman et al. 2014, p. 285). All of the parameters from the estimated model had a value of between 1.0 and 1.01.

Table 2 presents the 0.05 percentile, median, and 0.95 percentile of the estimated distributions of the entire-sample coefficients. This includes the invariant coefficients for shared industry as well as the other variables whose coefficients vary by exchange. These entire-sample estimates broadly support the hypotheses that director interlocks, direct ownership, shared ownership, family group comembership, and government ownership are each associated with an independent increase in pairwise synchronicity. All of the corresponding parameters do not include zero in their 90% posterior intervals. The median of the distribution of the family group comembership parameter is large compared to director interlocks and comparable to shared industry, but shows much greater uncertainty.

However, these mean estimates only show one part of the picture. It remains to assess how the estimated parameters vary by exchange. Figure 5 displays the magnitude and uncertainty of the director interlock and family group coefficients for each exchange. Focusing first on director interlocks, the exchanges can be divided into three rough categories: Kuwait, Egypt, Abu Dhabi, and Morocco each have relatively large median values and 90% credible intervals that exclude zero; Jordan, Palestine, and Qatar each have somewhat smaller medians and 90% credible intervals that barely include zero; and finally Dubai, Saudi Arabia, Oman, Iraq, and Tunisia have estimated values that are either smaller,

more uncertain, or both such that zero is closer to the middle of the interval. The coefficients for family-group comembership can be split into two groups. Seven of the 12 exchanges, Dubai, Kuwait, Egypt, Saudi Arabia, Jordan, Palestine, and Oman, have median values between 0.03 and 0.02 and—with the exception of Oman—90% credible intervals that exclude zero. The other 5 exchanges have estimated values that are largely indistinguishable from zero. A challenge in interpreting these estimates is distinguishing between the noisiness of the data and the underlying mechanisms that structure price synchronicity in the region, but one or more exchanges consistently lacking clear signals that are observed for the majority of the others can be taken as evidence, not of the irrelevance of the network measure, but rather of low information content of their stock prices.


Figure 5: The symbols indicate the median coefficient estimate for each exchange, and the lines show the extent of the 90% posterior credible intervals



As shown in Figure 6, the association between the various ownership measures and price comovement also varies between the different exchanges. The coefficient for the log of shared ownership is significant across all exchanges except Tunisia. Direct ownership, on the other hand, is less consistent and shows a higher degree of uncertainty. Egypt, Kuwait, Dubai, Oman, Tunisia, Jordan, and Palestine each have 90% credible intervals that do not include zero, while Qatar, Saudi Arabia, Morocco, and Iraq have median values that are closer to null or negative. Abu Dhabi has a median that is more positive, but zero is well within its 90% confidence interval. Finally, Qatar, Saudi Arabia, Egypt,

Kuwait, and Morocco all show strong evidence of greater price comovement associated with shared government ownership. Oman, Iraq, Tunisia, Abu Dhabi, Jordan, and Palestine have median coefficients that are close to zero, and Dubai is more ambiguous, with a median close to that of Morocco, but with a wider 90% credible interval.

Figure 6: The symbols indicate the median coefficient estimate for each exchange, and the lines show the extent of the 90% posterior credible intervals



At least three factors might explain the null results for some of the exchanges mentioned above. First, the community detection method might not succeed in capturing the relevant groups in each exchange. There is evidence for this explanation in at least two cases. Morocco has a few well-documented family business groups, the largest of which is associated with the royal family, and the results of the community detection analysis only partially agree with other characterizations of this group (Saadi 1989; Oubenal 2016). Qatar appears to be uniquely unsuitable for the method of looking for shared surnames among individuals associated with firms in the same community because members of the ruling Al Thani family hold chairman or director positions in a majority of public firms. This is reflected in the abnormally high density of shared director ties reported in Table 1. A second explanation is that family business groups might have disproportionately poor data, and hence the estimated coefficient is could be inaccurate due to the noisiness of the price similarity measure. For example, there are 54 firms on the Egyptian exchange that belong to cluster containing a family that is

involved in more than one firm, but only 28 of these have over 100 days with non-zero returns. Finally, it might also be the case that the prices on a particular exchange simply do not react to group-specific events, and given the fact that it shows no clear association between price synchronicity and most relational measures, Tunisia appears to fit this category.

While the interpretation of the family group and director interlock coefficients is simply the mean difference between two groups of dyads, the coefficients for ownership measures are more complicated since the model was fit using the log of each percent measure. Table 3 displays the median predicted difference between an ownership level of 20% versus 0% for each exchange. Government ownership appears to be the most consequential for Saudi Arabia and Qatar and to a less extent for Egypt and Kuwait.

Table 3: Predicted median predicted difference between 20% and 0% ownership

Exchange	Direct Ownership	Shared Ownership	Government Ownership
Abu Dhabi	0.0228	0.0142	-0.0003
Dubai	0.0591	0.0331	0.0086
Egypt	0.0423	0.0180	0.0130
Iraq	-0.0175	0.0497	0.0039
Jordan	0.0140	0.0241	-0.0016
Kuwait	0.0707	0.0420	0.0123
Morocco	0.0040	0.0116	0.0089
Oman	0.0575	0.0145	0.0042
Palestine	0.0344	0.0283	-0.0034
Qatar	-0.0065	0.0195	0.0246
Saudi Arabia	0.0077	0.0342	0.0218
Tunisia	0.0215	0.0042	0.0012

5.3. Predictive Checks


In addition to model specification and estimation, a third pillar of Bayesian analysis is an evaluation of the fit of the estimated model (Gelman et al. 2014, p. 139). Such evaluation is typically performed through posterior predictive checks that compare the observed values of a test statistics with those derived from simulated data generated by the posterior distribution of model parameters (Lynch and Western 2004; Gelman et al. 1996). In other words, after specifying the form of a relevant statistic the researcher compares its observed value with the distribution of simulated values.

The central claim underlying this paper is that community detection analysis can infer family business groups as a latent feature of inter-firm networks in the Middle East. The multilevel model in Equation 1 was designed to test this claim by searching for heightened synchronicity between pairs of firms in the same inferred business group beyond that which we would expect from the observed connections alone. Hence, the difference between the mean synchronicity of family business group dyads and that of dyads with other observed connections can be used to assess the model's ability to capture the

impact of these inferred relationships. I calculated this difference using two subsets: dyads with family group ties and dyads with at least one direct ownership, director-interlock, or shared ownership tie but less than 4.5% government ownership⁹ and no family group comembership. Figure 7 shows the distribution of this statistic for each exchange derived from 500 sets of predicted synchronicity values generated by random draws from the model's posterior distribution relative to the value from the observed data. The observed values fall within the main body of the simulated distribution and are within the 90% credible interval for all cases except Kuwait, where the observed difference of 0.107 is greater than 95.3% of the simulated values. This indicates that the model perhaps underestimates the importance of group comembership.

I followed a similar procedure to test the model's ability to capture the additional synchronicity among dyads with significant government ownership relative to dyads with other non-family business group ties. Figure 8 shows the observed versus simulated values for this statistic. For all exchanges except Egypt, the observed statistic again falls within the 90% credible interval. In the Egyptian case the observed value is both outside the main distribution and less than over 98% of the simulations. This suggests that the model is overestimating the impact of government ties in this case.

Figure 7: For each of the twelve exchanges, the red line shows the observed difference in the mean synchronicity of family group dyads and dyads with other types of observed, non-government connections relative to the distribution of this statistic, derived from 500 sets of simulated values. The values at the top of each figure give the observed difference and its percentile relative to the simulated values.



These posterior predictive checks provide confidence that the model in Equation 1 is capable of reproducing relevant aspects of the observed data. This assessment is possible because Bayesian estimates provide the full distribution model parameters, and each sample from their joint distribution can be used to provide a simulation of the expected value of the outcome variable. These simulations can also be used to check a crucial assumption of the model, namely that the various dyadic observations are independent conditional on the inclusion of firm- and exchange-level random intercepts. Figure 9 evaluates this assumption by plotting the mean correlation between the residuals of all pairs of dyads involving each firm across 500 random draws from the posterior distribution. More specifically, I calculated the mean residual correlation using a matrix with 500 rows, each corresponding to a draw from the posterior, and n columns, where n is the number of dyads involving the firm. This yields a correlation matrix with $n * (n - 1)/2$ values that were then averaged. A large number of points not clustered at zero would indicate that the observations are not conditionally independent. Overall, 93.9% of all firms have a mean residual correlation of less than 0.025, and only 26 firms have a correlation of greater than 0.05. Concerning this latter category, the figure suggests that this failure of the conditional independence assumption is a direct result of the sparse number of observations involving these firms, rather than the specification of the model itself.

Figure 8: For each of the twelve exchanges, the red line shows the observed difference in the mean synchronicity of government-owned dyads and dyads with other types of observed, non-family connections relative to the distribution of this statistic, derived from 500 sets of simulated values. The values at the top of each figure give the observed difference and its percentile relative to the simulated values.





Figure 9: The mean residual correlation of the dyads involving each firm are plotted relative to the number of dyads involving the firm, i.e. the number of dyads with more than 25 days of non-zero returns.



6. Conclusion

This study has combined direct observations and inferred ties in order to study the role of corporate governance relationships in structuring price synchronicity between publicly traded firms in the Middle East and North Africa. In doing so, it makes two contributions to the literatures on interfirm networks, financial development, the Middle East and North Africa, and emerging markets more generally. First, it provides estimates on the value relevance of family business groups, government ownership, other inter-firm connections in 12 exchanges throughout the region. These results are substantively important as direct measures of the importance of business groups and state control, but also have a more technical interpretation in that they can serve as a proxy for the ability of prices in a market to reflect detailed information about firm relationships. Second, it confirms the validity of using community detection methods from network analysis to measure business groups for at least 7 of the 12 exchanges. These results help make the case for using this unsupervised learning technique to study the role of business groups in the Middle East and North Africa and regions where direct observations are difficult to obtain. Furthermore, they also show that pairwise price synchronicity constitutes an additional network that can be used to improve the results of such procedures.

As a whole, these findings have the potential to encourage more systematic research on the economic sociology of financial development in the region. It has shown how readily available data on public firm ownership, boards of directors, and price variations can be combined with basic information from individual names and other sources to generate a rich description of the political economies of the

region, and this synthesis has been enabled by new computational methods as well as the reform of financial markets. The main limitations of this study are related to the quality of the data and the practicalities of the model. Concerning the former, the informativeness of each pairwise synchronicity observation is a direct function of the number of nonzero returns upon which it is based, and the low liquidity of many markets in the region means that the desire to estimate the model with as many data points as possible must be tempered by the realization that adding dyads with decreasing numbers of non-zero returns will eventually contribute more noise than reliable signal. As for the model, its complexity is limited by the time needed to compute draws from the posterior distribution, but it could potentially be improved by allowing for greater heterogeneity in the role of industry in shaping the comovement of prices, for example by disaggregating the simple binary measure into more specific categories indicating that both firms are involved in key sectors like finance or resource extraction.

References

- Abu-Ghunmi, Diana, Adel Bino, and Mohammad Tayeh. 2015. "Idiosyncratic Risk and Corporate Governance: Evidence from Jordan." *Emerging Markets Finance and Trade* 51:S40–S50.
- Abuzayed, Bana, Philip Molyneux, and Nedat Al-Fayoumi. 2009. "Market Value, Book Value and Earnings: Is Bank Efficiency a Missing Link?" *Managerial Finance* 35:156–179.
- Al-Hadi, Ahmed, Grantley Taylor, and Khamis Hamed Al-Yahyaee. 2016. "Ruling Family Political Connections and Risk Reporting: Evidence from the GCC." *The International Journal of Accounting* 51:504–524.
- Al-Hares, Osama M, Naser M AbuGhazaleh, and Ayman E Haddad. 2012. "Value Relevance of Earnings, Book Value and Dividends in an Emerging Capital Market: Kuwait Evidence." *Global Finance Journal* 23:221–234.
- Al-Sehali, Mohammed and Nasser Spear. 2004. "The Decision Relevance and Timeliness of Accounting Earnings in Saudi Arabia." *The International Journal of Accounting* 39:197–217.
- Almeida, Heitor, Chang-Soo Kim, and Hwanki Brian Kim. 2015. "Internal Capital Markets in Business Groups: Evidence from the Asian Financial Crisis." *The Journal of Finance* 70:2539–2586.
- Anandarajan, Asokan and Iftekhar Hasan. 2010. "Value Relevance of Earnings: Evidence from Middle Eastern and North African Countries." *Advances in Accounting* 26:270–279.
- Bae, Kee-Hong and Seok Woo Jeong. 2007. "The Value-relevance of Earnings and Book Value, Ownership Structure, and Business Group Affiliation: Evidence From Korean Business Groups." *Journal of Business Finance and Accounting* 34:740–766.
- Bae, Kee-Hong, Jun-Koo Kang, and Jin-Mo Kim. 2002. "Tunneling or Value Added? Evidence From Mergers by Korean Business Groups." *The Journal of Finance* 57:2695–2740.
- Benhassine, Najy, Andrew Stone, Philip Keefer, Youssef Saadani Hassani, and Sameh Neguib Wahba. 2009. "From Privilege to Competition : Unlocking Private-Led Growth in the Middle East and North Africa." Technical report, World Bank, Washington, DC. MENA Development Report.
- Bertrand, Marianne, Paras Mehta, and Sendhil Mullainathan. 2002. "Ferretting Out Tunneling: An Application to Indian Business Groups." *The Quarterly Journal of Economics* 117:121–148.
- Biggart, Nicole W. 1990. "Institutionalized Patrimonialism in Korean Business." *Comparative Social Research* 12:113–133.
- Boubaker, Sabri, Hatem Mansali, and Hatem Rjiba. 2014. "Large Controlling Shareholders and Stock Price Synchronicity." *Journal of Banking and Finance* 40:80–96.
- Buerkner, Paul-Christian. 2017. "brms: An R package for Bayesian Multilevel Models Using Stan." *Journal of Statistical Software* 80:1–28.
- Carney, Michael, Eric R Gedajlovic, Pursey Heugens, Marc Van Essen, and J Hans Van Oosterhout. 2011. "Business Group Affiliation, Performance, Context, and Strategy: A Meta-Analysis." *Academy of Management Journal* 54:437–460.
- Carpenter, Bob, Andrew Gelman, Matt Hoffman, Daniel Lee, Ben Goodrich, Michael Betancourt, Michael A Brubaker, Jiqiang Guo, Peter Li, and Allen Riddell. 2016. "Stan: A Probabilistic Programming Language." *Journal of Statistical Software* 20:1–37.

- Chekir, Hamouda and Ishac Diwan. 2014. "Crony Capitalism in Egypt." *Journal of Globalization and Development* 5:177–211.
- Chu, Johan SG and Gerald F Davis. 2016. "Who Killed the Inner Circle? The Decline of the American Corporate Interlock Network." *American Journal of Sociology* 122:714–754.
- Colli, Andrea and Asli M Colpan. 2016. "Business Groups and Corporate Governance: Review, Synthesis, and Extension." *Corporate Governance: An International Review* 24:274–302.
- Colomer, Jordi. 2012. "Appendix 2: Name-Matching Algorithms." *The Small Worlds of Corporate Governance* p. 337.
- Colpan, Asli M, Takashi Hikino, and James R Lincoln. 2010. *The Oxford Handbook of Business Groups*. Oxford University Press.
- Dieleman, Marleen and Wladimir M Sachs. 2008. "Coevolution of Institutions and Corporations in Emerging Economies: How the Salim Group Morphed Into an Institution of Suharto's Crony Regime." *Journal of Management Studies* 45:1274–1300.
- Diwan, Ishac, Philip Keefer, and Marc Schiffbauer. 2016. "Pyramid Capitalism: Cronyism, Regulation, and Firm Productivity in Egypt." Technical report, Inter-American Development Bank.
- El-Sayed Ebaid, Ibrahim. 2011. "The Value Relevance of Accounting-Based Performance Measures in Emerging Economies: The Case of Egypt." *Management Research Review* 35:69–88.
- El Shamy, Mostafa A and Metwally A Kayed. 2005. "The Value Relevance of Earnings and Book Values in Equity Valuation: An International Perspective-The Case of Kuwait." *International Journal of Commerce and Management* 15:68–79.
- Fan, Joseph PH, Feng Guan, Zengquan Li, and Yong George Yang. 2014. "Relationship Networks and Earnings Informativeness: Evidence From Corruption Cases." *Journal of Business Finance and Accounting* 41:831– 866.
- Ferreira, Miguel A and Paul A Laux. 2007. "Corporate Governance, Idiosyncratic Risk, and Information Flow." *The Journal of Finance* 62:951–989.
- Fisman, Raymond. 2001. "Estimating the Value of Political Connections." *The American Economic Review* 91:1095–1102.
- Fortunato, Santo and Marc Barthelemy. 2007. "Resolution limit in community detection." *Proceedings of the National Academy of Sciences* 104:36–41.
- Galanti, Sébastien, Zahra Ben Braham, et al. 2017. "Information Efficiency on an Emerging Market: Analysts' Recommendations in Tunisia." *Economics Bulletin* 37:377–390.
- Gelman, Andrew, John B Carlin, Hal S Stern, David B Dunson, Aki Vehtari, and Donald B Rubin. 2014. *Bayesian Data Analysis*. CRC Press, 3 edition.
- Gelman, Andrew and Jennifer Hill. 2007. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. New York, NY: Cambridge University Press.
- Gelman, Andrew, Xiao-Li Meng, and Hal Stern. 1996. "Posterior Predictive Assessment of Model Fitness via Realized Discrepancies." *Statistica Sinica* pp. 733–760.
- Good, Benjamin H, Yves-Alexandre de Montjoye, and Aaron Clauset. 2010. "Performance of Modularity Maximization in Practical Contexts." *Physical Review E* 81:046106.
- Granovetter, Mark. 1985. "Economic Action and Social Structure: The Problem of Embeddedness." *American Journal of Sociology* 91:481–510.

- Granovetter, Mark. 1995. "Coase revisited: Business groups in the modern Economy." *Industrial and Corporate change* 4:93–130.
- Granovetter, Mark. 2010. "Business Groups and Social Organization." In *The Handbook of Economic Sociology*, edited by Neil J. Smelser and Richard Swedberg, pp. 429–450. Princeton University Press.
- Guillen, Mauro F. 2000. "Business Groups in Emerging Economies: A Resource-Based vCew." *Academy of Management Journal* 43:362–380.
- Gul, Ferdinand A, Jeong-Bon Kim, and Annie A Qiu. 2010. "Ownership Concentration, Foreign Shareholding, Audit Quality, and Stock Price Synchronicity: Evidence From China." *Journal of Financial Economics* 95:425–442.
- Haddad, Ayman E, Wasim K AlShattarat, Naser M AbuGhazaleh, and Haitham Nobanee. 2015. "The Impact of Ownership Structure and Family Board Domination on Voluntary Disclosure for Jordanian Listed Companies." *Eurasian Business Review* 5:203–234.
- Hamilton, Gary G and Nicole Woolsey Biggart. 1988. "Market, Culture, and Authority: A Comparative Analysis of Management and Organization in the Far East." *American Journal of Sociology* 94:S52–S94.
- Haunschild, Pamela R and Christine M Beckman. 1998. "When Do Interlocks Matter?: Alternate Sources of Information and Interlock Influence." *Administrative Science Quarterly* pp. 815–844.
- Hearn, Bruce. 2014. "The Political Institutional and Firm Governance Determinants of Liquidity: Evidence From North Africa and the Arab Spring." *Journal of International Financial Markets, Institutions and Money* 31:127–158.
- Hearn, Bruce, Lars Oxelheim, and Trond Randøy. 2017. "The institutional determinants of private equity involvement in business groups: The case of Africa." *Journal of World Business* .
- Heydemann, Steven. 2004. *Networks of Privilege in the Middle East: The Politics of Economic Reform Revisited* . Springer.
- Hoffman, Matthew D and Andrew Gelman. 2014. "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo." *Journal of Machine Learning Research* 15:1593–1623.
- Johnson, Simon and Todd Mitton. 2003. "Cronyism and Capital Controls: Evidence from Malaysia." *Journal of Financial Economics* 67:351–382.
- Khanna, Tarun and Jan W Rivkin. 2006. "Interorganizational Ties and Business Group Boundaries: Evidence From an Emerging Economy." *Organization Science* 17:333–352.
- Khanna, Tarun and Catherine Thomas. 2009. "Synchronicity and Firm Interlocks in an Emerging Market." *Journal of Financial Economics* 92:182–204.
- Khanna, Tarun and Yishay Yafeh. 2007. "Business Groups in Emerging Markets: Paragons or Parasites?" *Journal of Economic Literature* 45:331– 372.
- Kim, Eun Mee. 1991. *The industrial Organization and Growth of the Korean Chaebol: Integrating Development and Organizational Theories*. Centre of Asian Studies, University of Hong Kong.
- La Porta, Rafael, Florencio Lopez-de Silanes, and Andrei Shleifer. 2008. "The Economic Consequences of Legal Origins." *Journal of Economic Literature* 46:285–332.

- Leuz, Christian and Felix Oberholzer-Gee. 2006. "Political Relationships, Global Financing, and Corporate Transparency: Evidence from Indonesia." *Journal of Financial Economics* 81:411–439.
- Lewandowski, Daniel, Dorota Kurowicka, and Harry Joe. 2009. "Generating Random Correlation Matrices Based on Vines and Extended Onion Method." *Journal of Multivariate Analysis* 100:1989–2001.
- Lincoln, James R, Michael L Gerlach, and Christina L Ahmadjian. 1996. "Keiretsu Networks and Corporate Performance in Japan." *American Sociological Review* pp. 67–88.
- Lincoln, James R, Michael L Gerlach, and Peggy Takahashi. 1992. "Keiretsu Networks in the Japanese Economy: A Dyad Analysis of Intercorporate Ties." *American Sociological Review* pp. 561–585.
- Lynch, Scott M and Bruce Western. 2004. "Bayesian Posterior Predictive Checks for Complex Models." *Sociological Methods and Research* 32:301–335.
- Mizruchi, Mark S. 1996. "What Do Interlocks Do? An Analysis, Critique, and Assessment of Research on Interlocking Directorates." *Annual Review of Sociology* 22:271–298.
- Mnasri, Khadija and Dorra Ellouze. 2015. "Ownership Structure, Product Market Competition and Productivity: Evidence From Tunisia." *Management Decision* 53:1771–1805.
- Morck, Randall, Daniel Wolfenzon, and Bernard Yeung. 2005. "Corporate Governance, Economic Entrenchment, and Growth." *Journal of Economic Literature* 43:655–720.
- Morck, Randall, Bernard Yeung, and Wayne Yu. 2000. "The Information Content of Stock Markets: Why Do Emerging Markets Have Synchronous Stock Price Movements?" *Journal of Financial Economics* 58:215–260.
- Mucha, Peter J, Thomas Richardson, Kevin Macon, Mason A Porter, and Jukka-Pekka Onnela. 2010. "Community Structure in Time-Dependent, Multiscale, and Multiplex Networks." *Science* 328:876–878.
- Nagy Eltony, M and Mustafa Babiker. 2005. "Arab Capital Markets Development and Institutions." *Journal of Economic and Administrative Sciences* 21:42–63.
- Nee, Victor and Sonja Oppen. 2012. *Capitalism from Below: Markets and Institutional Change in China*. Harvard University Press.
- Newman, Mark EJ. 2006. "Modularity and Community Structure in Networks." *Proceedings of the National Academy of Sciences* 103:8577–8582.
- Orrú, Marco, Gary G Hamilton, and Mariko Suzuki. 1989. "Patterns of inter-firm control in Japanese business." *Organization Studies* 10:549–574.
- Oubenal, Mohammed. 2016. "Crony Interlockers and the Centrality of Banks: The Network of Moroccan Listed Companies." Technical Report 1066, Economic Research Forum Working Paper.
- Palmer, Donald. 1983. "Broken Ties: Interlocking Directorates and Intercorporate Coordination." *Administrative Science Quarterly* pp. 40–55.
- Ragab, Aiman A and Mohammad M Omran. 2006. "Accounting Information, Value Relevance, and Investors Behavior in the Egyptian Equity Market." *Review of Accounting and Finance* 5:279–297.

- Rajab, Amraha Muneer, Hafiza Aishah Hashim, and Akmalia Mohamad Ariff. 2015. "The Moderating Effect of Family Control on the Relationship Between Board of Directors Effectiveness and Cost of Debt: Evidence From Oman." *International Journal of Economics, Management and Accounting* 23:217.
- Rajan, Raghuram G and Luigi Zingales. 1998. "Which capitalism? Lessons form the East Asian Crisis." *Journal of Applied Corporate Finance* 11:40–48.
- Saadi, Mohamed Said. 1989. *Les groupes financiers au Maroc*. Rabat, Morocco: Ed. Okad.
- Steers, Richard M, Yu-gun Sin, and Gerardo R Ungson. 1989. *The Chaebol: Korea's New Industrial Might*. Harper and Row, Ballinger Division.
- Strachan, Harry W. 1976. *Family and Other Business Groups in Economic Development: The Case of Nicaragua*. Praeger Publishers.
- Uddin, Md Hamid, Sawsan Halbouni, and Mahendra Raj. 2014. "Performance of Government-Linked Firms Listed on Two Stock Exchanges of the United Arab Emirates: An Empirical Study." *Emerging Markets Finance and Trade* 50:212–236.
- Uzzi, Brian. 1997. "Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness." *Administrative Science Quarterly* pp. 35–67.
- Van Duijn, Marijtje AJ, Tom AB Snijders, and Bonne JH Zijlstra. 2004. "p2: A Random Effects Model With Covariates for Directed Graphs." *Statistica Neerlandica* 58:234–254.
- Van Duijn, Marijtje AJ, Jooske T Van Busschbach, and Tom AB Snijders. 1999. "Multilevel Analysis of Personal Networks as Dependent Variables." *Social Networks* 21:187–210.
- Zijlstra, Bonne JH, Marijtje AJ Van Duijn, and Tom AB Snijders. 2006. "The Multilevel p2 Model." *Methodology* 2:42–47.