NETWORKS, BUSINESS GROUPS, AND CORPORATE GOVERNANCE IN EMERGING FINANCIAL MARKETS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Michael Charles Siemon

August 2018

NETWORKS, BUSINESS GROUPS, AND CORPORATE GOVERNANCE IN EMERGING FINANCIAL MARKETS

Michael Charles Siemon, Ph.D.

Cornell University 2018

Chapter 2: Business groups are an essential feature of economic and social organization across the world. As such, they provide a compelling empirical and theoretical tool for improving our understanding of the ownership and control of organizations outside of the Anglo-American tradition. However, widely incorporating business groups into analyses of regions like the Middle East has been limited by the availability of information on group membership. This paper seeks to advance the comparative study of corporate governance by demonstrating a method for inferring these groups with data on director-firm and ownership networks of publicly traded companies using data for 1336 listed firms from 12 Arab countries in the Middle East and North Africa (MENA) region. Group-level measures are then calculated based on information aggregated from firms affiliated with each potential group. These data are used to support the validity of the communities by showing that—based on scores of family influence, geographic heterogeneity, and government ownership—it is possible to identify groups corresponding to one of these specific and measurable axes of solidarity.

Chapter 3: Although business groups have been intensely studied in regions like Latin America and East Asia, the literature on the Middle East and North Africa is relatively less developed; the sparsity of this research is a pressing concern because of the variety of ways in which business groups have been

shown to influence the political allocation of economic resources in these other parts of the world. This study presents evidence for the value-relevance of family business groups, government ownership, and other inter-firm relationships among 1110 publicly traded firms in 12 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. Next, I apply a Bayesian multilevel model to estimate the associations between group-comembership(as well as other relationships), and pairwise stock returns correlations. Using the results of the model, I perform an exploratory analysis of the correlations between exchange-level attributes and estimated coefficient values, finding that, despite the small number of exchanges, minority-investor protections appear to be associated with stronger value-relevance of business groups.

Chapter 4: We use a novel data set of the outcomes of 18125 loan guarantee transactions attempted by Chinese public firms between 2008 and 2015 to study banks' reactions to related party transactions in the context of ongoing marketization reforms. We find that guaranteed loans with better governed guarantors from more marketized provinces are more likely to succeed, that the influence of political ties is contingent on the ownership structure of the guarantor, and that the embeddedness of the transaction in multiplex relationships between the guarantor and bank increases the likelihood of approval.

BIOGRAPHICAL SKETCH

Michael Siemon graduated from the University of Chicago in 2008 and currently resides in Illinois.

In memory of my mother, Dinah

ACKNOWLEDGEMENTS

I would like to thank my advisor, Victor Nee, and other committee members, Richard Swedberg and Benjamin Cornwell, for their support.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

This dissertation is composed of three papers that investigate the relationships between inter-organizational networks, business groups, and emerging financial markets. More specifically, Chapter 2 presents a procedure for using existing community detection methods to infer the business group affiliations of over 1100 firms in the Arab countries of the Middle East and North Africa. Chapter 3 seeks to confirm these results by using the inferred groups to predict the degree of pairwise stock price co-movement between firms. Chapter 4 analyzes the interactions between networks, business groups, and financial transactions using a unique data set of attempted loan guarantee transactions in China from 2008 to 2015.

These interrelated papers make several contributions to the economic sociology of emerging financial markets. First, Chapter 2 builds on the definition of business groups provided by Granovetter (2010) as well as empirical work by Khanna and Rivkin (2006) in seeking to establish whether it is possible to link mathematically derived clusters in inter-firm networks with the socially constructed idea of business groups. In doing so, it exposes broad regularities in the types of traces that business groups create in formal networks across a variety of countries, especially in terms of family control, government ownership, and geographical location. Chapter 3 takes as its point of departure the finding that firms within the same business group have correlated stock returns (Khanna and Thomas 2009). The motivating idea is similar to Chapter 2 in that just as business groups leave their mark on formal networks, their influence should also be evident in a weighted network of pairwise price correlations. The results

support this hypothesis, and finding traces of the same inferred groups in multiple, independently generated networks supports the validity of the method. Through the use of a multilevel model it also sheds light on the cross-national variation in the relevance of family business groups, which is important given the scarcity of such broad analyses (Davis 2012; Aguilera and Jackson 2010; Colli and Colpan 2016).

Finally, Chapter 4, which was co-authored with Lisha Liu, switches focus from the inductive study of business groups in the Middle East and North Africa to a deductive analysis of guaranteed loans in China. Our model seeks to integrate findings from three strands of literature. First, we use our transaction-level data set to analyze whether banks respond to the corporate governance of guarantor firms, providing a further characterization of the evolving relationship between controlling and minority investors in China (Liu and Lu 2007; Jia et al. 2013; He et al. 2013). Second, we use a variety of measures to analyze the role of firms' political connections in securing loan guarantees. This has directly relevance for the large literature on political capital in China's transition economy (e.g. Nee and Opper 2010; Haveman et al. 2017). Finally, our most novel contributions builds on the idea of the embeddedness of economic and financial transactions within a wider set of social relationships (Granovetter 1985; Uzzi 1999; Zhou et al. 2003). Our question is thus whether and under what conditions firms are more likely to receive a guarantee in the presence of such multiplex ties or a history of successful transactions. We also investigate whether the above relationships are contingent on the type of bank involved.

CHAPTER 2

IDENTIFYING AND CLASSIFYING BUSINESS GROUPS IN 12 ARAB COUNTRIES

2.1 Introduction

Business groups, like firms or states, are an essential feature of the organization of economic life, and a wealth of case-studies has demonstrated their diversity and geographical extent (Khanna and Yafeh 2007; Young et al. 2008; Granovetter 2010; Carney et al. 2011). Nevertheless, analysis of business groups can be limited by the availability of data, and the frequent use of country case-studies reflects the difficulty of observing business group membership in aggregate. As side effect of this practical difficulty, researchers' preference for countries where they are more likely to obtain data necessarily paints a selective picture of how groups operate, and little work has been done to analyze and describe business groups in the Middle East and Africa (though see Hearn 2014; Hearn et al. 2017, for two recent exceptions).

This paper advances business group research by using network analysis to identify business groups across the Arab states of the Middle East and North Africa (MENA) region. First, I explore the clustering of publicly traded firms in the region with a community-detection algorithm that searches for the common structure latent in two separate inter-firm networks. Next, I focus on the communities generated in the first stage and classify each set of firms according to an measurable basis of solidarity. Finally, I discuss specific examples of the business groups I identify and conclude by presenting possibilities for employing and improving the method.

This two-pronged analysis provides the first descriptive and comparative account of business groups in the contemporary MENA region and thereby allows for a rough estimate of the importance and diversity of business groups in the area. The broader purpose of the method is to enable the inclusion of business group membership in future analyses of firm outcomes. This has important implications for policy, especially in the field of corporate governance, since it allows researchers and practitioners to avoid the assumption that firms behave as independent or isolated actors and instead consider the latent structural factors that influence firm behavior. Furthermore, this study is an example of applying the tools of big data research to policy-relevant problems, especially in terms of increasing the transparency of important social, political, and economic structures.

The rest of the paper is organized as follows. The next section describes previous work on business groups. The method section presents the data and outlines the procedure for using publicly-available information on firm networks to infer group membership. The results section combines the exploratory partitioning of these firm networks into overlapping clusters with an analysis of cluster-level properties to show that this approach reveals potential groups characterized by varying levels of family control, government ownership, and foreignness. Finally, I conclude with a discussion of implications, limitations, and next steps.

2.2 Business Groups, Networks, and Corporate Governance

2.2.1 Defining Business Groups

A major achievement of the business group literature has been its transformation from scattered 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 business groups 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.

Overall, however, the findings of this literature have been mixed, and two recent 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 this contingency is frustrating from a policy perspective, it 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. Instead, available evidence suggests that they tend to emerge as a response to local con-

ditions 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 are obliged to identify areas of common ground.

Shared definitions tend to emphasize three features in particular. 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, whose exact content can take variety of forms. One key regularity, however, is that ownership is 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 scholars 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 the studies cited above, the analysis of this paper does not begin with direct observations of group membership, and as a result a clear working definition of what can plausibly be defined as a business group is even 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 consensus definition presented above. Key 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, the first part this 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 an additional aspect of business groups that is often ignored in formal analyses, the underlying basis of solidarity in each group responsible for maintaining the permanence and cohesion of group ties.

In his summary of early business group research, Granovetter (1995) gives a definition that covers the three areas noted by later authors but goes on to argue that business groups are distinguished from other collections of firms by "the existence of social solidarity and social structure among component firms". This solidarity can take a wide variety of forms, and while family is the most widely studied example (Morck and Yeung 2004; Luo and Chung 2005; Masulis et al. 2011), others include religion, ethnicity, geography, and state control. This insight serves as a foundation for the second phase of my analysis. Generat-

ing measures for three of these sources—family, geography, and state control—provides a crucial opportunity to confirm or reject the salience of the clusters from the first stage based on measuring potential sources of that group's cohesion.

2.2.2 Business Groups and the Economy

Research on business groups is motivated by both their ubiquity and importance. Recent efforts to collect and synthesize our knowledge on the topic paint a rich, if incomplete picture. For example, Carney et al. (2011) include 141 studies in 28 countries as part of their meta-analysis, and this sample overlaps with many of the most commonly studied countries covered by Khanna and Yafeh (2007) and Colpan et al. (2010). As for their economic importance, reliable global estimates are not available, and we rely on country-level studies. In Turkey, which along with Israel is the only state in the MENA region where business groups have been studied in detail, business group members employed over 400,000 people in 2005 (Colpan 2010). At the firm level, estimating the proportion of affiliated firms in a country is usually only possible for the population of listed firms, and Khanna and Yafeh (2007) report ratios of approximately 50% for Argentina, Brazil, Indonesia, South Korea, Taiwan, Thailand, and Turkey, while Chile, Israel, Mexico, and the Philippines range from around 30% to 20%.

These figures raise an important question. If business groups are so widespread and influential, why do measurements tend to be so selective? On a practical level, much of our knowledge about organizations and their relationships is derived from legally mandated disclosures. While a comparative

account of the legality and regulation of business groups across the world is beyond the scope of this paper, the simplest explanation is that corporate law often ignores business groups and hence makes no demand for their widespread disclosure (Granovetter 1995). Of course, this ignores the substantial variation in country-level practices, but on an anecdotal level researchers report using a combination of interviews, surveys, and private data to obtain data on group membership (Guillen 2000; Fisman 2001; Khanna and Rivkin 2006).

Situated in the nebulous middle-ground between firm- and country-level analyses, the difficulty of measuring group affiliation contrasts with well established traditions of studying these two fundamental aspects of the organization of economic, political, and social life. Nevertheless, research comparing group members and unaffiliated firms continues to show that business groups should not be ignored. For example, a recent study of Korean firms before and after the 1997 Asian financial crisis shows how intra-group capital transfers enabled high-growth members to mitigate the effects of the crisis relative to their unaffiliated peers (Almeida et al. 2015). The transfer of resources among a set of tightly knit firms is a fundamental mechanism largely unavailable to firms outside of the trust, solidarity, and obligation engendered by business groups. Still, the impacts of such coordination are not always positive. Sociological studies of the embeddedness of economic behavior within webs social relations have demonstrated the costs of such redistribution to successful firms, and business groups are no exception (Lincoln et al. 1996; Uzzi 1997). This type of leveling might be consensual for firms who value the stability and other benefits of group membership more than their cost in redistributed profits, but business groups are nonetheless notorious for their role in undermining legally sanctioned and formalized institutions which may or may not reflect their informal patterns of

coordination.

Indeed, the recent impeachment of South Korean president Park Geun-hye—the daughter of the dictator credited with establishing the chaebol as the key institution of the modern South Korean economy—over her top aide's solicitation of bribes from the country's major business groups demonstrates the potentially explosive nature of these informal networks of political and economic power. The public airing of such corruption is the exception rather than the rule, but business groups are also accused of other, more easily observed abuses. Chief among these is the expropriation of the returns owed to minority shareholders by the controlling interests within a group (Bertrand et al. 2002; Bae et al. 2002). The prevalence of this type of behavior is an empirical question, but its implications for financial development are dire (Porta et al. 2002; Young et al. 2008). To wit, if investors conclude that they are likely to experience such treatment when investing in a particular firms or country, they will reduce the amount of capital they are willing to provide for a given amount of ownership, perhaps to the point of being unwilling to invest at all.

2.2.3 Business Groups and Corporate Governance

Compared with firm-level analyses of business group membership and performance, this last point touches on a deeper and less immediately tractable problem facing the students of business groups, namely their relationship with country-level institutional environments. Scholars have long maintained that business groups arise in response to poor market institutions that make it difficult for firms to raise capital and hire talented employees (Guillen 2000; Khanna

and Palepu 2000; Yiu et al. 2005), but even if this is generally true, the increased capabilities of group members might reduce their incentive to improve the general environment.

The plausibility of this type of causal mechanism has important implications for how we perceive business groups. A further consequence of the business group literature's reliance on single country studies is that there is little room for analyzing variation in group- or country-level outcomes, and as such this work tends to operationalize business groups at the firm-level with a simple binary indicator of group membership. In contrast, the goal of this paper is to remove the onerous data-collection constraints faced scholars who wish to investigate the role of business groups outside of commonly studied countries like Korea, Japan, India, or Chile.[1] In this sense, whatever uncertainty we introduce into our analyses by inferring rather than directly observing business group would be balanced by new opportunities for large-scale, cross-national research, perhaps up to the point of fitting hierarchical models of business group effects, which have the desirable property of allowing coefficient estimates to vary among different subpopulations (see Gelman and Hill 2006).

Assessing group- and country-level variation in firm behavior and outcomes could illuminate previously obscure facets of financial and economic development, especially in the realm of corporate governance. Defined variously as "the ways in which suppliers of finance to corporations assure themselves of getting a return on their investment" (Shleifer and Vishny 1997b, p.737) or more broadly as "the way in which corporations are owned, controlled, and coor-

[1]This concern is all the more pressing given that the availability of business group data is linked to issues of disclosure and transparency. Hence, we should expect that studies of countries with relatively high-quality institutional environments are more likely to obtain necessary data and thus that focusing solely on these cases produces insights that might not be generalizable to under-studied areas.

dinated and set their goals" (Colli and Colpan 2016, p.275), corporate governance is intimately concerned with the competing interests of individuals with de facto versus de jure control over organizations, and given that the analysis of such conflict already involves firms' relationships with external actors, there is no reason why it could not be extended to include group-level processes. Indeed, situating research on corporate governance with reference to the broader social and political context has been identified as a promising area of future investigation (Morck et al. 2005; Davis 2012; Aguilera and Jackson 2010; Kogut 2012). However, in their exhaustive review of research on business groups and corporate governance Colli and Colpan (2016, p.297) note that such work is still "developing", and they identify cross-national studies as a key area for future research.

Although corporate governance's concern with the rights of investors might seem limiting, in fact its implications for financial and economic development are well established. Despite the successful diffusion of market-based policies during and after the conclusion of the Cold War (Lee and Strang 2006; Weber et al. 2009), the wide range of outcomes experienced by countries undergoing such reform brought attention to the quality, rather than quantity, of institutional change, and the case of financial liberalization and the development of stock markets is an especially instructive example. In assessing the economic impact of these changes, studies at both the firm- and country-level agree that institutional quality, not market size, was the key predictor of successful financial development (Demirgüç-Kunt and Maksimovic 1998; Levine and Zervos 1998; Henry 2000; Hail and Leuz 2006; De la Torre et al. 2007). Hence, corporate governance research outside of the world's most developed economies has focused on better understanding the causes and consequences this differential

institutional development.

In the MENA region, the growth of stock markets (Figure 2.1)has provided an opportunity to test existing corporate governance theories in a new empirical context while also suggesting alternative mechanisms that appear to improve cooperation between controlling owners, managers, and minority shareholders. Concerning the economic impact of financial liberalization, a comprehensive study of listed firms in 11 countries in the Middle East and North Africa reports that in their sample "liberalization has no effect on economic and investment growth" (Naceur et al. 2008, p.673). Rather, they concur with the studies cited above that the specific institutional environment achieved prior to liberalization is the best predictor. Lagoarde-Segot (2013) reinforces this message in his study of Tunisia, which appears as an example of a country not implementing the types of institutional reforms emphasized in so much of the above work. Other corporate governance studies in the region provide focused views of specific issues like market reactions to the 2008 financial crisis (Farooq et al. 2013) and the Arab Spring (Hearn 2014). This latter study is especially notable as being one of the only studies of which I am aware that directly includes business groups in a study of firms in an Arab MENA country, although there are several studies focusing on Turkey or Israel, and Hearn et al. (2017) furthers this important work with an analysis of IPO financing on 22 exchanges throughout Africa. The rest of the literature focus on areas like the benefits associated with better governance (Al-Khouri 2006; El Mehdi 2007; Omran 2009), or the predictors of board composition (Elsayed 2010) and salary disclosure Hearn (2013).

In summary, business groups present a paradox. They are a fundamental means of coordinating large formal organizations throughout the world, yet

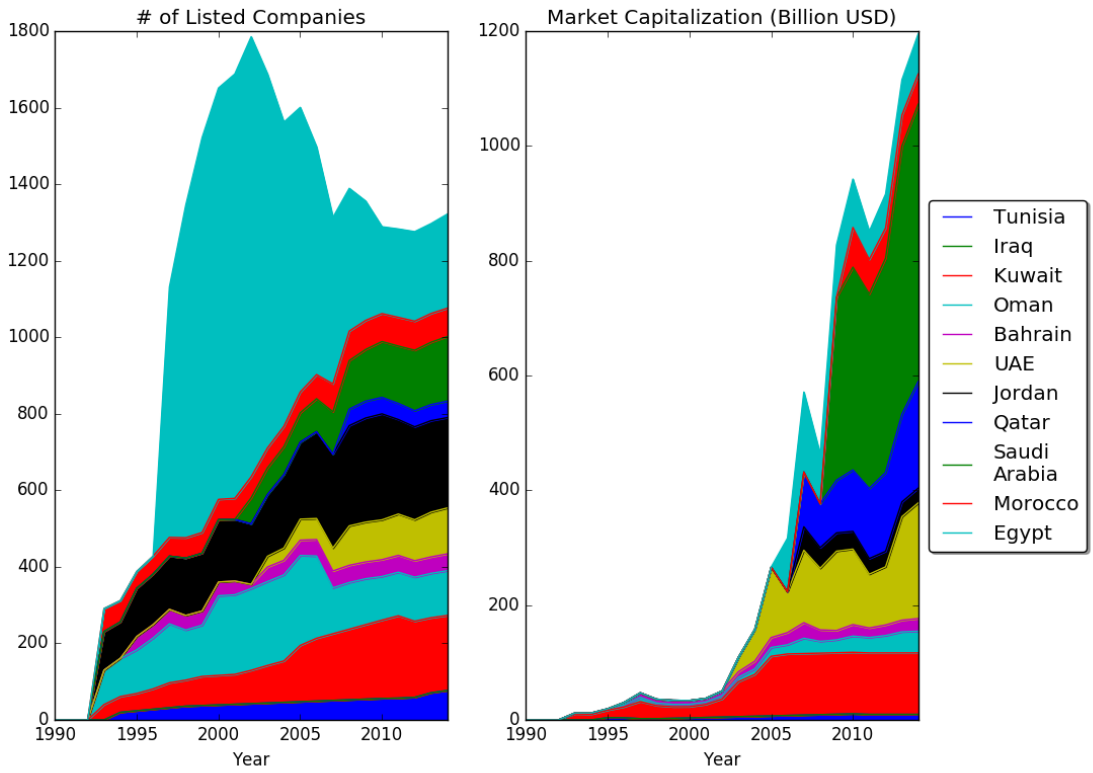

Figure 2.1: Growth of equity markets in the MENA region. Data are taken from the World Bank Development indicators.

they operate through combinations of social relationships that not only vary along numerous dimensions, but are also difficult to observe. This section has argued that despite these challenges, the impact of business groups on economic and institutional development is too powerful to ignore.

2.3 Methods

2.3.1 Data

The data are from three primary sources. First, information on corporate networks comes 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. Overall, I collected data on 1336 unique publicly traded firms in the region from firm profiles constructed from 2015 annual reports. These firms are listed on at least one of 13 exchanges in 12 countries. These are: Morocco, Tunisia, Egypt, Palestine, Jordan, Saudi Arabia, Oman, the United Arab Emirates with exchanges in both Abu Dhabi and Dubai, Qatar, Bahrain, Kuwait, and Iraq. I supplemented this information with SIC codes from Bureau van Dijk's Orbis database, and data on government control from Zawya.com.

2.3.2 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.

2.3.3 Networks

Matching names across appearances in different firm profiles reveals the social structure of publicly traded firms. For example, observing that an individual who owns large blocks of several companies and serves as the chairman in others demonstrates the influence of that particular actor, but it also signals the relatedness of these companies since common ownership and employment would provide an incentive and a means to cooperate. Networks, also known as graphs, formalize this structure using two fundamental elements, nodes and edges. A node can be any entities of interest, in this case firms and individuals. Edges represent the relationships between different nodes, and different types of relationships can be studied by constructing different networks, each defined by a unique type of tie or ties measured in a particular time period. I focus on two types of edges, ownership of companies by individuals and other firms and the employee-employer relationship between directors or managers and the firms they serve.

In the former, edges have an inherent direction and thus distinguish between an edge representing Firm A owning part of Firm B and vice versa. They are also weighted in the sense that we can measure an edge's importance by the percent of a firm held by that particular owner. In contrast, a director-firm network contains uniformly weighted edges between nodes of complimentary types; there

are no employment relationships between firms or between individuals. This leads to a "bipartite" structure that can be analyzed in its own right or condensed to describe relationships between only firms or only individuals. The latter option interprets the fact that a single person sits on the boards of two companies as representing a connection between the two firms, for example by communicating important information from the management of one firm to the other. This is known as a "director-interlock" and has been a key focus of the literature on interorganizational networks (Mizruchi 1996). These two types of networks have formal properties that inform the community detection method described below.

Before analyzing the graphs, I make three alterations to the raw data. First, I delete all nodes that are not a publicly traded firm or else tied to at least two such companies. This is to prevent the community detection algorithm from producing a degenerate solution where most firms are simply grouped alone with their large numbers of exclusive connections. Removing any director, manager, or owner who was only tied to a single firm across both networks focuses attention on the observed structure between public companies, which is appropriate given the limited information on individuals or private firms. Next, I transform the ownership edges to make the sum of their weights on the same scale as the number of edges in the director-interlock network and to better reflect the qualitative importance of different ownership levels[2].Finally, I analyze the networks one connected component at a time, which means that I assume that groups of firms that share no connections are in separate communities and focus on identifying clusters within each set of connected firms.

[2]Much of the important variation in ownership is within the range $0 - 0.2$ because after this point the control conferred by additional stock diminishes (La Porta et al. 1999; Claessens et al. 2000).

2.3.4 Community Detection

Inferring group membership from network data rests on the assumption of a close connection between the mechanisms that form business groups and those that generate the observed networks. Support for this assumption comes from three sources. First, as described in Subsection 2.2.1 the definition of a business group relies on the informal and formal ties that bind member firms into a cohesive unit. This dovetails with basic tenets of sociological theory which argue that ties tend to arise around important focal points (Feld 1981). Second, empirical treatments of business groups often rely on private information kept by regulators, field interviews, or purchasing data sets from private firms. Studies making use of this last source often note that these providers rely on subjective assessments of director interlocks (Fisman 2001; Khanna and Rivkin 2001). Finally, the strongest support for this assumption comes from Khanna and Rivkin's (2006) analysis of the predictors of business group co-membership in Chile. In their analysis, the authors find that the presence of ownership and director-interlock ties between two firms are strong predictors of co-membership, while familial connections are of lesser importance. Using networks to study business groups is a well established practice, and both anecdotal reports and statistical analyses indicate that ownership and interlock networks contain enough information to approximate business groups as a form of community structure.

Communities in this sense refer to clusters of nodes that exhibit more internal and fewer external ties than we would expect in a random graph with similar properties, and the development of methods to detect these communities has grown into an important subfield of network analysis (Newman 2006;

Fortunato 2010). In general, community detection relies on the concept of a network clustering, which assigns each node in the network to at least one community, though most methods limit nodes to exactly one. Community detection also requires a means of assessing the quality of a given clustering, for example a function that measures the total intra-community connectivity relative to a baseline that reflects the properties of the network. The most commonly studied such quality function is called modularity. The basic form is described in Equation 2.1 where Q is the modularity value, A is an adjacency matrix such that A_{ij} equals the weight of an observed edge between nodes i and j, $2m = \sum_{ij} A_{ij}$, $k_i = \sum_j A_{ij}$, c_i is the community assigned to node i, and δ_{xy} is the Kroneker delta, which equals 1 if $x = y$ and 0 otherwise.

$$Q = \frac{1}{2m} \sum_{ij} \left(A_{ij} - \frac{k_i k_j}{2m} \right) \delta_{c_i c_j} \tag{2.1}$$

The two inputs to the modularity function are the network in the form of the modularity matrix A and a vector of community indices, c. The essential feature of modularity is its choice of expected value for each edge, which is based on the values k_i and k_j, which are called the degree of nodes i and j and used to measure their connectivity in the network. Modularity then uses the expected value for an edge between nodes i and j, which is based on their degrees and the total number of edges in the network, as a baseline against which to assess intra-community connectivity. This specification assumes that edges are observed between nodes only based on each node's connectivity and thus that edge formation does not depend on other forces that encourage links between some nodes and not others. Moreover, this assumption can be expressed differently for different types of networks. In this analysis, I adapt the multiplex mod-

ularity function of Mucha et al. (2010) according to the specification for directed and bipartite networks developed by (Leicht and Newman 2008) and Barber (2007), respectively. The specification, parameterization, and derivation of the full modularity functions used in this paper are provided in the Appendix.

In addition to being widely studied, modularity's flexibility makes it well suited to analyzing ownership and interlock networks. It does, however, also present important challenges. First, modularity has an inherent "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.

Accounting for the resolution limit depends on our prior expectations of group size. On the one hand, we might expect from prior studies that the mean number of public firms be less than ten with the common occurrence of a few disproportionately large groups in each country (He et al. 2013). One approach is to include a parameter γ which increases the penalty of grouping unconnected nodes, and it is possible to set γ to return results consistent with our prior expectations, in this case requiring value of at least 6.0. However, this approach is limited in that it imposes a uniform increase on the penalty for including unconnected nodes in the same community. In many cases, this fits the way data have been generated or observed, but analyzing nested but interconnected networks presents a special challenge. 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 countries, and then to treat each top-level community as a separate network.

In addition to the necessity of specifying community size, modularity methods also struggle with degenerate solutions and local optima (Good et al. 2010). This necessitates caution in interpreting communities based only on relational information, and my solution to this challenge is the further analysis of community-level attributes. Simultaneously analyzing both corporate networks helps to mitigate the noise in the data by searching for areas of agreement between the two networks[3], but not all areas exhibit shared patterns. Strategies for confronting this degeneracy by looking for areas of stability or consistency across multiple solutions have recently been proposed and represent an area for improving the method (Lancichinetti and Fortunato 2012; Bassett et al. 2013). More will be said on this issue below in the section on community-level attributes and in the conclusion.

After specifying an appropriate modularity function, I use an adaptation of Newman's 2006 spectral method to repeatedly bisect the two types of networks based on a weighted combination of their edges until doing so does not improve the overall modularity of the network. Details are presented in the Appendix. I chose this method due to the ease of implementation, but numerous other algorithms could be used that are suited to millions of nodes and better identify

[3]For example, an OLS regression found that nodes sharing a direct link in one network are on average three degrees closer in the other after controlling for node centrality.

area of significant versus incidental clustering (Blondel et al. 2008; Lancichinetti et al. 2011; Lancichinetti and Fortunato 2012; Bassett et al. 2013). An important feature of the multiplex method developed by Mucha et al. (2010) is that a node may be assigned a different community in each network. In the results presented below, 14% of nodes were place in two separate communities. I treat a firm as belonging to a community if it is present in at least one network.

To reiterate, community detection methods can translate the information stored in networks into an inference about business group structure. Doing so requires tailoring both the networks and methods to enhance the visibility and interpretation of these latent communities, but even so modularity maximization often fails to deliver an optimal or intuitive solution. Below I describe the assessments I perform on the results of the community detection algorithm, which hinges on using the aggregate properties of communities to tie them to one of three organizing axes of solidarity which characterize business groups.

2.3.5 Measures of Community Attributes

The results of the community detection algorithm are thus an informed guess about the latent forces represented by the two corporate networks. There is no guarantee that they truly represent the global optimum. Even if they did, there is doubt whether this optimum would correspond to an intuitive idea about business groups. To mitigate these concerns, I attempt to establish the social relevance of each community by measuring the presence of three potential sources of solidarity: family control, government ownership, and geography in addition to the potentially confounding influence of industry homophily.

Recalling the definitions presented above, business groups are more than random collections of firms cooperating on an ad hoc or transactional basis. Rather, they are cohesive and enduring communities of organizations, and such solidarity can in some cases be measured. I consider measure family control by grouping individuals by shared or highly similar surnames, government by hand-coding owners as state or non-state, and geography by accounting for the appearance of the same individuals in records from different exchanges. These measures are based on all individuals and firms tied to a community, including those deleted before running the community detection algorithm.

Family measures are based on assessing the extent and exclusivity of a family's involvement with a particular community. It includes the following measures: the percent of firms with at least one family member as an owner, director, or manager; mean family ownership; the family's share of managers, chairmen, and directors; the percent of firms in which the family collectively owns a 10% stake[4]; and the average number of unique individuals associated with the major surnames in that community. For each group, I assess family control based on the combined scores for the top two families measured in terms of their overall involvement. I attempt to break ties in favor or larger families. Note that these measures cannot account for family control of private companies and are thus likely to understate family ownership through these intermediaries.

I measure government ownership by hand-coding all owners based on their names and web searches. This includes clear government entities like ministries of investment or finance as well as state-owned enterprises. Just as in the case of family ownership, I construct measures of mean government ownership and

[4]This cutoff represents the amount of ownership necessary for significant control over a company (La Porta et al. 1999; Claessens et al. 2000)

the percentage of firms in which government entities own at least 10% or 20%. I do not attempt to code the political affiliation of individuals.

Location is another potential source of solidarity, and I develop a series of location scores for each node based on its affiliations with companies listed on different exchanges. Using an indirect proxy for true country of origin is not likely to be accurate in every case, but on the whole it captures key patterns indicating geographic diversity at the community level. For example, a director for a single Egyptian firm would be classified as Egyptian while a director for two Kuwaiti firms, one Egyptian firm, and an Iraqi firm would be classified as Kuwaiti. This allows us to assign each node one or more top countries while also measuring its geographic homogeneity.

Using these measures for each listed firm, I generate three of geographic homogeneity measures. I focus on geographic homogeneity rather than a binary foreign-domestic variable due to the difficulty of maintaining this classification for multi-country data. A geographically heterogeneous community is characterized by ownership and control relationships that span national boundaries, and given the otherwise strong tendency for edges to occur within close geographical proximity, this type of anomaly might indicate a strong relationship. The first measure, top exchange score, is the percentage of firms that are listed on the most common stock market within that community. The next measure includes information from each firm's activity in multiple markets and is the mean strength of each listed firm's affiliation with its top market, hereafter referred to as mean firm domesticity. The final statistic is the mean strength of each listed firm's affiliation with the the most common stock market from the top exchange score, i.e. that group's "home" exchange. The difference between

these two measures is that this mean home exchange score includes the same broader set of information as mean firm domesticity.

Finally, industry is also a likely source of inter-firm connections due to homophily, or the tendency of nodes to form ties with similar alters (McPherson et al. 2001). I broadly assess absolute homophily at the group level as the percent of public firms in the community belonging to the most common top-level SIC code in that community. Since financial firms are over-represented on the stock markets of the region, I also include the the relative degree of industry concentration in the form of the euclidean distance between a vector of the percent of community firms within each category and the same vector defined for all public firms in that community's home country.

2.3.6 Corroborating and Classifying Business Groups

The final stage in my analysis is using these measure investigate the correspondence between the inferred communities and socially meaningful groups. I explore the structure of these data with a Gaussian Mixture Model (GMM), which assumes that the data were generated by a given number of distinct normal distributions and further clusters the data along the dimensions described above. GMMs require the researcher to specify the number of underlying distributions. I chose four, one for each category of variables. Using these clusters as a guide, I briefly describe two examples from each category.

2.4 Results

2.4.1 Communities

The first step in identifying business groups from corporate networks was to simultaneously divide the two graphs into a relatively small number of top-level communities. I use node location as a first guess about the broad structure of the data and assign each node to exactly one exchange. This includes cross-listed firms, and ties were broken randomly. Next, I employ a Kernighan-Lin algorithm to incrementally switch the community assignments of nodes until no modularity-increasing moves can be made (Kernighan and Lin 1970). Figure 2.2 and Figure 2.3 provide a direct visualization of the minor differences between country of origin and top-level community. Nodes are colored by their top exchange and refined top-level community, respectively.

These figures are based on a composite network where a single undirected tie between two nodes indicates the presence of at least one type of connection. Large nodes represent public firms, and all other individuals and organizations are shown by the smaller nodes. In terms of modularity, the refined boundaries slightly outperform the raw exchange-based communities, 0.632 to 0.630. Table 2.1 describes some of the properties of each community in the ownership and director-interlock networks. These 13 communities vary in the number of unique public firms they contain across both networks, the density of connections between public firms, their modularity, and in the extent to which they are contained within one or two exchanges.

Finally, Figure 2.4 shows same network as Figure 2.3 with nodes colored by

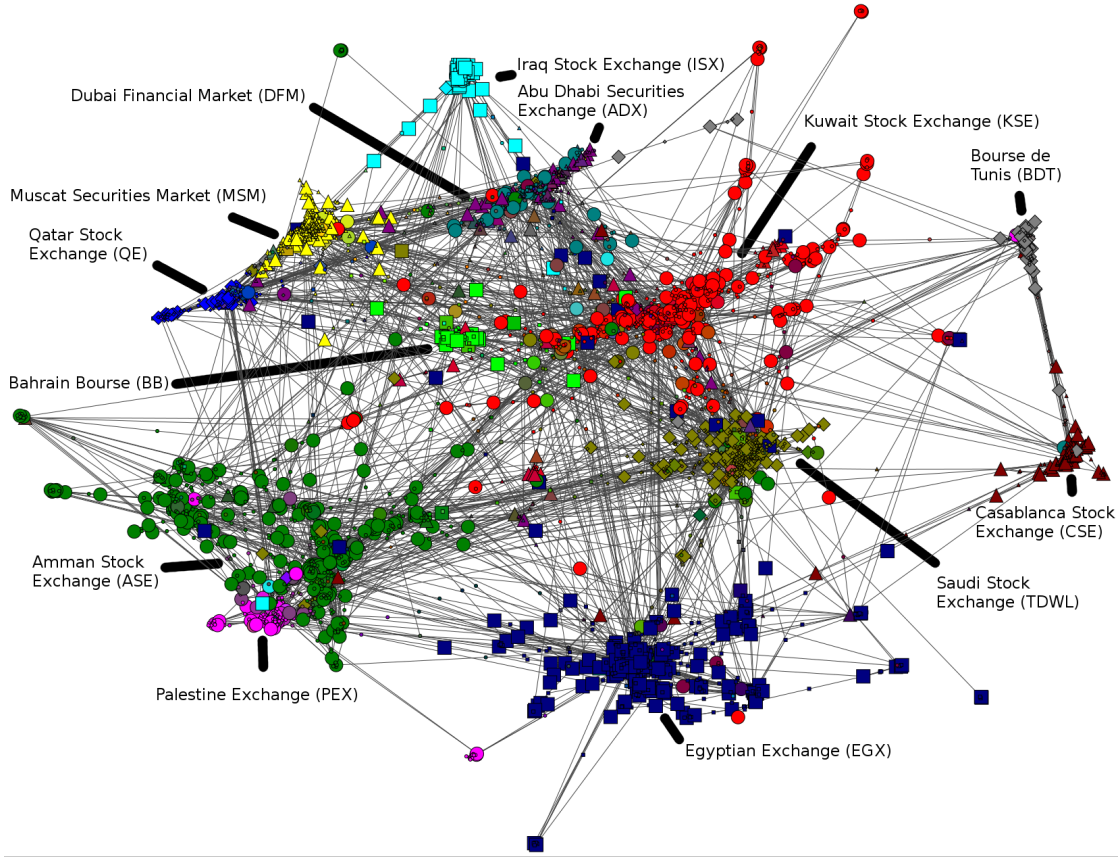

Figure 2.2: A visualization of the network of publicly-traded firms in the MENA region. Nodes are colored and shaped by inferred country of origin. Edges indicate the presence of at least one interlock or ownership connection.

Top exchange Second Exchange	KSE DFM	QE BDT	MSM DFM	PEX CSE	ISX	BDT PEX	CSE	EGX KSE	TDWL ASE	DFM CSE	ADX DFM	BB KSE	ASE PEX
No. of unique public firms	195	43	114	42	75	58	57	195	168	40	74	35	238
No. of interlock edges	505	214	397	184	115	108	143	437	659	112	270	98	794
Mean interlock degree	2.58	4.97	3.48	4.38	1.53	1.86	2.50	2.24	3.92	2.80	3.64	2.80	3.33
No. of ownership edges	628	178	417	223	187	304	265	743	875	108	170	83	834
Mean ownership degree	3.92	4.44	4.31	6.73	3.74	6.31	5.24	4.17	5.57	3.10	2.51	2.88	4.31
Modularity	0.579	0.502	0.533	0.635	0.676	0.533	0.535	0.635	0.617	0.670	0.626	0.573	0.640
Number of lower communities	25	12	20	9	18	13	11	29	29	15	18	8	28
% of firms top exchange	0.90	0.95	0.96	0.95	1.0	0.97	1.0	0.95	0.94	0.92	0.78	0.87	0.92

Table 2.1: Network statistics of each higher-level community

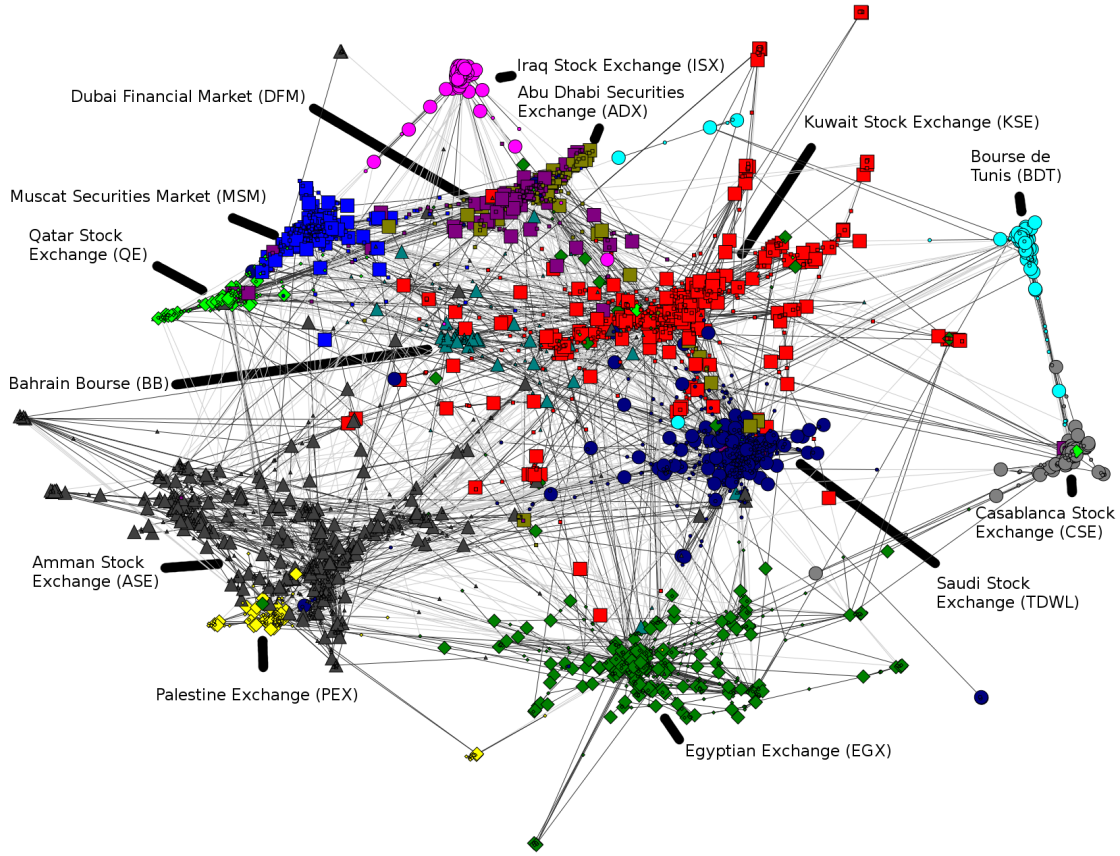

Figure 2.3: This visualization has the same layout as Figure 2.2. Nodes are colored and shaped by location-based community. Inter-community edges have been lightened.

their lower-level community in the director-interlock network. The algorithm identified total of 237 lower-level communities. Of these, 186 have two or more public firms. The rest of the paper focuses on assessing the validity of these communities.

2.4.2 Classifying Communities

Figure 2.5 shows the results of clustering each community-level data point with more than one listed firm. The axes are defined by the mean score for standardized variables belonging to the three sources of solidarity described above.

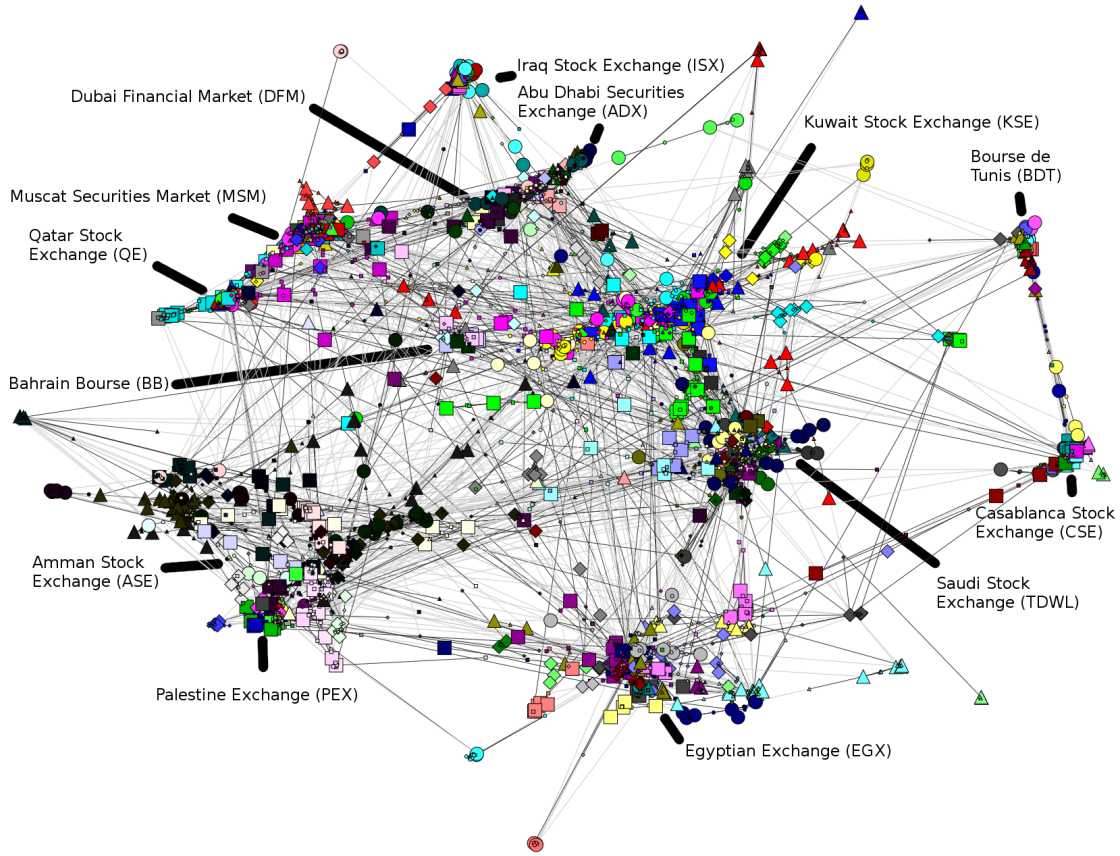

Figure 2.4: This visualization has the same layout as Figure 2.2. Nodes are colored and shaped by their lower-level community assignment. Inter-community edges have been lightened.

The four clusters are shown by different colors and shapes with one along each axis and a fourth near the origin. The GMM result classify 17 groups as family-dominated, 26 as largely government-owned, and 37 as spanning multiple exchanges. The remaining 106 data points are concentrated at the far center corner of the figure, indicating communities that lack distinctive scores in any category. This lack of corroboration stems from at least three sources. First, it is possible that the community does in fact correspond to a business group, but that it relies on an unobserved form of solidarity. Second, the algorithm may have incorrectly identified the boundaries of a group, preventing the full observation of family or political ties. This could be the result of incorrect data or one of the many degeneracies to which the modularity function is vulnerable. Finally, the

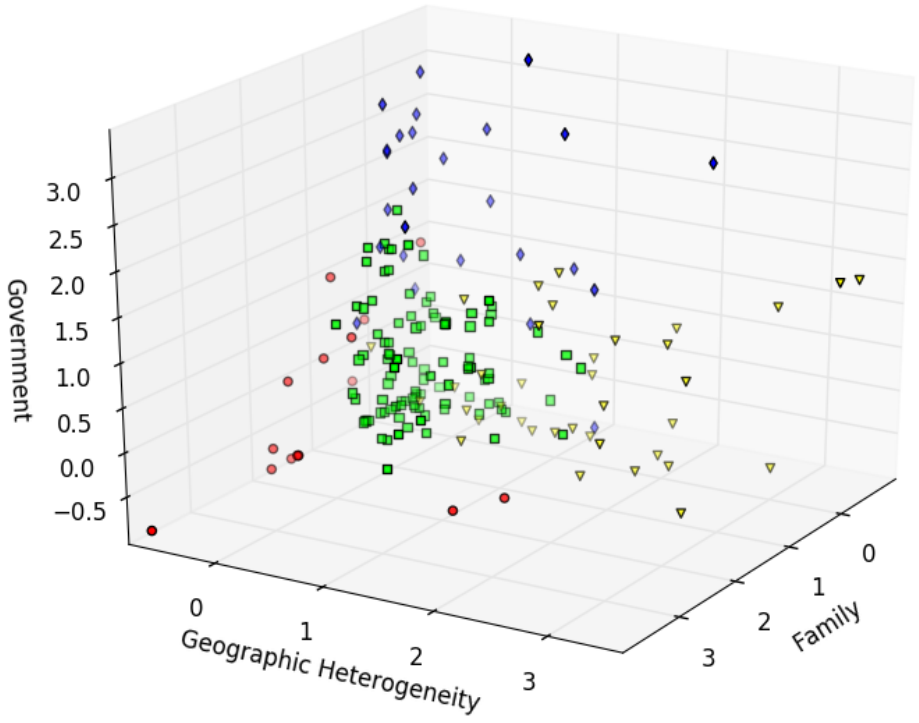

Figure 2.5: A three-dimensional scatter plot showing the results of the GMM clustering algorithm. The x, y, and z axes are based on the average z-score of the community measures in each category. Family groups are the red circles, government groups are shown as blue diamonds, transnational groups are yellow triangles, and the remaining communities are shown as green squares.

most common cause is likely that many firms are simply not members of any group.

Figure 2.6 provides a more detailed assessment of this pattern. Here, each community is represented as a dot overlain on box plots that describe the distribution of a variable within each group. Category 1 includes family-dominated groups, which are especially strongly distinguished by their elevated level of family ownership, 2 represents uncorroborated communities with low scores across all three dimensions and as well as slightly less geographic homogeneity, and groups in the third category are distinguished by violating the general ten-

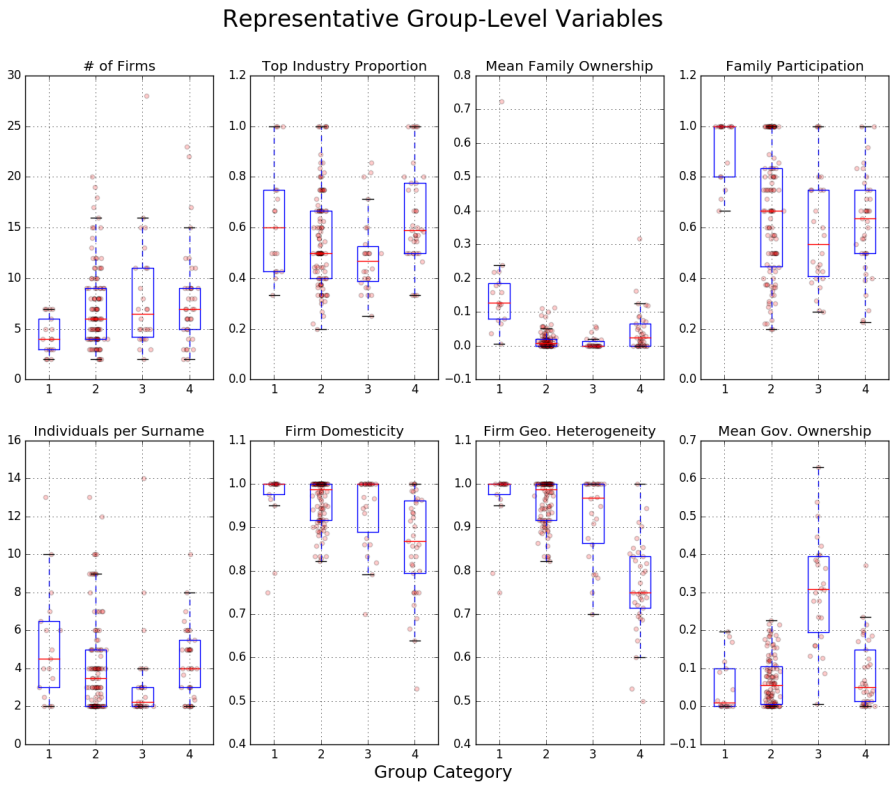

Figure 2.6

dency toward geographic homogeneity within inferred communities, perhaps indicating a group of firms operating in different markets but bound by a common origin . Finally, the fourth type corresponds to government-owned groups and is characterized by high government ownership.

2.5 Examples

The correspondence between the results of the GMM analysis and the three observed axes of solidarity helps to establish the relevance of the community-level measures. Without observed group labels, however, it is not possible to establish its overall accuracy. Instead, this section presents two examples from each

category that show some of the important patterns in the data. The visualizations accompanying each example show all nodes within a community as well as adjacent external nodes. Intra-community edges are darkened.

Beginning with family firms, Figure 2.7 shows the Jordanian Abo Khadiga business group. All seven firms in this community have at least one of ten family members as an owner or director, and the family owns on average 15.9% of each company. This results in an aggregate of family score of 1.8 standard deviations above the mean for all communities. However, the results are not always so clear, and countries like Saudi Arabia and Qatar present a special challenge for the method due to the extraordinary influence of a small number of families. Figure 2.8 shows a community of Saudi firms that feature eight individuals from the al-Subaie and al-Serafi families, but this high number of individuals and overall participation is offset by a relatively low level of ownership yielding an aggregate family z-score of 0.85. Furthermore, members of the al-Subaie family also appear in other groups, as can be seen in the three interconnected green-colored nodes.

Moving on to communities of government-owned firms, Figure 2.9 shows group of seven publicly traded Omani firms with strong signs of government control. State organizations, such as the Ministry of Finance and the Civil Service Employees Pension Fund, own 37.2% of the entire group, yielding an aggregate government z-score of 2.1 after taking into account variation across companies. Other communities reflect more ambiguous situations. In the case of the group of four UAE companies shown in Figure 2.10, a relatively low level of overall ownership of 15.6% is offset by the fact that the two banks in the group are closely tied to the Emir of Sharjah, Mohamed Al-Qassemi, although this is

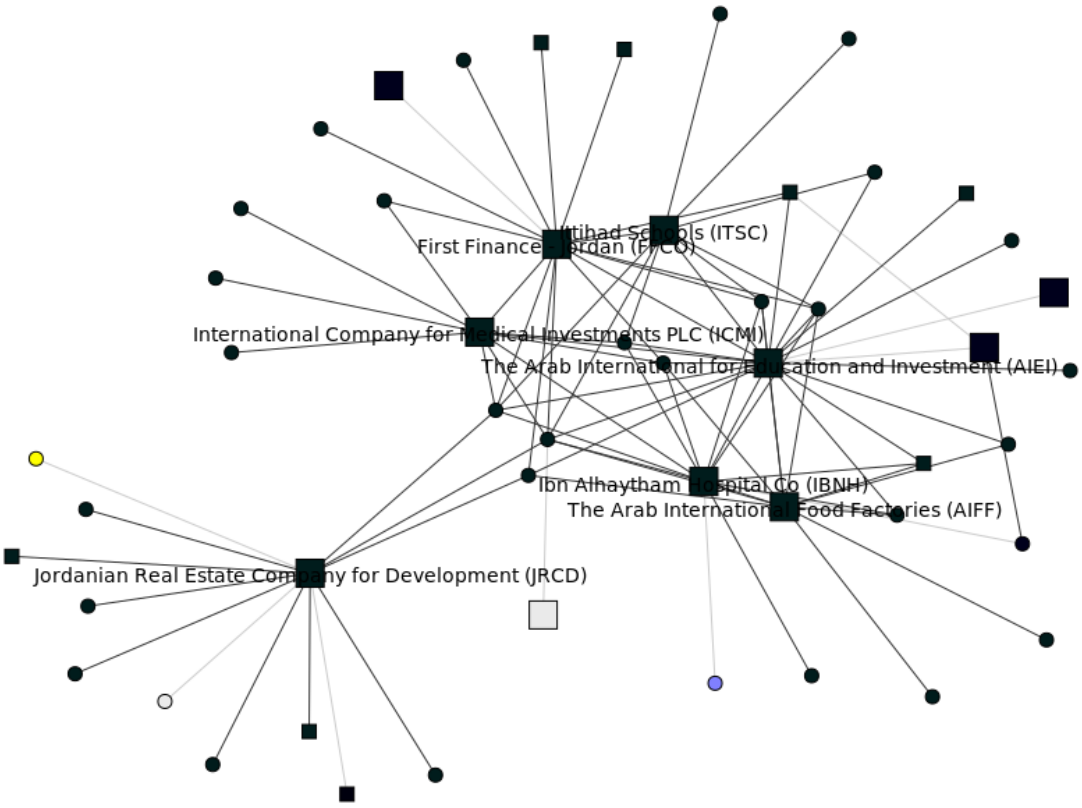

Figure 2.7: The Jordanian Abo Khadiga business group. Note the density of internal ties.

not reflected in a lower government score of 0.88.

Many communities of government-owned firms also have a transnational character. Though assigned to the cluster of the transnational communities by the GMM, the group shown in Figure 2.11, is largely owned—along with several other firms visible in the picture—by the Bahraini Social Insurance Organization, yielding a government z-score of 2.4. Its high geographic heterogeneity z-score of 2.1 stems from the subsidiary telecommunications company listed on the Saudi Exchange. Figure 2.12 provides an example of a group of mainly Kuwaiti firms that are transnational based on their extensive investments in other countries, yielding an aggregate z-score of 2.0.

Rather than containing industry-based groups, the final cluster is composed

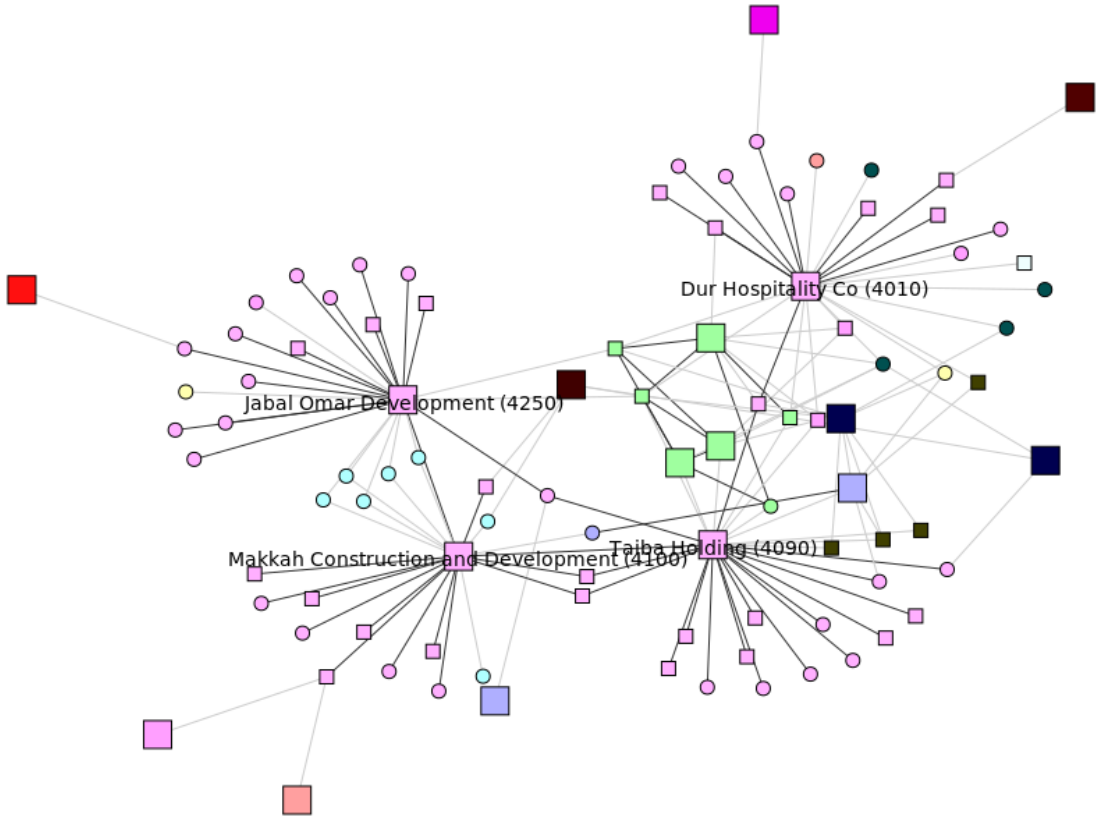

Figure 2.8: A set of Saudi firms with strong involvement by the al-Subaie and al-Serafi families. Note the interconnected group of green firms adjacent to this community, which also feature heavy involvement by the al-Subaie family.

mostly of communities with low scores across all four categories, but there are exceptions. Figure 2.13, for example, shows a dense group of Palestinian firms connected with the al-Masri family with each firm associated with at least one of seven different individuals, yielding an overall ownership of 10% and an aggregate family z-score of 1.1. The group of four Tunisian firms in Figure 2.14 is in a similar situation with regard to government ownership with a z-score of 1.4.

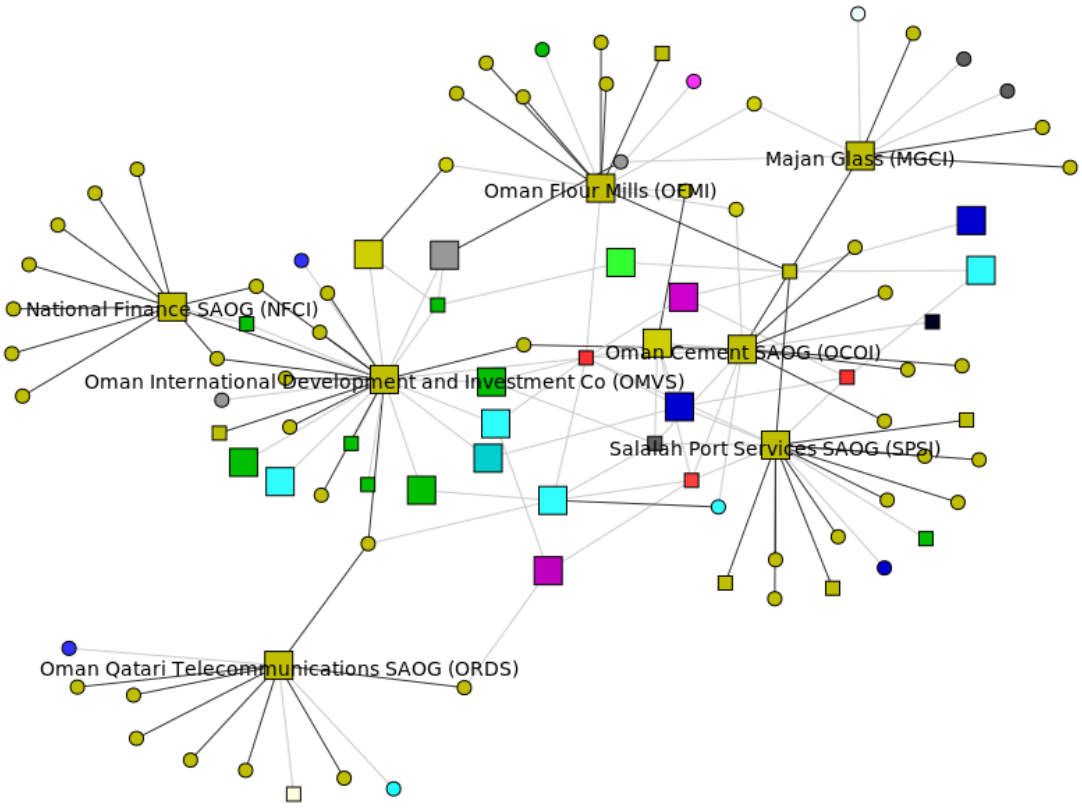

Figure 2.9: A group of Omani firms with high-government ownership.

2.6 Discussion

The results in this paper rely on publicly available information on corporate networks to make inferences about business group structures, which are otherwise difficult to observe. The above analysis of community properties and brief examples show that the community detection algorithm does succeed in identifying sets of firms that meet a broad definition of business groups. More specifically, by measuring the aggregate properties of inferred communities we can identify socially meaningful groups.

Nevertheless, there is room for improvement. In order to provide accurate information on the relative prevalence of business groups it is necessary to bet-

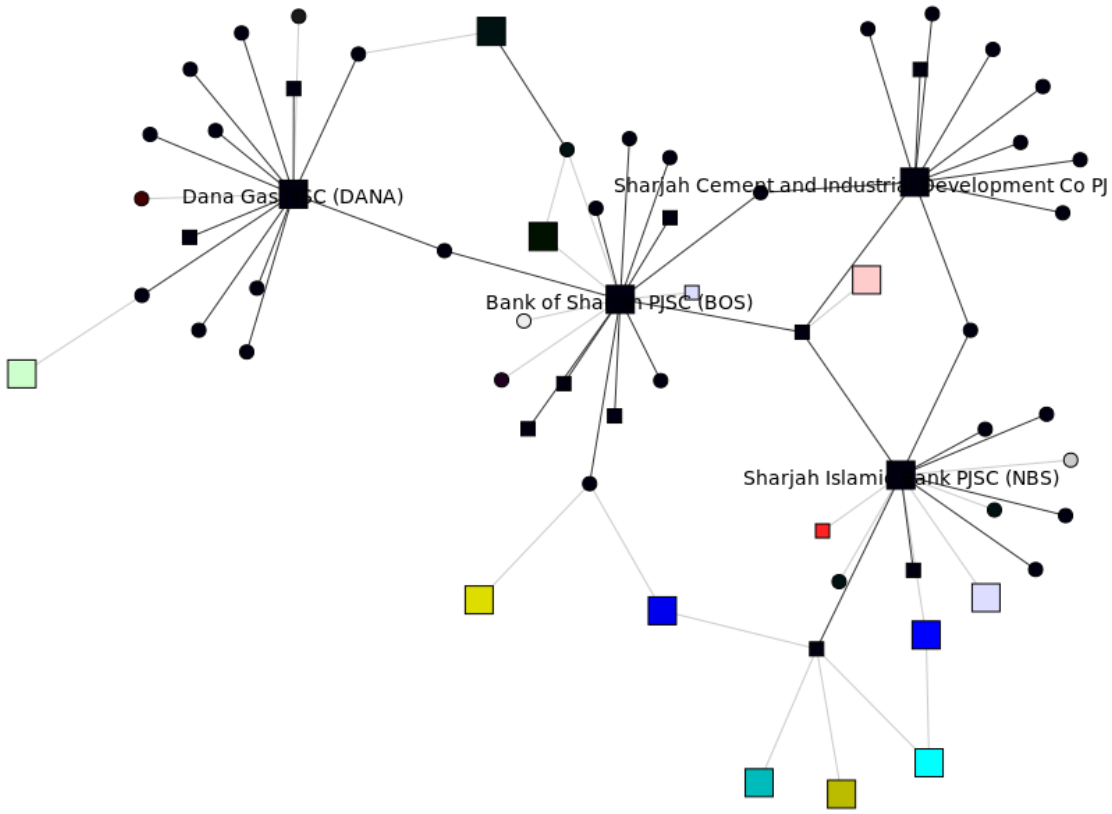

Figure 2.10: An Emirati business group with close ties to the Emirate of Sharjah, one of the constituent states of the federal UAE.

ter distinguish member and non-member firms. This is frustrated by the fact that in the current method the combination of limited data and degenerate results of the modularity algorithm might prevent the identification of all groups. Moreover, there are many other possible sources of solidarity, which might distinguish further communities if they were observed. For example, membership in a religious or ethnic minority could provide also strong basis for cohesion.

Two general strategies could mitigate these concerns. First, incorporating more data could improve the accuracy of the algorithm while also giving a better window into the social relevance of each group. This paper uses information on location, family, and government affiliation as means of checking the validity of patterns derived from ownership and director-interlock networks. However,

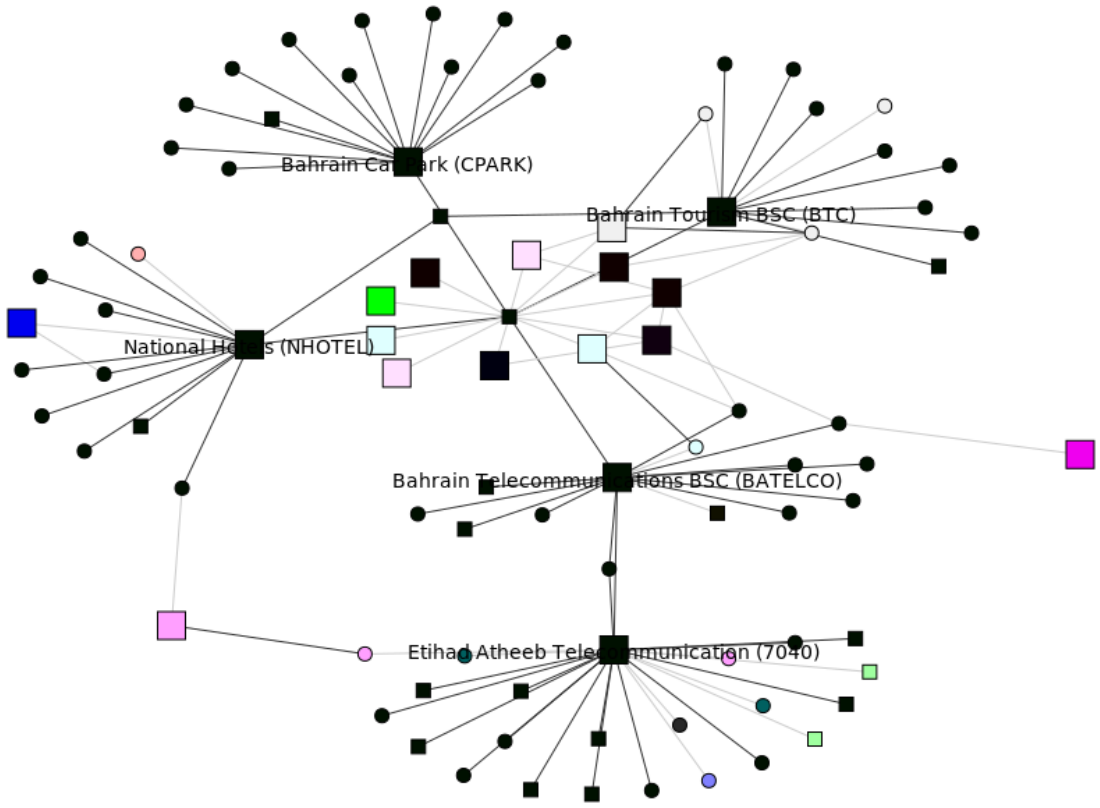

Figure 2.11: A group owned by the Bahrani government with a Saudi subsidiary.

these concepts can also be represented as bipartite networks where the owner and director nodes are linked by common affiliation with a government, surname, or country. Hence, they could be incorporated into the multiplex algorithm in order to model our knowledge of these structures. Furthermore, it is also possible to construct a third corporate network using data on stock price comovement, and Khanna and Thomas (2009) have shown that belonging to the same group is a significant predictor of price correlation, even when also accounting for direct ownership and director-interlock relations. Thus, a further means of confirming the salience of a set of communities is to test their ability to predict similarly priced firms. This weighted undirected network could also serve as an input for the community detection algorithm. This analysis is the subject of the second chapter of this dissertation.

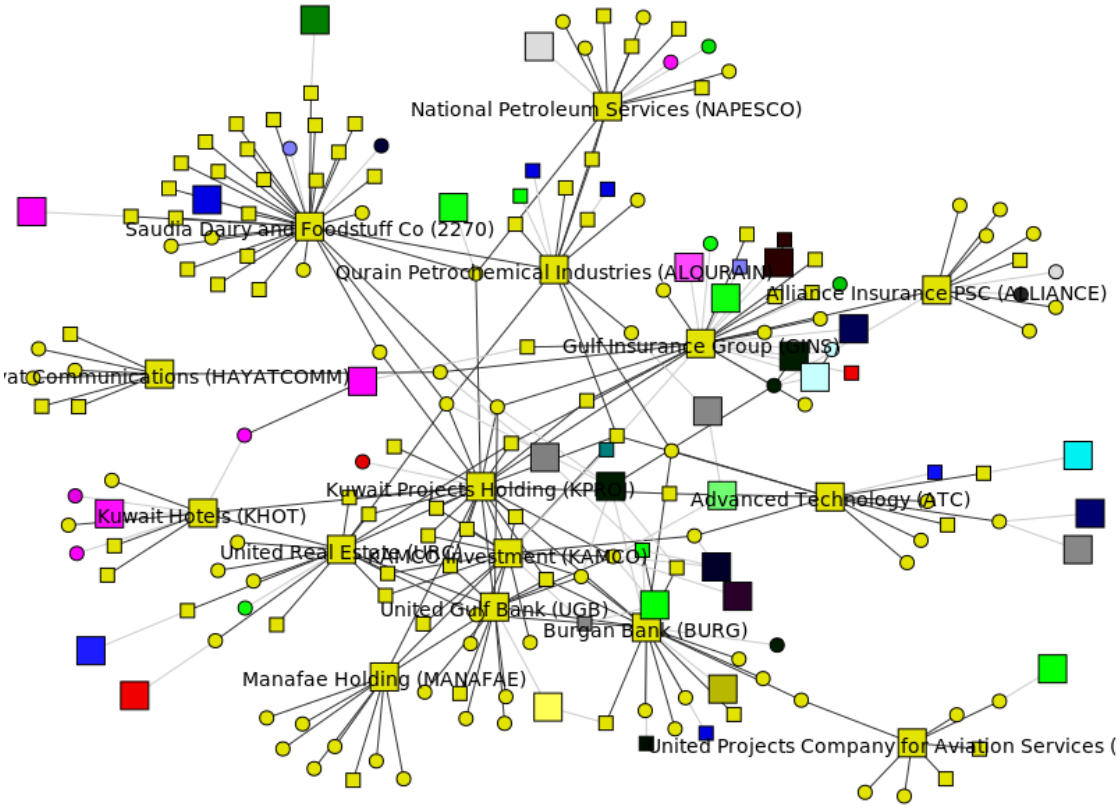

Figure 2.12: A dense core of Kuwaiti firms surrounded by geographically diverse affiliates.

2.7 Conclusion

Business groups are sets of firms bound by an intermediate level of formal and informal connections that coordinate their actions by means of several forms of intra-group solidarity. Using this definition as a guide, this paper has employed a community detection algorithm and a Gaussian Mixture Model to learn about the latent structures in the network of corporate ownership and director-interlock networks in the Middle East and North Africa. The results show that using a multiplex modularity-based algorithm developed by Mucha et al. (2010), it is possible to identify 50-75 sets of firms with strong family, government, or transnational ties.

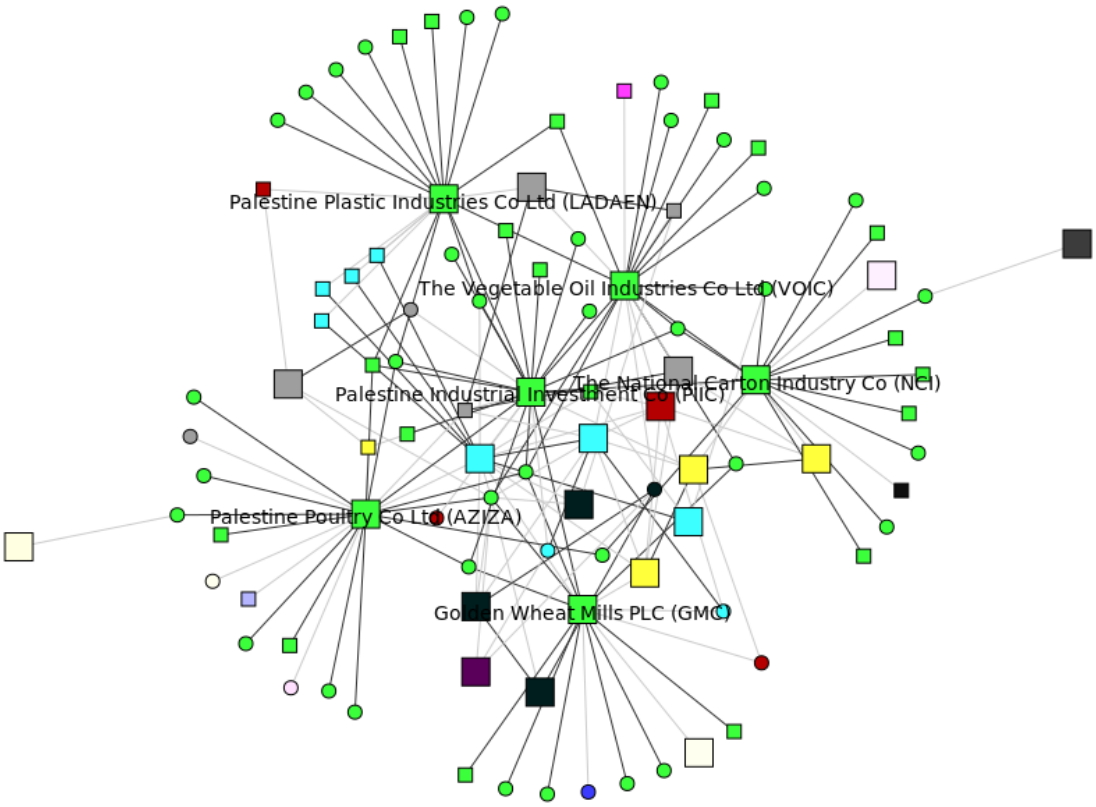

Figure 2.13: The Palestinian al-Masri business group.

The implications of this work are twofold. First, most studies of business groups have been conducted in a small number of countries. Using network analysis to infer these groups reduces the need for extensive field work or to pay for expensive supplementary data. Especially after implementing the refinements discussed in the previous section, using community information as a proxy for directly observed business groups would allow for analyses across a wider range of times and locations. Second, the potential for using computational methods to increase transparency extends beyond business groups, and as such this paper contributes to emerging work on using machine learning and other methods to uncover consequential economic, political, and social structures (Hicks et al. 2015).

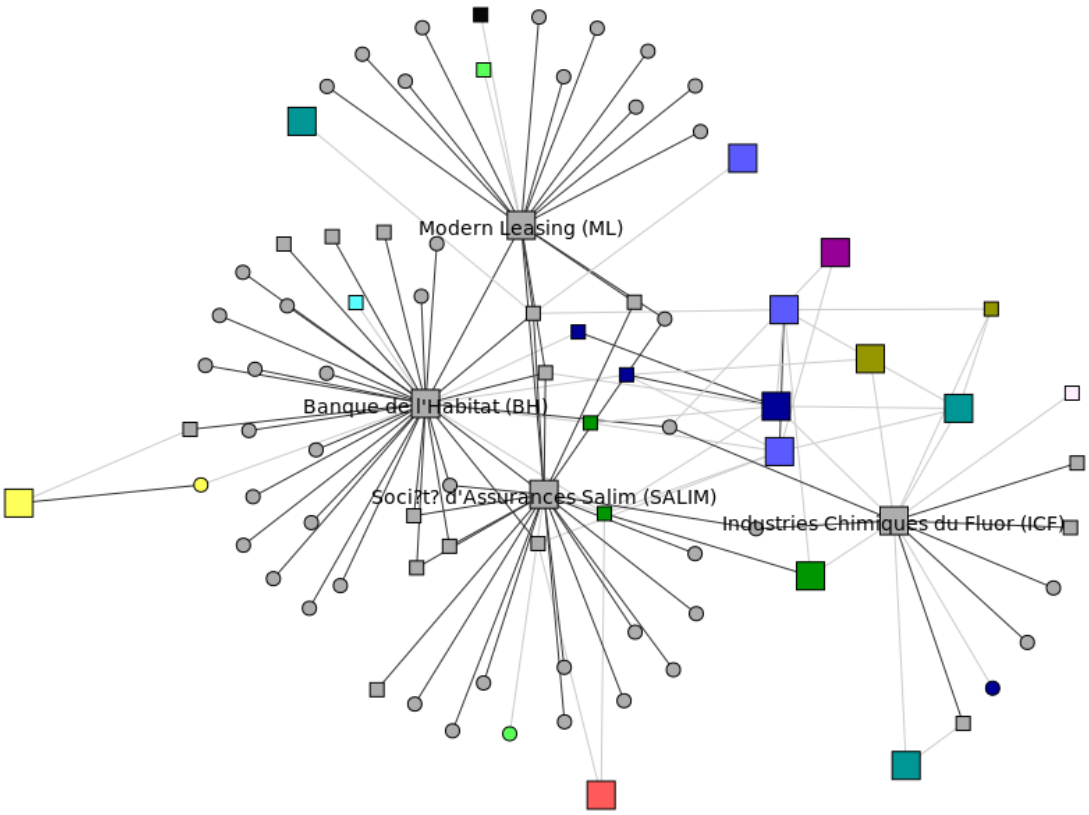

Figure 2.14: A group of Tunisian firms with significant state ownership.

CHAPTER 3

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

3.1 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 proliferation of research on business groups, few studies have investigated their role in regions like the Middle East and North Africa. Lack of available data is one potential reason for this neglect, and I am not aware of any public data on group membership for Arab countries in the region[1], and in previous work I show how community detection techniques can be used to infer business group membership (Siemon 2017). This paper extends that approach and further verifies its results by providing evidence that these latent groups influence the comovement of stock prices among public firms in the Middle East and North Africa.

Studying the price mechanism in developing stock exchanges is essential because financial markets are useful only insofar as they improve the matching of investment opportunities with the capital required for their realization. Economic theory suggests that in equity markets, which provide a forum for the buying and selling of ownership stakes in large corporations, the price of a firm's shares should reflect market participants' expectations about the value of the future cash-flow to which the ownership entitles them (Wurgler 2000).

[1]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 that may either be purchased or obtained from government agencies.

While country- and industry-specific environments certainly inform these expectations, studies have found that these factors are overrepresented in firm prices beyond what might be expected from their influence on other measures of performance (Pindyck and Rotemberg 1993). In other words, stock prices are often synchronous in that a large part of variation over time can be explained by the overall movement of an entire exchange or within a particular industry (Roll 1988; Morck et al. 2000). Taken at face value, such comovement implies that investors have limited access to firm-specific information and thus rely on broader knowledge of a firm's environment when deciding whether to buy or sell a stock at a given price.

Studies of synchronicity in financial markets typically focus on the comovement of individual firms with a relevant exchange or industry index. This generates an informative measure of the sensitivity of each firm's price to its aggregate environment but necessarily ignores other factors that might make some firms more synchronous with one another. For example, numerous studies have found that political connections are related to distinctive price movements in response to politically charged events (Fisman 2001; Faccio 2006; Leuz and Oberholzer-Gee 2006; Chekir and Diwan 2014). Rather than studying synchronicity at the firm- or exchange-level, this paper adopts the approach of Khanna and Thomas (2009), who disaggregated synchronicity into an analysis of pairwise correlations in the returns of listed firms in Chile, in order to study the value relevance of inter-firm relationships and business groups in thirteen exchanges in the Middle East and North Africa.

By combining price data with directly observed and inferred measures of ties among listed firms in the region, this study provides further evidence on the

link between inter-firm networks and price similarity. In doing so, it also confirms the validity of using community detection methods from network analysis to identify meaningful business groups by comparing the inferred clusters with an independently generated but related data set. Furthermore, it uses a Bayesian multilevel modeling framework to assess how the strength of this association varies between exchanges by measuring the substantive importance of family business group ties and government ownership in determining firm value in different contexts. Finally, although the number of exchanges in the region limits the strength of country-level inferences, it is also useful to inspect the relationships among this type of variation and important exchange-level measures, and the analysis concludes by presenting evidence on the relationship between estimated coefficients and liquidity, regulatory quality, and other variables.

The paper is structured as follows. Section 2 reviews the relevant literature and presents three hypotheses. Section 3 describes the data. Section 4 describes the model and methodology. Section 5 presents results, and Section 6 concludes.

3.2 Literature Review and Hypotheses

The goal of this paper is to provide evidence pertaining to interdependence among stock-price comovement, inter-firm relationships, and business groups. This section describes the relevant literature on these topics and establishes the rationale for three hypotheses relating to price synchronicity and business groups. As noted in the introduction, I do not observe business groups directly but instead infer them from clusters of observed network ties. This section fo-

cuses only on the relation between synchronicity and inter-firm relationships. A discussion of the literature linking network clusters to business groups is presented in Chapter 2.

3.2.1 Synchronicity, R^2, and Firm-Specific Information

Neoliberal reforms in the Middle East and North Africa have been justified by the ability of market-mechanisms to reduce the inefficiencies that arise from politically motivated economic decisions, and no mechanism is more central to market-based systems than price. In the case of equity markets, determining the value of a given company is complex but at a fundamental level is shaped by market participants' expectations about the future profits of the firm to which the owners will have a claim (Fama and Jensen 1985). This expectation in turn depends on information relevant to the firm's future performance, as well as investors' ability to collect their due profits, and the efficiency of financial markets is correlated with both of these factors (Wurgler 2000; Durnev et al. 2004). In other words, investors with better information are able to make better decisions about where to invest and will do so more freely when they feel confident about receiving a return.

Directly measuring the quality of information, however, is often impractical, and the synchronicity literature instead approaches the issue by decomposing the temporal variation in stock prices into market-, industry-, and firm-specific components (Roll 1988). This is achieved by first estimating a linear regression for each firm where the dependent variable is the firm's change in price during a given time period. The predictors are an intercept and the average change in

price of one or more indices of stocks listed on the same exchange or operating the the same industry. The proportion of explained variance, or R^2, under these simple models is then taken as a measure of synchronicity, or the dependence of that firm's price on the price of other firms. The variance of the residual is interpreted as the firm-specific idiosyncrasy.

The accumulation of results linking R^2 synchronicity to theoretically important phenomena has supported its widespread adoption. For example, Wurgler (2000) and Morck et al. (2000) provide evidence of its association with efficient capital allocation and protections for minority shareholders, respectively. Specific evidence linking R^2 to the frequency of informative events has been provided by Piotroski and Roulstone (2004), Chan and Hameed (2006), and An and Zhang (2013), and these studies suggest that trades by firm insiders and dedicated, long-term institutional investors have the expected effect of increasing the firm-specific variance, while greater analyst coverage increases the dependence on industry or market trends. This latter result reflects the fact that analysts tend to specialize in a particular area and hence have a comparative advantage in assessing the implications of industry-wide events. A complementary set of analyses investigate the role transparency and financial disclosures and again find the expected relationship; countries and firms with more opaque reporting practices tend to have greater price synchronicity (Jin and Myers 2006; Haggard et al. 2008; Hutton et al. 2009).

Even with these results, synchronicity remains an indirect proxy for firm-specific information, and it is not always clear to what extent firm-specific variation reflects meaningful information rather than uninformative noise (De Long et al. 1990; Shleifer and Vishny 1997a). Although Durnev et al. (2003) address

this issue with data on US firms, by demonstrating an association between low synchronicity and the ability of current returns to predict future earnings, other countries might have extremely poor information environments characterized by both low synchronicity and a weak connection between firm fundamentals and share prices. Nevertheless, the important point for this paper is that both firms and exchanges vary substantially in the types of information reflected by their stock prices, and that this variation can be observed by analyzing the interdependence between different time-series of returns. For example, Figure 3.1 shows the stock prices of all Saudi and Jordanian firms as a percent of their mean value during the 2015 calendar year. The Jordanian firms have little common pattern besides a slight decrease, but the Saudi exchanges shows a strong tendency for prices to move in unison. The corresponding R^2 values for this time period are 0.41 for Saudi Arabia and 0.05 for Jordan[2].

3.2.2 Inter-firm Relationships and the Price Mechanism

This study differs from the firm- and country-level studies noted above in that it focuses on synchronicity between pairs of companies, as opposed to comovement between one firm and the rest of the market. As such, this section turns to the literature on firm relationships and stock-prices in order to motivate hypotheses about synchronicity within business groups, which for the purposes of this study are defined as collections of firms which share a variety for formal and informal connections reinforced by an underlying source of solidarity (Gra-

[2]The R^2 values were calculated following Morck et al. (2000) by regressing the daily percent change for each firm on the market by the mean percent change for all firms on that market and then calculating the average of all firm-specific R^2 values weighted by the variance in their returns. Because some exchanges in the region are dominated by a small number of highly valued equities, I do not weight the mean market return by firm value.

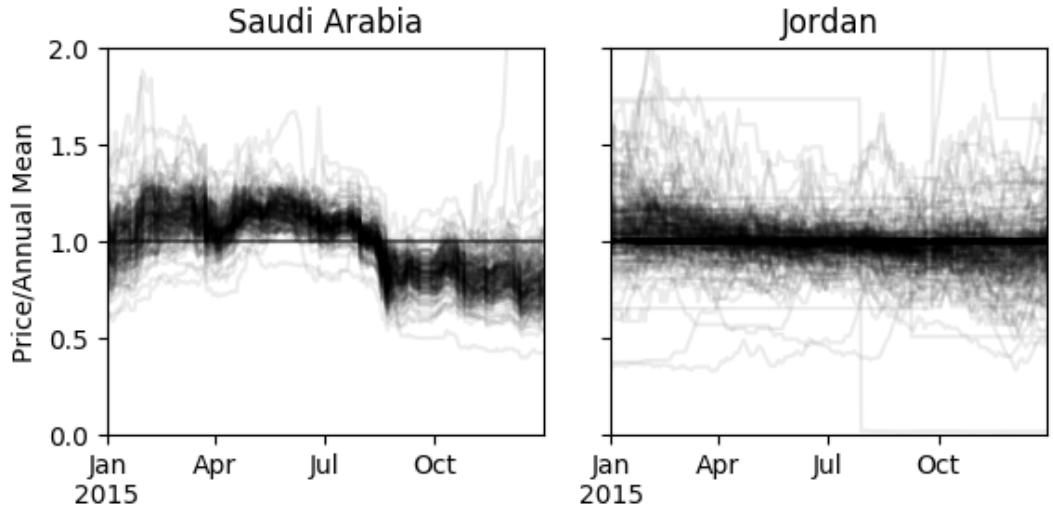

Figure 3.1: Prices of Saudi and Jordanian firms in 2015 plotted relative to each firm's mean price for the entire year.

novetter 2010). See Chapter 2 for a more complete discussion of this definition.

The idea that firms within the same business group should 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 either case entails 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 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 in conflict with the logic of a public market for company 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 most listed companies in the world, and studies of corporate governance in emerging markets tend to 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 synchronicity and whereby the positive association between the largest owner's share and synchronicity begins to decrease once that share passes around 50%. 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.

3.2.3 Stock Markets in the Middle East and North Africa

The Middle East and North Africa are an attractive area for studying the associations between networks, business groups, and stock prices for three interrelated 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 recent work by 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 Opper 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 in-

formation 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 countries[3], 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 non-financial 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

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

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 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.

Finally, another set of studies focuses on the associations between different governance structures and directly observed measures of firm performance in addition to 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 the firm pays on its debt, and find that creditors appear to be 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-dominated firms in Tunisia are more productive, but only in less-competitive sectors of the economy.

Taken together, these studies are suggestive of the underlying 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 manifest 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). As described in Chapter 2, there is evidence that an identifiable subset of these network communities can serve as proxies for the business group memberships of listed firms. 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 importance 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 the market price. 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 prices in the region do depend at

least somewhat on firm-specific information.

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 algorithm from Chapter 2 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 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, chairmen, and top management as follows.

Hypothesis 3: *Firms that are members of the same inferred network community and both have at least one owner, chairman, or top manager from that community's domi-*

nant family have greater pairwise sychronicity than that associated with their directly observed shared 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.

3.3 Data

This study investigates the association between inter-firm relationships and price synchronicity using data from 1373 publicly traded firms on 13 stock exchanges in 12 countries in the Middle East and North Africa. Of these firms, 1110 have sufficient non-missing data. The units of analysis are all pairs of firms listed on the same exchange. Including cross-listings, this yields a total of 71,677 observations.

3.3.1 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 year[4]. I calculated each firm's daily

[4]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

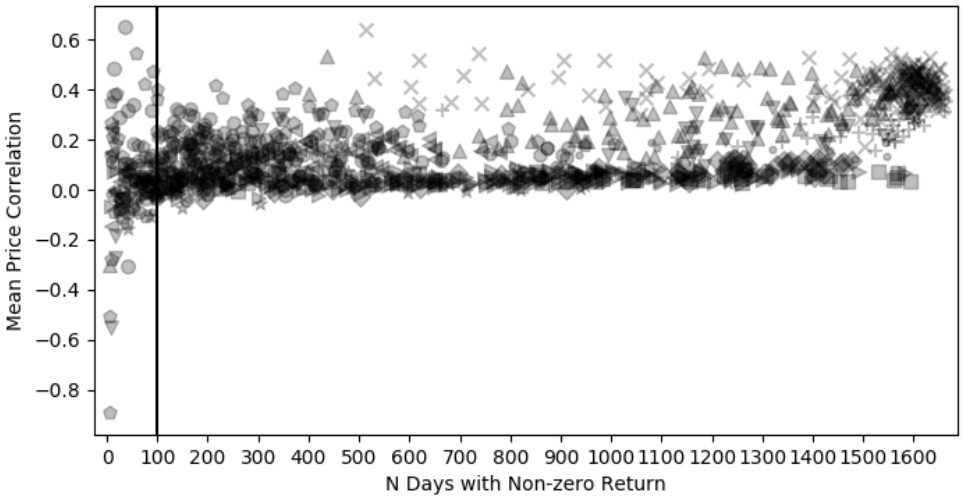

Figure 3.2: The number of days with non-zero returns for each firm plotted against its mean price correlation with other firms on the same exchange. Note the extreme mean values for firms with fewer than 100 days of non-zero returns. Different exchanges are indicated by shapes to distinguish the influence of small sample size and high exchange-level synchronicity.

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 firms with at least 100 non-zero returns. This accords with the finding that non-zero returns are a convenient

exchange, and estimating $N - 1$ parameters from fewer than N observations will necessarily produce noisy estimates.

proxy for liquidity in the region (Hearn 2014). I chose the cutoff of 100 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 firms with abnormally high or low values[5]. Figure 3.2 shows a plot of these values with firms on different exchanges represented by different shapes.

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 transformation[6]. The resulting data have a higher proportion of extreme values than would be expected from a normal distribution. Figure 3.3 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.

[5] 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. For example, while a correlation measure based on 10 observations is obviously much less reliable than one based on 110, it is less clear how much preference to show between 300 and 400 or 1000 and 1500.

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

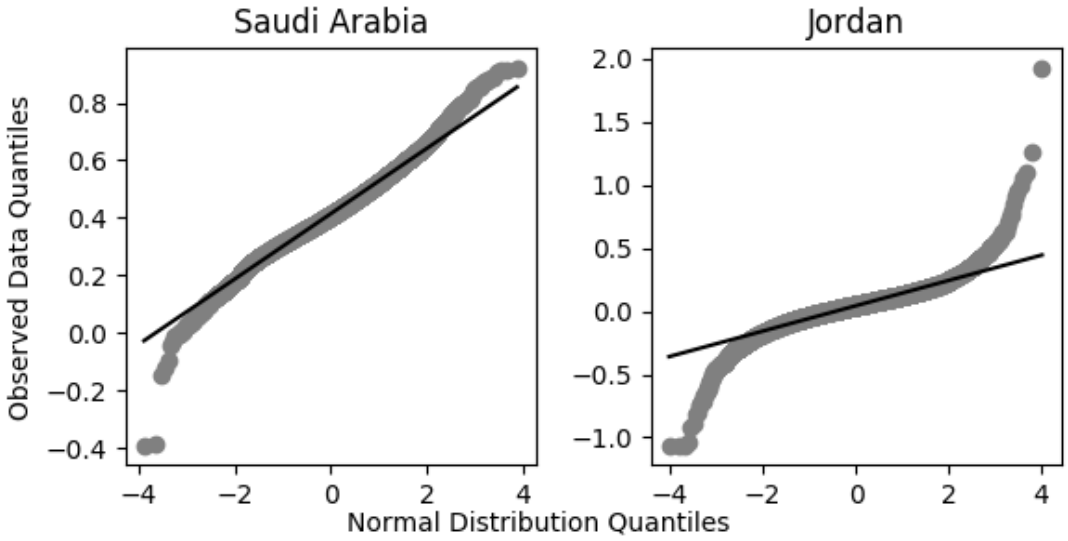

Figure 3.3: Quantile-quantile plots for Jordan and Saudi Arabia demonstrate the high proportion of extreme values in the data relative to a normal distribution.

3.3.2 Observed Network Ties

The measures of inter-firm relationships are based on lists of the directors, managers, and owners of each firm in 2015[7]. These data were taken from firm profiles on mubasher.com, which has been used as a source in other studies of public firms in the region (e.g. Hearn 2014). Individual names appear in Arabic and as transliterations into Latin characters, and company names are similarly available in Arabic and as either English or French translations. In order to construct measures of inter-firm relationships, names were cleaned and matched across occurrences in different firm profiles using an adaptation of the algorithm described in Colomer (2012). In general, I gave more weight to Arabic company names and ignored transliterated individual names unless no Arabic name was available. The core of the matching algorithm assigns scores to different types of similarity between two words and aggregates these values across all words in

[7]The analysis thus compares observations from a single year to correlations based on a seven-year window, though I believe this is justified not only because of the need for more non-zero return observations but also because the ties under consideration are relatively stable over time.

a name. Names with a score above a threshold are considered a match. Chapter 2 provides full details on this procedure.

Next, I used the matched names to construct four types of networks, which are defined as a set of vertices linked by edges. The two principle graphs 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.

I supplemented these direct representations of ownership control relationships with information on the surnames of individuals and the government ownership of non-publicly traded firms, creating two additional bipartite networks representing possible family and government ties, respectively. The first is composed of surnames linked to company or individual names where they appear. The second includes non-public firms, which were hand-coded as state-owned and linked to a rough proxy for their government. I coded firms based on their names, as well as through online information on zawya.com and bloomberg.com. I created one government node for each exchange and linked

owners to the government where they owned the greatest number of firms[8].

3.3.3 Community Detection

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. Chapter 2 provides full details on the algorithm.

[8]This is an over-simplification for some cases: for example, in the UAE government authority is split among seven federated states.

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, by taking the communities themselves as a unit of analysis, I distinguished family- and government-dominated communities from those lacking a clear basis of affiliation, and in the previous chapter this is done without using the supplemental family and government information as an input to the algorithm (Siemon 2017). Finding such distinctive communities thus proves that at least some of the groups are meaningful. Second, adding data that reflects expectations about where groups might be found should improve the reliability of the method. Although the supplemental family and government networks contain a limited amount of structural information, when included as slices of the multiplex structure they provide greater certainty to the algorithm if they correspond to similarly located concentrations of director-interlock and ownership ties; and are easily ignored if they do not. Finally, testing the results against independently generated but related data would provide perhaps the strongest verifica-

tion, short of comparing them against true group labels. While family and government measures are useful insofar as they leverage extra information about the specific attributes of vertices in the network, this information is still derived from the same networks that were used to infer the communities. Pairwise price correlations, on the other hand, provide a completely separate weighted network with which to compare the results.

3.3.4 Variable Construction

In order to test the hypotheses described above, I use a measure derived from the community detection analysis, as well as direct observations of joint government ownership. The latter is defined as the geometric mean of the percent of each firm's shares owned by government-affiliated organizations. 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 chairman, deputy chairman, 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.

Exchange	Abu Dhabi	Bahrain	Dubai	Egypt	Iraq	Jordan	Kuwait	Morocco	Oman	Palestine	Qatar	Saudi Arabia	Tunisia	Entire Sample
N Firms	52	20	44	130	64	218	168	68	61	37	43	167	60	1110
N Dyads	1176	153	982	8367	2013	23616	14176	2210	1826	664	903	13821	1770	71677
Mean Price Correlation	0.083	0.081	0.235	0.312	0.154	0.042	0.144	0.044	0.252	0.036	0.236	0.413	0.073	0.181
N Shared Directors	97	13	47	110	52	425	239	117	104	72	249	423	63	2011
Shared Dir. Density	0.082	0.084	0.047	0.013	0.025	0.017	0.016	0.052	0.056	0.108	0.275	0.030	0.035	0.028
N Direct Own. Ties	4	7	13	25	41	153	89	32	27	41	11	51	55	549
Direct Own. Density	0.003	0.045	0.013	0.002	0.020	0.006	0.006	0.014	0.014	0.061	0.012	0.003	0.031	0.007
N Shared Own. Ties	92	41	32	2217	1106	1640	580	325	437	164	177	1398	226	8435
Shared Own. Density	0.078	0.267	0.032	0.264	0.549	0.069	0.040	0.147	0.239	0.246	0.196	0.101	0.127	0.117
Mean Shared Ownership	0.012	0.039	0.006	0.010	0.012	0.006	0.006	0.015	0.029	0.027	0.034	0.011	0.017	0.009
N Family Group Group Dyads	13	1	4	7	10	96	43	6	19	25	47	68	8	347
Family Group Density	0.011	0.006	0.004	0.000	0.004	0.004	0.003	0.002	0.010	0.037	0.052	0.004	0.004	0.004
Mean Shared Gov. Own.	0.081	0.133	0.023	0.083	0.031	0.008	0.008	0.028	0.085	0.009	0.096	0.027	0.028	0.027
N Same 2-Digit SIC	127	55	125	522	233	1665	1547	130	101	42	75	1076	186	5884
N Same 3-Digit SIC	109	55	93	210	167	992	824	54	49	33	54	779	106	3525

Table 3.1: Mean and count statistics for observations with sufficient price data for analysis are presented by exchange and the entire sample.

3.3.5 Data Description

Table 3.1 presents a summary of the data at the exchange level as well as for the entire sample. It includes only dyads involving the 1128 firms with 100 or more days of non-zero returns during the 2010-2016 period. Of the 1373 public firms with information on mubasher.com, I matched 1293 with one or more Datastream records. Of these 80 unmatched firms, 63 were listed on the Egyptian Exchange with no other exchange having more than 5. The number of firms on each exchange is thus a product of its size as well as liquidity. The sum of firms listed on each exchange is slightly lower than the total number of firms because of cross-listings. The mean price correlations also vary substantially by exchange. These quantities mirror the R^2 measures described above for the cases of Jordan and Saudi Arabia, and more data will be presented about the relationships between exchange-level characteristics and the dyadic analysis in the next section. The density measures refer to the number of observed connections on an exchange divided by the total number of possible dyads.

Table 3.2 describes the dropped observations and compares them to all dyads involving firms that were matched with a Datastream record. Excluded data are most significant for Bahrain and Oman, but as noted above, Egypt also had a disproportionate number of firms that were not covered at all by Datastream and hence not added to the overall data set. In comparison with the entire sample, excluded observations tend to have lower mean price correlation and have higher standard deviations among the correlation measures, which reflects the greater noise among the dropped data. Relational measures are broadly similar, with the key exception of family group dyads. A majority of these ties are excluded in Oman, Dubai, Abu Dhabi, and Egypt. An alternative approach

would be to keep these data and compensate for their greater noise by weighting them according to a function of the inverse of the number of non-zero daily returns. Doing so produces results with similar magnitude, but greater uncertainty compared with those presented below. I plan to explore this approach more in a future version of this paper. In the meantime, I prefer simply to exclude them on the grounds that prices with fewer than 100 non-zero returns out of over 1600 trading days are not likely to be very informative. Nevertheless, 100 is an arbitrary cutoff, and choosing an appropriate function for generating sample weights would allow for more flexibility than the binary choice of inclusion or exclusion.

3.4 Methods

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 distributions given the data and prior expectations about their likely values.

Exchange		Abu Dhabi	Bahrain	Dubai	Egypt	Iraq	Jordan	Kuwait	Morocco	Oman	Palestine	Qatar	Saudi Arabia	Tunisia	Entire Sample
N Firms	All	65	41	64	151	82	230	201	68	112	45	43	168	60	1330
	Dropped	16	23	19	21	18	12	32	1	51	8	0	1	0	202
Mean Price	All	0.067	0.091	0.173	0.301	0.123	0.040	0.123	0.042	0.191	0.020	0.236	0.413	0.073	0.165
Correlation	Dropped	0.035	0.099	-0.035	-0.126	0.027	0.004	0.050	-0.041	0.112	-0.023	0	0	0	0.040
Std. Dev. of	All	0.218	0.431	0.293	0.209	0.212	0.132	0.292	0.090	0.340	0.231	0.114	0.113	0.079	0.246
Price Corr.	Dropped	0.312	0.518	0.372	0.332	0.336	0.337	0.418	0.320	0.425	0.387	0	0	0	0.395
N Dyads	All	1760	346	1274	8586	2661	25253	18218	2272	3212	898	903	13821	1770	80974
	Dropped	584	193	292	219	648	1637	4042	62	1386	234	0	0	0	9297
Interlock	All	0.076	0.075	0.041	0.012	0.018	0.017	0.015	0.051	0.042	0.094	0.275	0.030	0.035	0.027
Density	Dropped	0.068	0.073	0.035	0.011	0.007	0.014	0.012	0	0.036	0.067	0	0.024	0	0.023
Direct Own.	All	0.003	0.020	0.009	0.003	0.025	0.006	0.006	0.014	0.010	0.059	0.012	0.003	0.031	0.007
Density	Dropped	0.004	0.014	0.005	0.003	0.034	0.004	0.005	0	0.008	0.055	0	0	0	0.008
Shared Own.	All	0.081	0.151	0.031	0.219	0.486	0.069	0.037	0.147	0.156	0.194	0.196	0.099	0.127	0.112
Density	Dropped	0.085	0.124	0.030	0.090	0.390	0.065	0.027	0.161	0.122	0.088	0	0.004	0	0.092
N Family	All	28	4	17	20	10	101	49	6	73	34	47	68	8	465
Group Dyads	Dropped	15	3	13	13	0	5	6	0	54	9	0	0	0	118
Mean Gov.	All	0.080	0.091	0.020	0.069	0.026	0.008	0.007	0.027	0.058	0.007	0.096	0.027	0.028	0.027
Ownership	Dropped	0.079	0.081	0.017	0.031	0.017	0.012	0.005	0	0.046	0.001	0	0.004	0	0.026

Table 3.2: Adjacent columns compare statistics for the entire data set versus the observations that were dropped from the analysis.

3.4.1 Bayesian Analysis

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. 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 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 carefully 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 difficulties 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.

3.4.2 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 data 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. For example, firm-level intercepts .

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 the Fisher-transformed stock-price correlations are a continuous measure of firm similarity, they can be predicted with a linear regression model. However, as noted in the previous 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:

$$y_i = \boldsymbol{\beta}\boldsymbol{x}_i^T + \boldsymbol{\gamma}_j\boldsymbol{z}_i^T + \delta_k + \delta_l + \epsilon_i \tag{3.1}$$

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; $\boldsymbol{x}_i$ is a vector containing the variables whose associations with price comovement does not vary by exchange in the model; and $\boldsymbol{z}_i$ contains the other variables whose coefficients are allowed to vary. I chose to constrain direct ownership to be constant across exchanges because it has the straightforward economic implication that one firm is entitled to a share of profits of another and thus contrasts with the more ambiguous relationships implied by shared ownership or business group comembership. 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. Both cases, however, were foremost a practical decision since each additional coefficient that is allowed to vary by group significantly increases the computation time. Their variation is simply less interesting in this context than the other variables, but I do plan to test alternative specifications in a future version of this paper.

Returning to the model, $\boldsymbol{\beta}$ is a vector of group-invariant coefficients and the other parameters of the model are defined as follows:

$$\boldsymbol{\gamma}_j \sim N(\boldsymbol{\mu}_\gamma, \boldsymbol{\Sigma}_\gamma) \tag{3.2}$$

$$\delta_k \sim N(0, \sigma_\delta^2) \tag{3.3}$$

$$\delta_l \sim N(0, \sigma_\delta^2) \tag{3.4}$$

$$\epsilon_i \sim t_\nu(0, \sigma^2) \tag{3.5}$$

where $\boldsymbol{\gamma}_j$ is a vector of varying coefficients corresponding to exchange j, which includes an intercept term; $\boldsymbol{\mu}_\gamma$ is a vector of the mean values of each coefficient in $\boldsymbol{\gamma}_j$ for the overall sample; $\boldsymbol{\Sigma}_\gamma$ 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; ϵ_i is the residual for dyad i; and ν 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 3.1 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 3.2, Equation 3.3, Equation 3.4,

and Equation 3.5. These values are referred to as hyper-parameters 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 z_i, and the coefficient for any particular exchange is thus a deviation from this value. The magnitude of this deviation depends on they hyper parameter as well as the data Σ_γ.

3.4.3 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 hyper-parameters, 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 interference 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 with in 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 coefficients[9]. Another example of this type of weak constraint involves variance hyper-parameters, which must be non-negative. More formally, the various prior distributions for the model parameters described above are as follows:

$$\beta_n \sim N(0, 5^2) \tag{3.6}$$

$$\mu_{\gamma_m} \sim N(0, 5^2) \tag{3.7}$$

where n and m index the exchange-invariant and varying coefficients, respectively. Σ_γ is decomposed into a vector of coefficient standard deviations, $\boldsymbol{\sigma}_\gamma$, and off-diagonal correlations Ω_γ such that:

$$\Sigma_\gamma = D(\boldsymbol{\sigma}_\gamma)\Omega_\gamma D(\boldsymbol{\sigma}_\gamma) \tag{3.8}$$

where $D(\boldsymbol{\sigma}_\gamma)$ is a square diagonal matrix such that $D_{mm} = \boldsymbol{\sigma}_{\gamma m}$ where m indexes the parameters that vary by exchange and all elements are 0. Ω_γ and $\boldsymbol{\sigma}_\gamma$ are a vector and matrix, respectively, defined by the following prior distributions:

$$\sigma_{\gamma_m} \sim \text{half-t}_3(1) \tag{3.9}$$

$$\Omega_\gamma \sim LKJ(1) \tag{3.10}$$

where $LKJ(\zeta)$ 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

[9]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)

parameters are given the following weakly-informative priors:

$$\sigma_\delta \sim \text{half-t}_3(1) \tag{3.11}$$

$$\nu \sim gamma(2, 0.1) \tag{3.12}$$

$$\sigma \sim \text{half-t}_1(\sigma_y) \tag{3.13}$$

where σ_y is the standard deviation of the dependent variable. With the exception of the priors for the coefficients in Equation 3.6and Equation 3.7, all of the above priors are the defaults suggested by the software used to fit the model (Buerkner 2017, see).

3.4.4 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 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 devel-

oped 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).

3.5 Results

This section presents the results of the analysis in four parts. First, it briefly discusses the community detection procedure with a focus on how the inter-slice strength parameters were tuned to best fit the data. Next, it presents the distributions of the estimated parameters from the multilevel model. It then attempts to verify the validity of the model by comparing predicted values to the observed data. Finally, it briefly compares the variation in coefficient estimates with characteristics of the exchanges.

3.5.1 Community Detection Parameters

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 theoretical expectation of the

stability of the underlying communities relative to frequency of noisy variations in the observed ties.

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.

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 as described in Chapter 2. I then compared a range of strengths of connections among the different networks across the four measurement schemes according to their ability to iden-

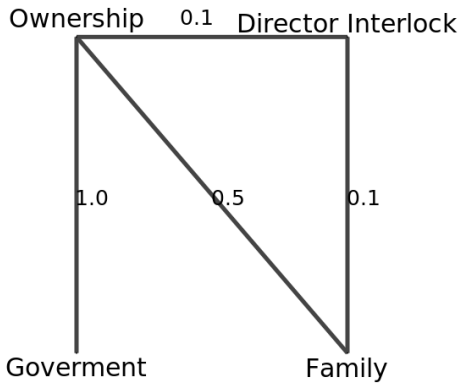

Figure 3.4: The four networks used in the multiplex community detection analysis are shown as nodes in a multi-network structure. The number by each edge is the strength of the connection between each network. Networks with no edge are not connected.

tify 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 3.4. 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.

3.5.2 Multilevel Regression

The model in Equation 3.1 was estimated using four simultaneous HMC chains, each with a total of 2000 sampling draws. The first 1000 draws in each chain

were considered a warm-up phase during which the algorithm searches for a high-probability area of the parameter space. The inferences reported below are thus based on a total of 4000 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, known as $\hat{R}$. An $\hat{R}$ 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 3.3 presents the 0.025 percentile, median, and 0.975 percentile of the estimated distributions of the entire-sample coefficients. This includes both the invariant coefficients for direct ownership and shared industry as well as the other variables whose coefficient vary by exchange. The latter category is accompanied by the median of the corresponding σ_γ parameter, which indicates the standard deviation of the coefficient among the different exchanges. These entire-sample estimates broadly support the hypotheses that director interlocks, shared ownership, family group comembership, and government ownership are each associated with an independent increase in firm synchronicity. With the key exception of family group comebership, all of the corresponding parameters do not include zero in their 95% 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. For example, zero falls at the 0.056 percentile, while at the other end the 0.944 percentile falls at 0.063, meaning that the overall mean of the parameter is just as likely to be above 0.063 as it is to be below zero.

However, these mean estimates only show one part of the picture. As the

Coefficient	2.5%	Median	97.5%	σ_γ
Intercept	0.071	0.150	0.231	0.132
Shared Director	0.002	0.009	0.018	0.009
Direct Ownership	0.056	0.090	0.128	-
Shared Ownership	0.074	0.153	0.241	0.105
Squared Shared Ownership	-0.311	-0.153	-0.020	0.170
Same Family Group	-0.009	0.030	0.074	0.057
Government Ownership	0.020	0.062	0.101	0.040
Squared Gov. Ownership	-0.131	-0.076	-0.024	0.025
Same 2-Digit SIC	0.015	0.018	0.022	-
Same 3-Digit SIC	0.028	0.033	0.038	-

Table 3.3: Median and 95% posterior intervals for entire-sample coefficients.

σ_γ column indicates, there is substantial variation in the parameters across the different exchanges. Figure 3.5 shows the exchange intercepts in the form of the exchange-specific deviation plus the overall mean from Table 3.3 plotted against the exchanges' overall R^2 value for the entire sample period. The strong linear relationship supports the similarity between traditional R^2 measures and dyadic synchronicity. The one exception is Bahrain, likely due to a combination of a small number of firms and low liquidity, and more will be said on this anomaly below. Each firm also has its own intercept parameter that applies to all dyads of which it is a member. Figure 3.6 and Figure 3.7 plot the median value of these intercepts relative to the number of days with non-zero returns and R^2 for each firm, respectively. Exchanges are indicated in both figures by different shapes in order to give a sense of how the relationship varies in different contexts. In Figure 3.6 the number of non-zero returns has a weakly positive relationship with the firm-intercept, but this relationship is stronger for some exchanges such, such as Egypt, represented by the triangles visible on the lower half of the figure. 3.7 again displays a strong relationship between R^2 measures and model intercepts. Taken together, these three figures provide evidence that

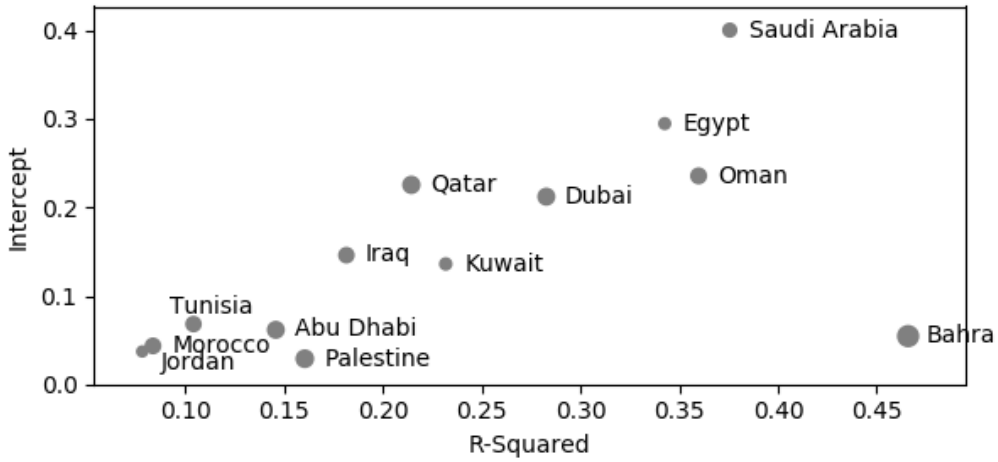

Figure 3.5: The total exchange-specific intercepts are plotted relative to their R^2 values for regressing daily non-zero returns on an unweighted daily index of the exchange's average return over the 2010-2016 period. Points are sized according to the relative uncertainty of the estimates.

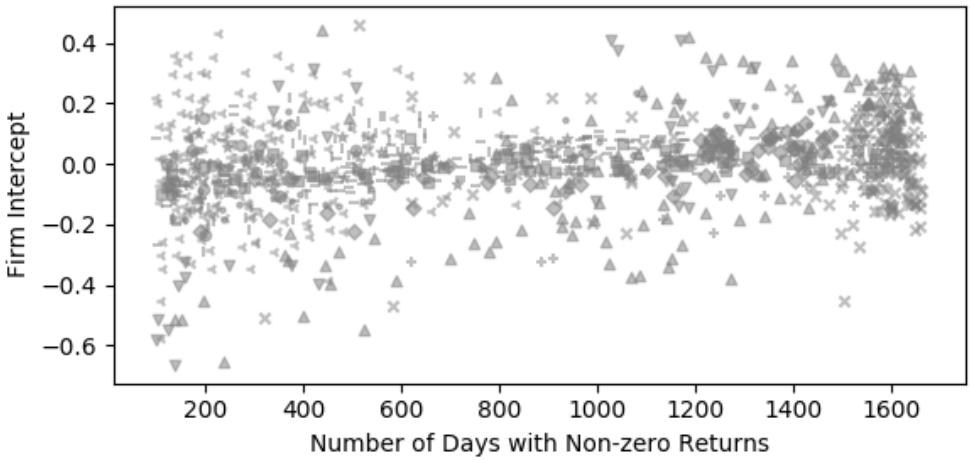

Figure 3.6: Median raw firm-specific intercepts are plotted relative to their number of non-zero daily returns. The exchange of a firm is indicated by the shape of each point in order to provide an impression of the overall variation in the relationship.

the model captures meaningful differences in the baseline synchronicity at both the firm- and exchange-level.

Table 3.3 provides initial evidence in support of the hypotheses predicting positive associations between firm relationships and price synchronicity, but it remains to assess how the estimated parameters vary by exchange. Figure 3.8

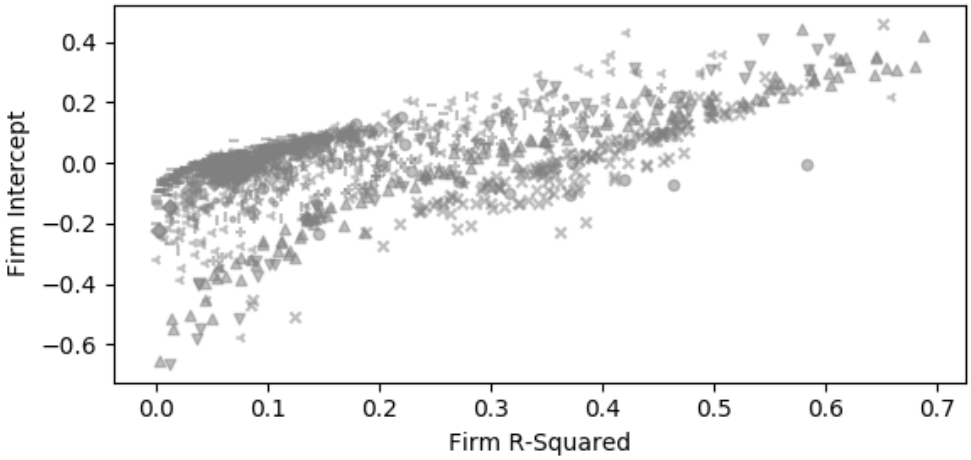

Figure 3.7: Median raw firm-specific intercepts are plotted relative to their R^2 values for regressing their daily non-zero returns on an unweighted daily index of the exchange's average return over the 2010-2016 period. The exchange of a firm is indicated by the shape of each point in order to provide an impression of the overall variation in the relationship.

displays the magnitude and uncertainty of the director interlock and shared ownership coefficients for each exchange. Focusing first on shared ownership, the exchanges can be divided into three rough categories: Dubai, Kuwait, Jordan, Saudi Arabia, Qatar, Iraq, and Egypt have estimated distributions that are far from zero; Palestine, Abu Dhabi, Bahrain, and to a lesser extent Oman have large median estimates coupled with 95% posterior intervals that marginally include zero; and finally Tunisia and Morocco have median estimates closer to zero and a greater share of negative values in their posterior intervals. 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 referring back to Table 3.1 shows that the first group includes the largest exchanges in the region as well as Qatar, Dubai, and Iraq, while the second group of exchanges have the fewest listed firms.

The relationship between synchronicity and director-interlocks appears less robust compared with shared ownership. Only Kuwait, Egypt, and Morocco

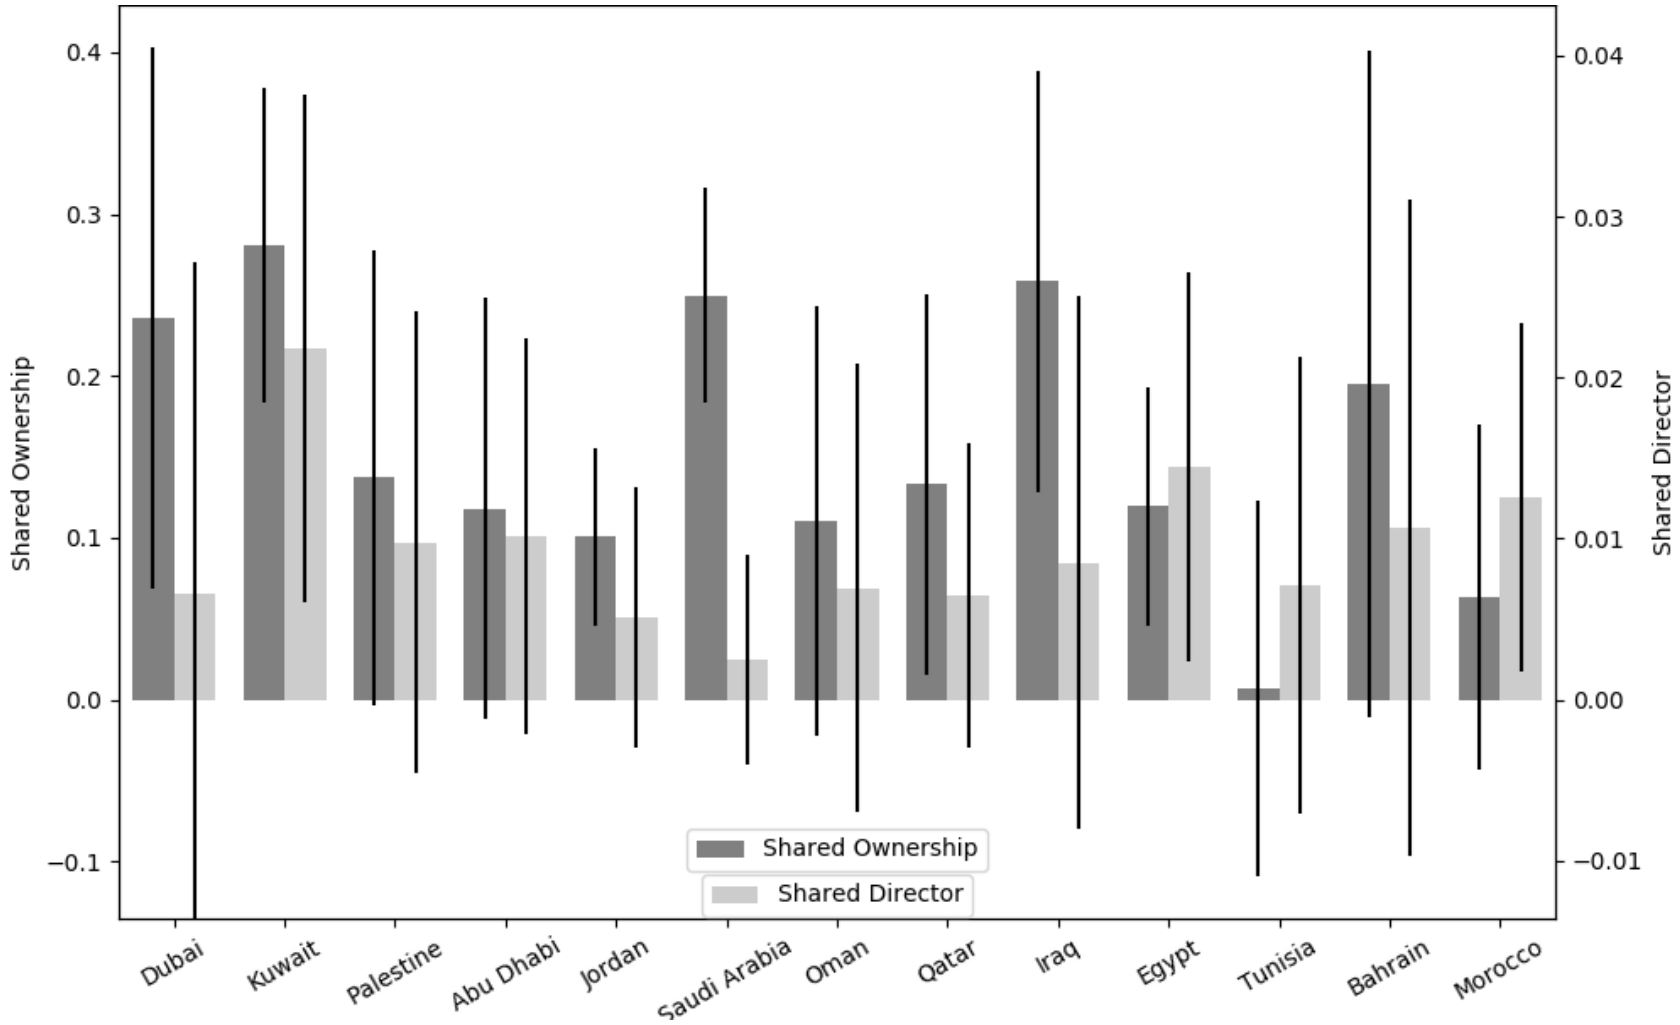

Figure 3.8: Bars indicate the median coefficient estimate for each exchange, and the lines show the extent of the 95% posterior intervals. The scales of the shared ownership and shared director coefficients are shown on the left and right y-axes, respectively.

have an estimated distributions with a 97.5% probability of being greater than zero. The other exchanges have both smaller median values and a large proportion of negative values in their 95% posterior intervals. When taken as a whole, all but two of the exchanges have a median value that falls between 0.005 and 0.015, which helps to explain why the distribution of the mean parameter for the entire sample presented in Table 3.3 nevertheless has a 95% posterior interval between 0.002 and 0.018.

Interpreting the substantive meaning of the coefficients for these predictors is simple in the case of director interlocks because as a binary variable it represents the mean difference between interlock and non-interlock dyads conditional on the other predictors in the model. Shared ownership, however, is not only continuous but also modeled as having a quadratic relationship with price synchronicity. Figure 3.9 displays the median and 95% posterior interval for the overall effect of increasing shared ownership for each exchange. Each figure also provides a histogram indicating the relative frequency of different shared ownership values to give a sense of the amount of data relevant to each part of the curve. As the cases of Dubai and Saudi Arabia show, there is little support for the decreasing areas of the curves, and the negative coefficients for the squared terms instead reflect that the marginal effect of shared ownership levels off at high values. Overall, a shared ownership value of 0.25 is associated with a median increase in correlation between 0.015 and 0.050 for all countries except Morocco and Tunisia.

Director-interlock and shared ownership ties are a basic part of the complex web of relationships that link firms to one another. Family business groups and government ownership are examples of the variety of other connections that

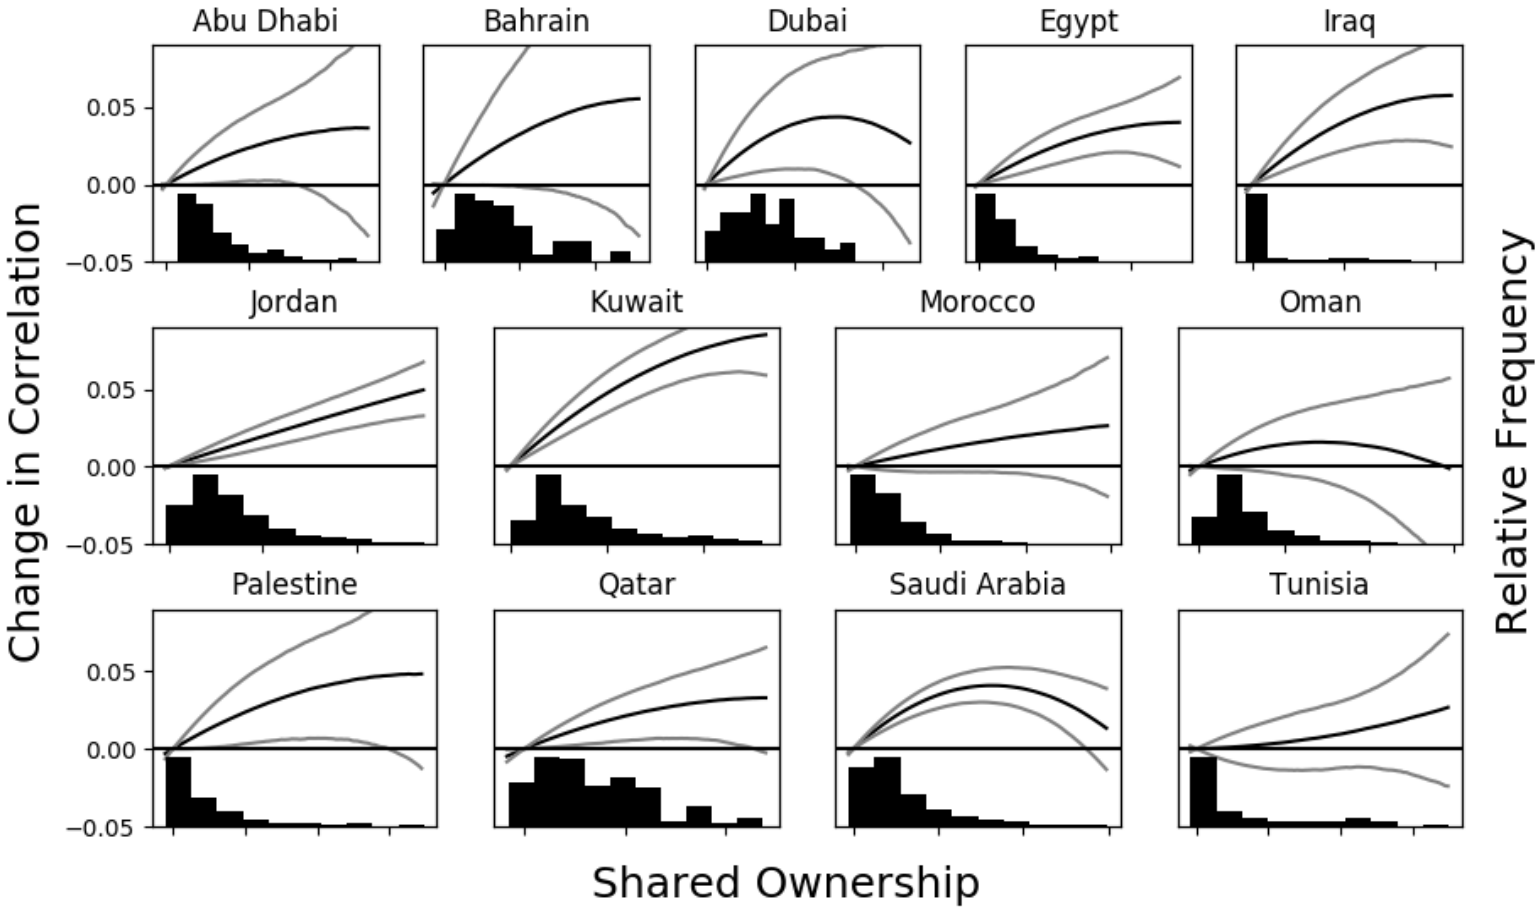

Figure 3.9: The dark curve toward the top of each plot shows the median increase in price correlation associated with a given level of shared ownership. The two lighter lines give the bounds of the 95% posterior interval. Below these is a histogram of the relative frequency of observed shared ownership values. In order to facilitate visibility, they are not shown on the same scale and exclude observations with no shared ownership.

might be more salient to the control and performance of large corporations. Figure 3.10 shows the coefficients for these two variables in each exchange. Starting on the left of the figure, Dubai, Kuwait, Palestine, Abu Dhabi, Jordan, and Oman all show strong evidence that family business group comembership, as measured by the presence of family members within the same network community, is an important predictor of stock price similarity. Indeed, compared with the results for director-interlocks in Figure 3.8 or the marginal effect of shared ownership in Figure 3.9, these coefficients are large, indicating that markets perceive a much closer relationship between these types of firms than between other pairs of firms that might have similar ownership or director-interlock connections without the underlying family business group. The other six exchanges, however, do not show a meaningful relationship between family business groups and price comovement. These six exchanges include only 79 of the 347 total family business group dyads, and of these 79, 47 are in Qatar. A smaller number of relevant observations should increase the uncertainty of the estimates, but does not account for the fact that many of the median values are near or substantially below zero.

At least three factors might explain these null results. 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, which has more family dyads that the rest of these six combined, appears to be uniquely unsuitable for the method of looking for shared surnames among individuals associated with firms in the same commu-

nity, 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 3.1. A second explanation is that family business groups might have disproportionately poor data, and hence the estimated coefficient is based on a non-random subset of ties. 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 enough data to be included in the model. In terms of family group dyads, however, this yields only 7 out of a total of 37 possible observations[10] because often only one firm in a group has enough data, which means that no pairwise observations are possible within that cluster. Furthermore, many of these dyads are from clusters that have weaker evidence of being associated with a true family business group, as measured by total family participation or mean family ownership. 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 any other relational measure, Tunisia appears to fit this category.

The association between government ownership and price comovement also varies between the different exchanges. Kuwait, Jordan, Saudi Arabia, Qatar, Egypt, Bahrain, and Morocco all have estimated coefficients with a 95% posterior interval that is greater than zero. Among the other exchanges, the greater uncertainty for Dubai and Abu Dhabi might reflect the need to disaggregate the governments of the seven emirates that make up the UAE. The lack of any evidence for the influence of government ownership in Palestine can perhaps be attributed to its unique status under Israeli occupation. The role of government

[10]The number of possible dyads is higher here than in Table 3.2 because many firms had no Datastream data at all and hence were not included in the table.

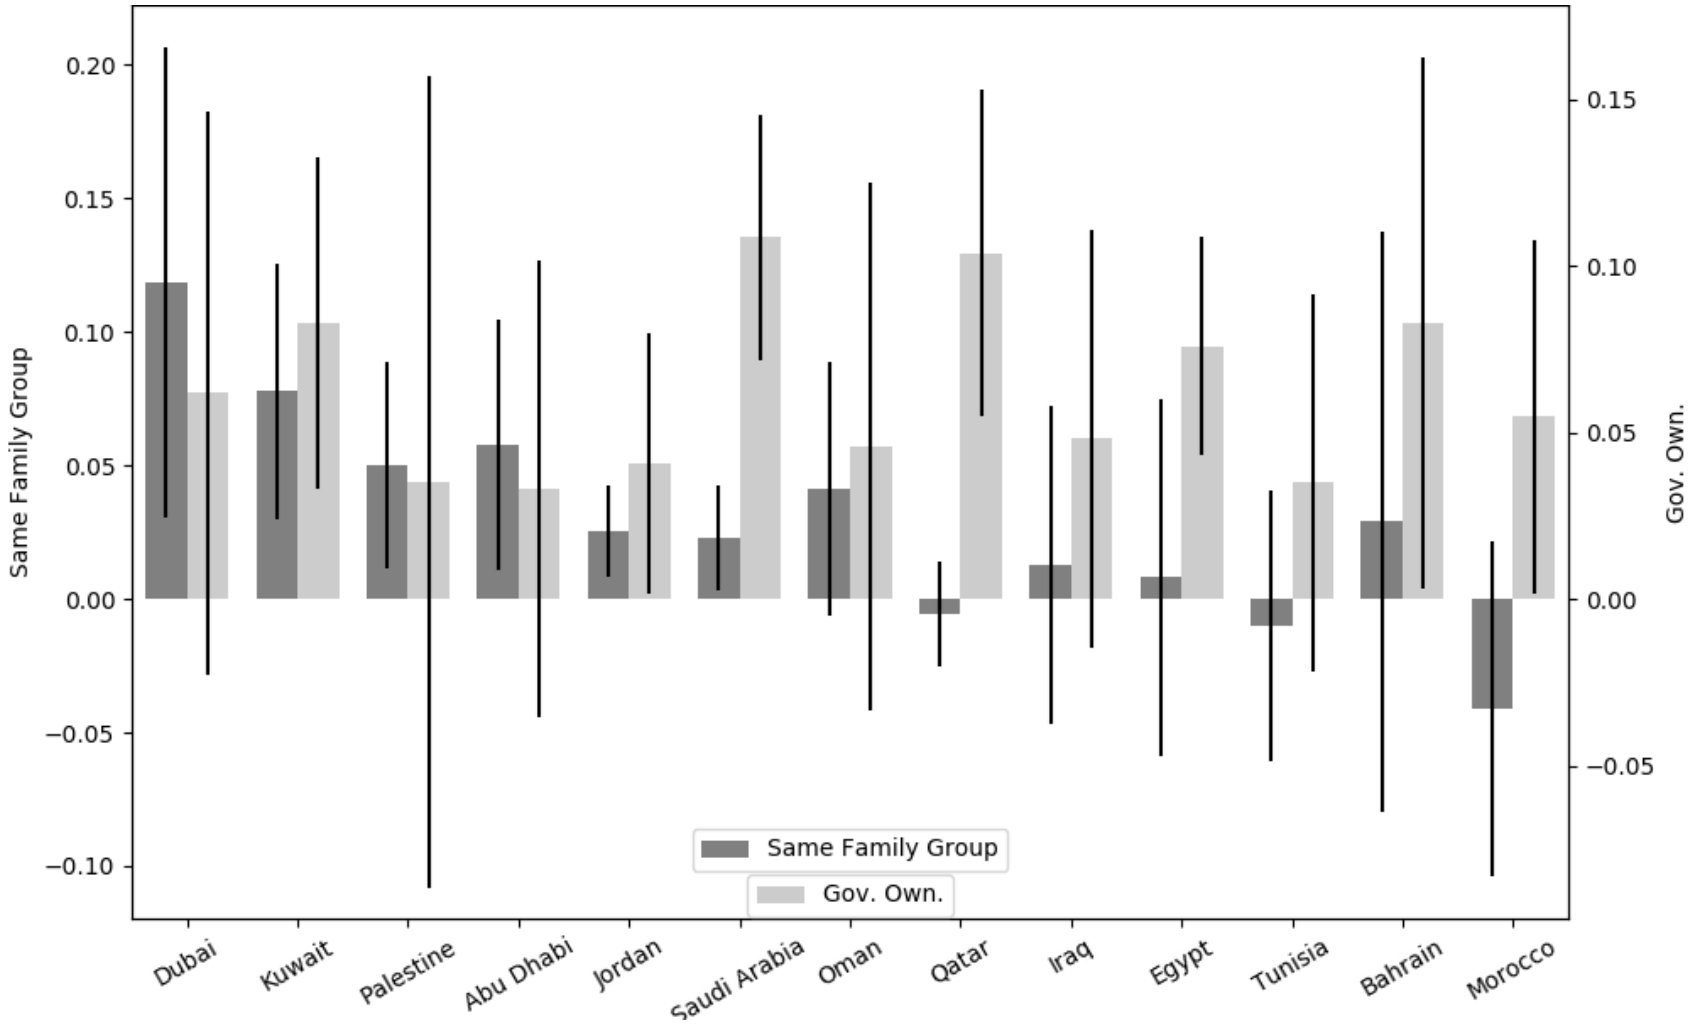

Figure 3.10: Bars indicate the median coefficient estimate for each exchange, and the lines show the extent of the 95% posterior intervals. The scales of the same family group and government ownership coefficients are shown on the left and right y-axes, respectively.

ownership in Tunisia is again consistent with the low information content of its price variations. Iraq and Oman are more ambiguous, and given the relatively high uncertainty for the other coefficients in the model, the results might reflect either the uneven information content of price changes or a true ambivalence on the part of market participants about the impact of state ownership on firm governance and performance.

Figure 3.11 shows the marginal effect of increasing government ownership according to the model. The layout is similar to that of Figure 3.9. Including the coefficients of the squared terms for each exchange shows the further nuance that among the exchanges with 95% posterior intervals that do not include zero, Saudi Arabia and Qatar have the highest marginal effect with low uncertainty even for high government ownership; while the posterior intervals for Egypt, Kuwait, Bahrain, Jordan, and Morocco dip below zero for high values. The marginal effect of government ownership is smaller than shared ownership because it reflects the additional synchronicity between firms owned by government organizations relative to those with any shared owner.

3.5.3 Predictive Checks of Model Fit

The model specification and estimation presented above are two essential features of Bayesian analysis but are incomplete without an evaluation of the fit of the estimated model (Gelman et al. 2014, p. 139). This evaluation is typically performed through posterior predictive checks that compare the observed values of the dependent variable with simulated data generated by draws from the fitted distributions of the model parameters (Lynch and Western 2004; Gel-

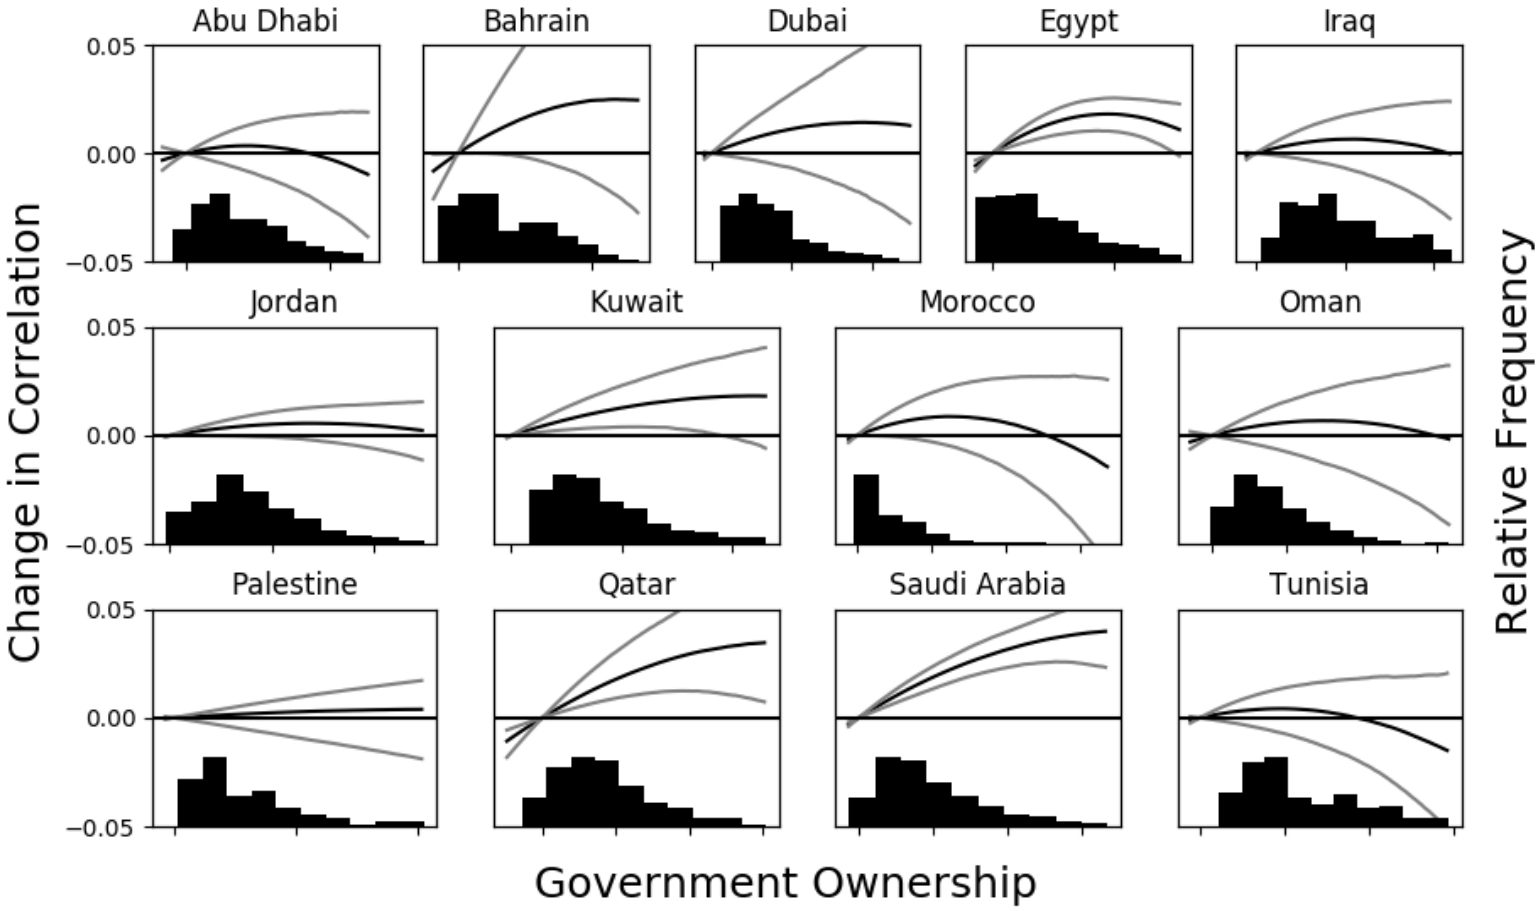

Figure 3.11: The dark curve toward the top of each plot shows the median increase in price correlation associated with a given level of shared ownership. The two lighter lines give the bounds of the 95% posterior interval. Below these is a histogram of the relative frequency of observed shared ownership values. In order to facilitate visibility, they are not shown on the same scale and exclude observations with no shared ownership.

man et al. 1996). This comparison can be made by specifying a test-statistic that describes a key feature of the data that the model is designed to capture. A distribution of this statistic under the model is then calculated using the simulated outcomes.

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 3.1 was designed to test this claim by searching for heightened synchronicity between pairs of firms in the same inferred business group, relative to other firms with observed connections but no inferred group comembership. 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 0.05 government ownership and no family group comembership. Figure 3.12 shows the distribution of this statistic for 500 replications generated by random draws from the posterior distribution given by the model relative to the observed data. The statistic from the observed data is greater than 79.8% of the simulated values derived from the model, and graphically falls within the main body of the distribution, indicating a reasonable fit.

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 3.13 shows the observed versus simulated values for this statistic. The observed statistic is

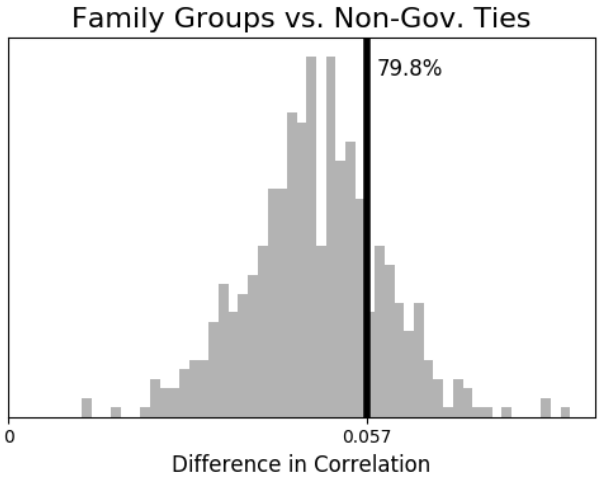

Figure 3.12: The black line shows the observed difference in the mean synchronicity of family group dyads and dyads with other types of observed, non-government connections relative to the distribution of this statistic, derived from 500 sets of simulated values. The percent at the top of the figure is the percent of simulated statistics less than or equal to the observed value.

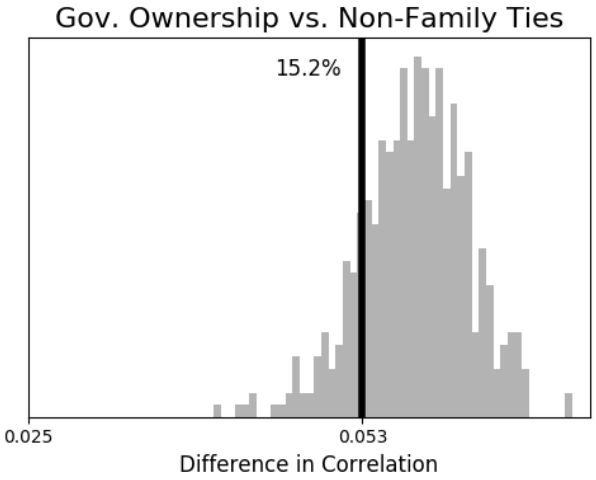

Figure 3.13: The black line shows the observed difference in the mean synchronicity of government-owned dyads and dyads with other types of observed, non-family connections relative to the distribution of this statistic, derived from 500 sets of simulated values. The percent at the top of the figure is the percent of simulated statistics less than or equal to the observed value.

greater than 15.2% of the simulated values and again falls within the main body of the distributions, as opposed to its tail.

Statistical models are necessarily oversimplifications of actual relationships, but these posterior predictive checks provide confidence that the model in Equa-

tion 3.1 is capable of reproducing relevant aspects of the observed data. These assessments are 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 expected values can also be used to check for dependence among the residuals by estimating the correlation between the errors of different observations over the replicated outcomes. Figure 3.14 evaluates the assumption that the dyad-specific residuals are independent conditional on firm random intercepts by plotting the mean correlation among between all dyad residuals involving each firm in the sample. Each point represents a single firm, and the mean residual correlation is plotted on the y-axis relative to the number of non-zero returns for that firm on the x-axis. A large number of points not clustered at zero would indicate that the observations are not conditionally independent. The exchange of each firm is indicated by its shape. Overall, 93.7% of firms have a mean residual correlation of less than 0.01, and 98.4% have a value of less than 0.02. The 18 firms with a mean of over 0.02 are visible in the upper left of Figure 3.14, and of these 2 are listed on the Dubai Financial Market, 3 are from the Muscat Securities Market in Oman, and the remaining 13 belong to the Bahrain Bourse. Moreover, the Bahraini firms are clearly visible due to their aberrantly high residual correlations and low number of non-zero returns. This echoes the pattern in Figure 3.5 where Bahrain was the only exchange whose intercept deviated from what might be expected from its R^2 value. Given these inconsistencies, and that there are only 20 Bahraini firms in the data with over 100 non-zero returns, I plan to drop Bahrain in a future version of this paper.

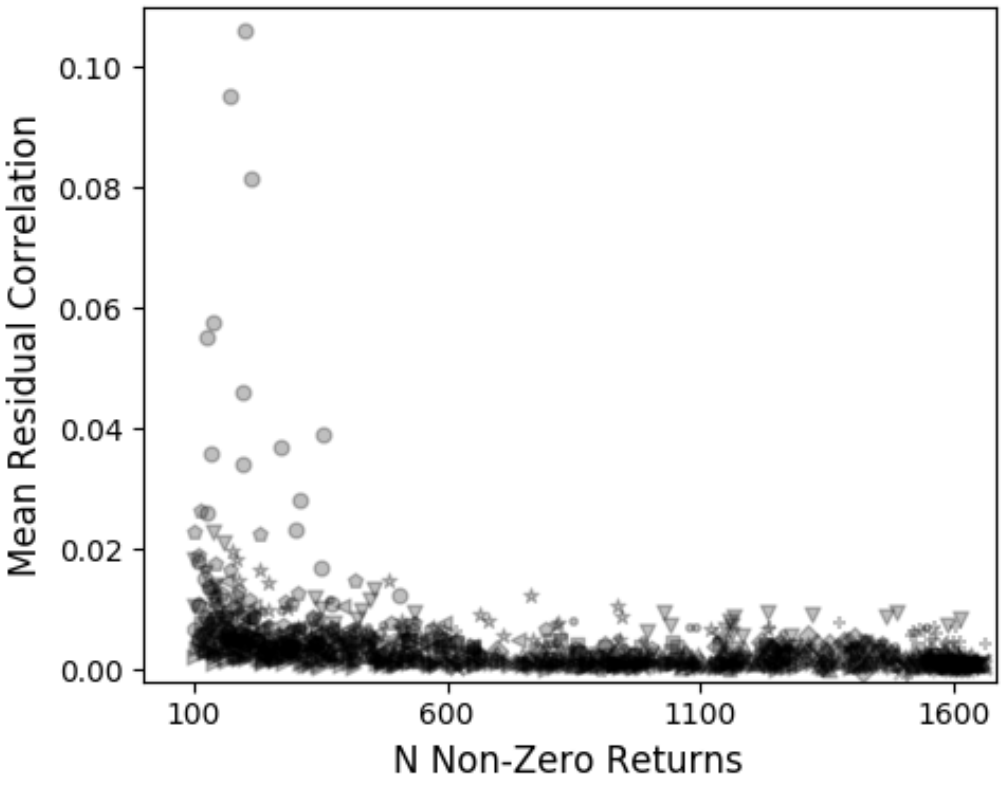

Figure 3.14: The mean residual correlation of the dyads involving each firm are plotted relative to that firm's number of non-zero returns. Exchanges are indicated by the shape of each point. Note the number of Bahraini firms represented by the circles in the top-left of the figure.

3.5.4 Coefficients and Exchange Level Characteristics

The model in this paper has focused on the intermediate pooling of coefficient estimates allowed by multilevel models, but another important strength of the method is that multilevel modeling allows for the simultaneous inclusion of predictors at different observation levels. Higher-level variables can be used not only to predict the outcome directly but also to test hypotheses about how the qualities of higher-level units influence the coefficients of lower-level predictors. Including group-level predictors can improve the estimation of a model, since the additional information might reduce the variance of the coefficients that are

allowed to vary between different groups. Despite these potential benefits, I chose to limit the model to dyad-level predictors for three reasons. First, there is little prior work to inform how exchange-level attributes might influence the relationship between inter-firm ties and price synchronicity. Furthermore, the small number of exchanges compounds this challenge by likely yielding uncertain estimates that might not be generalizable to exchanges in other regions. The number of firms, on the other hand, is much higher, but they are less relevant for the purpose of the model as long as their random intercepts allow for the conditional independence of the dyad observations.

This section presents a exploratory analysis by investigating the correlations between five exchange- and country-level attributes, and three varying coefficients from the model. While stopping short of including them in our full model, this approach allows us to at least evaluate the strength of possible bivariate relationship. The exchange-level variables are: market liquidity, measured as the average of the value of monthly trades on an exchange divided by its capitalization; an R^2 synchronicity measure as defined above; a governance index composed of four World Bank World Governance Indicators[11]; an index of three shareholder control measures from the World Bank Ease of Doing Business database[12]; and an index of four minority investor protection measures from the same source[13]. The correlations of these variables with the coefficients for family business group comembership, shared ownership, and government ownership were calculated for each of the 4000 samples from the posterior dis-

[11]Government effectiveness, regulatory quality, rule of law, and control of corruption. I excluded Voice and Accountability because it has relatively little variation in the region and stability because it is weakly correlated with the other four.

[12]Extent of shareholder governance, extent of ownership and control, and extent of corporate transparency

[13]Extent of conflict of interest regulation, strength of minority investor protection, extent of disclosure, and extent of director liability.

	Family Business Group			Shared Ownership			Government Onwership		
	2.5%	Median	97.5%	2.5%	Median	97.5%	2.5%	Median	97.5%
Market Liquidity	-0.101	0.212	0.427	0.059	0.384	0.660	-0.021	0.411	0.775
R^2	-0.174	0.227	0.566	-0.001	0.399	0.698	-0.037	0.462	0.759
Governance	-0.065	0.352	0.613	-0.442	-0.055	0.302	-0.340	0.099	0.492
Shareholder Control	-0.300	-0.063	0.230	-0.323	-0.023	0.279	-0.360	0.090	0.480
Investor Protections	0.273	0.592	0.761	-0.086	0.306	0.636	-0.359	0.069	0.515

Table 3.4: The distribution of each correlation measure is calculated based on the full distribution of the model coefficients.

tribution, in order to represent the uncertainty of the estimated correlation. The 95% posterior intervals and median of each correlation value are presented in Table 3.4.

As expected from the small sample size, the estimates are quite noisy, but two results are notable. First, minority investor protections appear to be strongly associated with the value relevance of family business groups. Governance quality shows a weaker relationship. Second, market liquidity and overall price synchronicity are moderately linked with the relevance of shared ownership and government ownership. These correlations suggest potentially fruitful areas for future research on the political economy of financial markets in the Middle East and North Africa. For example, do value-relevant family groups advocate for investor protections, or is it perhaps the case that their synchronicity is only observable because these regulations strengthen the price mechanism by encouraging outside investment in closely held firms? On the other hand, what are the links between liquidity, overall-synchronicity, and ownership? It seems plausible that exchange-level synchronicity would be buttressed by widespread co-ownership, and that higher trade volume would improve the value-relevance of interconnections relative to random noise, but what exactly is the role of government ownership in this process? In any case, the tenuousness of these results must emphasized. Not only are they based on 13 data points,

but at least two of these, Tunisia and Bahrain, appear to have poor enough data to warrant skepticism that they provide useful information.

3.6 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 three contributions to the literatures on inter-firm 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, shared ownership, and director-interlocks in 13 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 13 exchanges. These results help make the case for using this type of unsupervised learning technique to study the role of business groups in the Middle East and North Africa, as well as other 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. Finally, it presents initial findings on the links between the value relevance of family business groups and minority-shareholder protections, as well as among market liquidity and synchronicity and shared ownership, both

governmental and otherwise.

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 non-zero 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. Still, the cutoff described above is probably cruder than necessary, and I plan to replace it with an appropriate weighting scheme in a future revision. As for the model, its complexity is limited by the practicalities of computation, and the model could be improved by allowing for heterogeneity in direct ownership and industry effects, for example by disaggregating the latter into categories like finance or resource extraction.

CHAPTER 4

GUARANTEED LOANS, INTER-ORGANIZATIONAL RELATIONSHIPS, AND BANK DECISION-MAKING IN CHINA (WITH LISHA LIU)

4.1 Introduction

Business groups have incentives and opportunities to use inter-firm transactions in order to divert the resources from publicly-held firms toward their closely held affiliates. On the other hand, such transfers might also serve an important economic purpose beyond the enrichment of the group's ultimate controlling owners. The practice of firms guaranteeing the loans of other companies is an important example of this type of transfer whereby the guarantor firm agrees to assume some form of liability for the repayment of the loan in the event that the recipient of the guaranteed loan defaults. These transactions thus involve three distinct parties: the loan recipient, the guarantor who makes the guarantee, and the bank providing the loan.

From the perspective of the lender, guaranteed loans can be problematic. Although they can improve the distribution of capital by allowing smaller firms to borrow at lower rates, the guarantor is often related to the recipient and hence vulnerable to the same types of factors that might cause the latter's default. Furthermore, if the guarantor is publicly-traded, then the transaction will push the risk of default onto its stockholders, resulting in an observable decrease in firm value. This paper investigates the evolving role of guaranteed loans in the Chinese economy from the perspective of the lenders' varying sensitivity along three dimensions: corporate governance, relationships between the guarantor and bank, and the political capital of the guarantor.

Using Bayesian multilevel logistic regression models, we find that better governed firms are more likely to complete guarantees, that the influence of political ties is contingent on the ownership of the firm, and that banks prefer to make loans to well governed firms with which they have a history of interaction. These results have important implications for research on the role of political ties in emerging markets, financial embeddedness, and the interaction between business groups and their environment.

4.2 Literature Review and Hypotheses

Our analysis of loan guarantees in China draws heavily from the literature on related party transactions (RPTs) in China. As the name implies, these transactions typically occur between firms that belong to the same business group or are bound by other ties. However, loan guarantees can be distinguished from other RPTs that involve direct transfers of capital, goods, and services from one party to another in that they involve a third party, the bank making the loan. Hence, evaluating a lender's willingness to provide a guaranteed loan provides an opportunity to characterize how these transactions are perceived by an external actor.

We model a bank's decision-making strategy along three dimensions: the corporate governance structure of the guarantor firm, the degree of financial embeddedness between the bank and guarantor, and the political connections of the guarantor. We also take advantage of China's ongoing development into an advanced market economy to assess how temporal and spatial variations in marketization might influence the importance of these different factors. This

section describes the motivation for our approach by first describing key elements of the empirical context before summarizing prior research involving each of the three classes of predictors. It concludes by presenting a set of hypotheses.

4.2.1 Guaranteed Loans in China

Although not all loan guarantees are RPTs, related-party guarantees comprise over 90% of our observations. A loan guarantee to a related party refers to a guarantee issued by one entity that it will ensure repayment of a loan made to a related entity by a third party, usually a bank. In our data, the entity issuing the guarantee is a listed firm and the loan recipient is typically privately held. According to the Chinese Generally Accepted Accounting Principles (GAAP), firms are defined as related-parties if one party controls, jointly-controls, or exerts strong influence over another party or if two firms are under the control or strong influence of a common third entity (Huang 2016).

A guarantee transaction involving a public firm typically proceeds as follows. First, the guarantee firm, which is the primary debtor firm that plans to apply for bank loans, and the guarantor firm sign an intended guarantee contract that will be submitted with the bank loan application of the primary debtor firm. After receiving the loan application, banks further investigate the borrower's creditworthiness as well as the guarantor's information and decide whether to approve the loan. In the event of a default, banks will turn to the guarantor firm for the repayment according to the terms of the contract.

The borrowing activities of public firms and their affiliates are also influ-

enced by the unique history of business groups in China. Chinese groups have emerged not only as a response to market failures but also as a direct result of government policy (Keister 2001; Lee and Jin 2009). As a result, many business groups in post-reform China remain state-owned and are directed by professional managers, which contrasts with the well documented importance of family owned groups in other parts of East Asia and the world (Luo and Chung 2005). A further idiosyncrasy is that China tightly regulates a group's access to equity markets, and in order to meet the listing requirements of the China Securities Regulatory Commission (CSRC), a business group typically selects one of its strongest firms and spins off low-performing assets to make it more attractive to investors (Fisman and Wang 2010). This process is illustrative of the governance issues faced by Chinese firms since the resulting improvement in performance does not diminish firms' incentives to support their less-profitable affiliates rather than distribute surplus revenue to shareholders.

Over the past two decades, the Chinese government has implemented various measures in order to mitigate these types of problems. For example, in June 2000 the Chinese securities market regulator introduced a regulation prohibiting the issuance of any new debt guarantees to shareholders of listed firms or to subsidiaries of these shareholders (Berkman et al. 2005). In 2004 the State Council issued a directive specifically to address the expropriation of public firm assets through RPTs stating, "We must prevent controlling shareholders from embezzling listed company assets, and punish those who did." (Jiang et al. 2010). Finally, in 2005, China Securities Regulatory Commission introduced a regulation mandating the disclosure of loan guarantees provided by listed firms if a board or shareholders' meeting has approved the intention to offer the guarantee. This was soon supplemented by a 2006 revision of Arti-

cle 36 of the Generally Accepted Accounting Principles (GAAP) to improve the disclosure of RPTs more generally.

The Chinese Banking sector has also evolved significantly as China has become more market-oriented. Historically, the Peoples Bank of China (PBC) was the only bank in China, acting as a hybrid between a central bank and a commercial bank (Lin and Zhang 2009). As part of market economy reforms from 1979 to 1984, the big-four state-owned banks—the Bank of China, the China Construction Bank, the Agricultural Bank of China, and the Industrial and Commercial Bank of China—were established. Further reforms introduced a series of joint-stock or joint-equity banks (Liang et al. 2013). In 1994, the promulgation of bank laws (i.e. the Central Bank Law and the Commercial Bank Law) and the establishment of policy banks brought additional changes. Another exciting development has been the emergence of regional commercial banks in both rural and urban areas (Zhang et al. 2016). Nevertheless, the nature of state ownership and its associated soft budget constraints continues to contribute to a large proportion of non-performing loans in the Chinese banking sector (Allen et al. 2017).

4.2.2 Corporate Governance

This ongoing process of reform paralleled the development of a large literature on the relationship between corporate governance and RPTs. Although the traditional concern of the corporate governance literature has been on the principal-agent conflict between owners and managers, more recent work has emphasized that many nascent financial markets across the world struggle im-

plement any meaningful public control over openly traded firms (Porta et al. 2002; Morck et al. 2005; Aguilera and Jackson 2010; Davis 2012). This is because after the public offering of shares the original owners of the firms often retain an overwhelming large stake through mechanisms like ownership pyramids or cross-holdings. This has two important implications. The controlling owner is not obligated to seek the consider the opinions of other shareholders, and second, they are less concerned about market reactions to firm behavior since only a small fraction of the total equity is traded. This results in a reformulation of the classic principal-agent dilemma as a conflict between two principals where one has total control over the firm but both are entitled to a fair share of the profits.

The most widely studied manifestation of this conflict of interest is "tunneling", or the transfer of resources from a publicly held firm to a different company more closely held by the listed firm's controller (Bertrand et al. 2002). Indeed, much of the literature on RPTs in China is devoted to improving our understanding of this phenomenon (e.g. Jiang et al. 2010; Liu and Tian 2012). However, the empirical identification of such expropriation can be problematic. For example, redistribution among related parties has been shown to harm high-performing firms in the short term, but there is also strong evidence that group firms can reap important benefits from such mutual aid in times of crisis (Lincoln et al. 1996; Almeida et al. 2015). Thus the impact of RPTs on economic outcomes depends on the intent behind the transaction, which is largely unobserved. Work on Chinese RPTs has approached this issue from a variety of directions, for example by categorizing RPTs as efficient or harmful based on market reactions, firm governance, or the current performance of the firm and its affiliates (Fan et al. 2016; Berkman et al. 2010; Wong et al. 2015). By inves-

tigating banks' willingness to provide guaranteed loans, our study avoids this difficulty and instead uses the outcome of a transaction to study how different variables predict banks' decisions. For example, Haveman et al. (2017) find that political ties reduce firms' ratio of intra-group loans guaranteed or provided to those received. However, this could be the result of politically connected firms applying for fewer loans, being rejected for more, or a combination of both, and our study provides a direct test of this second mechanism for political ties as well as corporate governance and financial embeddedness variables

The literature on the relationship between corporate governance and access to capital in China has found a complex intersection of three dimensions: firm attributes, the institutional environment, and political connections. Focusing here on the first two, Firth et al. (2009) find that better governed private firms have more access to bank credit. This reflects the endogenous relationship between credit and governance since not only are better governed firms less risky, but firms that rely on debt rather than accumulated profits are also constrained by the obligation of repayment to pursue higher value projects. However, there is evidence that this latter mechanism does not apply to government-owned firms (Tian 2005). In the case of public firms, the discrepancy between ownership and control rights of the ultimate controller is a direct measure of their incentive to expropriate, and thus has important implications for their relationship with creditors and shareholders (Zhang et al. 2014). A further variable unique to China is the percentage of tradable shares, and Zhu and Zhu (2012) find that the value penalty suffered by public firms engaging in RPTs decreased after a series of government reforms designed to curtail firms' proportion of non-tradable equity and offer the explanation that the improved governance caused firms to engage in fewer tunneling transactions. However, market-based

governance mechanisms are not likely to be effective without a set of matching institutions. For example, Chen et al. (2015) present evidence that the relationship between governance mechanisms and guarantee issuance in family-controlled firms is moderated by the marketization of firms' provinces.

4.2.3 Political Connections

Ties between businesses and the state are of central importance in China, and this relationship is typically operationalized along two dimensions, direct ownership by state entities and formal or informal relationships between business leaders and government officials. From a bank's perspective, a firm's political status might influence their decision making through at least two mechanisms. First, political ties confer substantive advantages that might improve a firm's performance and make it less risky as a guarantor. For example, Li et al. (2008) survey private firms about their party membership and find that affiliated firms have higher performance, more faith in the legal system, and better access to credit, and Haveman et al. (2017) report a similar finding whereby political ties appear to benefit smaller publicly traded firms in competitive sectors.

This latter result echoes other studies that have noted the contingency of the benefits of political ties. While private firms might benefit from the access that such connections provide, they appear to enhance the ability of firms with substantial state ownership to avoid market discipline and engage in more extractive RPTs (Wang 2015; Berkman et al. 2009; Wang et al. 2008; Tian 2005). Furthermore, not all political connections are the same, and numerous studies have found heterogeneous results for ownership by or relationships with local

versus central government entities (Cheung et al. 2009; He et al. 2013; Liu et al. 2016). Indeed, direct government involvement in firm decision making can directly harm performance (Nee et al. 2007). Finally, the implications of a firm's relationship with the state also depends on their broader environment. It is a standard practice for studies of political relationship in China to evaluate their contingency on the marketization of a firm's province, and most studies find that the ongoing shift to market-based institutions appears to reduce the benefits of government ownership (Carney et al. 2009; Calomiris et al. 2010; Li et al. 2008). Indeed, Nee and Opper (2010) provide evidence that political relationships are only valuable to firms operating in industries with intense government regulation, indicating that marketization has proceeded to the point that such ties are no more beneficial than similar relationships in advanced capitalist economies.

4.2.4 Financial Embeddedness

Compared with corporate governance and state-business relationships, the embeddedness of financial transactions within ongoing inter-firm relationships has received less attention. This embeddedness is most commonly defined in terms of the tendency for economic relationships to occur within a wider set of multiplex ties but can also refer to the ability of repeated transactions to develop into a deeper connection (Granovetter 1985; Uzzi 1999; Sleptsov et al. 2013). Previous research on financial embeddedness has demonstrated the capacity of embedded ties to smooth firms' access to finance and subsequent performance. More specifically, entrepreneurs with embedded bank relationships can overcome information asymmetry problems and establish noncontractual gov-

ernance arrangements of trust and reciprocity that reduce banks' uncertainty and decrease their own cost of capital (Uzzi 1999; Uzzi and Gillespie 2002; Uzzi and Lancaster 2003).

Although the focus of this work has primarily been on the relational dimension, the characteristics of an organization and its institutional environment are also important to consider. Previous research on Chinas transitional economy shows the interplay among economic rationality, social networks and institutional links (Zhou et al. 2003). Within the context of the hybrid financial sector in China, banks are engaging with multiple and often competing realities points to bring about change and maintain stability in existing structures. For instance, political and commercial logics co-exist for Chinese banks and banks of different types value political capital and social capital differently. Carruthers and Kim (2011) studied the U.S. financial system and found the importance of politics for many financial market developments and the continuing significance of social factors within finance after the 2008-2010 financial crisis. To obtain comprehensive insights into the role of social embeddedness on firms' access to financial capital, researchers need to take into considerations the alternative channels through which firms can gain favors from banks.

Uzzi (1999) argues that it is important to study the development, use and operation of ties within a wider context. Mizruchi and Stearns (2006) examines the extent to which the effects of interfirm networks on the behavior of firms are historically contingent and shows that in the US there is a decline of the influence of interlocking directorate network on firms use of debt. Elfenbein and Zenger (2014) present an empirical analysis that suggests that relational capital plays a more important role in environments with more uncertainties. In a similar vein,

Rangan (2000) argues that social networks systematically influence efficiency when actors need to but cannot cost-effectively ascertain the identity and reliability of potential exchange partners. The varying institutional environments in China provide an opportunity for us to capture this dynamism.

4.2.5 Hypotheses

Loan guarantees, especially among related parties, have a checkered history in China. Despite their economic rationale of encouraging banks to lend to smaller, privately held firms, they are generally considered to reflect poor corporate governance and to provide an opportunity for political and economic elites to channel resources away from minority investors. However, following decades of marketization and more recent reforms aimed at curbing the abuse of guarantees and other RPTs, the availability of transaction-level information on guarantee outcomes provides an opportunity to assess how firm attributes, political connections, and embedded relationships impact the likelihood of the bank approving the guarantee.

The three dimensions of a bank's decision-making strategy—the corporate governance structure of the guarantor firm, the degree of financial embeddedness between the bank and guarantor, and the political connections of the guarantor—are important lines along which the bank assess the risk of a loan guarantee. We expect to find evidence on the influence of each of these dimensions. However, the relative importance of these dimensions vary based on the type of the bank that evaluates a loan guarantee application and the institutional environments. The hybrid financial sector and the regional differ-

ences of the marketization level in China provide us an opportunity to explore the relationships between the mechanisms . While political connections tend to play a prominent role when the political logic dominants a large proportion of decisions made by a bank, the quality of corporate governance may convey a stronger signal to banks that operate mainly under the commercial logic. The transition into a market economy could change the status quo of different actors and therefore the value of social and political capital, as well as the signaling effect of good corporate governance practices. The role of financial embeddedness has been documented in advanced capitalist economies and it may consistently facilitate access to formal finance for Chinese firms as the economy became more market-oriented. Based on the above discussion we present three sets of broad hypotheses:

Hypothesis 4: *a) Banks are more likely to approve guarantees with better governed guarantors. b) The relationship in a) is moderated by marketization. c) The association will be weaker for state-owned banks.*

Hypothesis 5: *a) Banks are more likely to approve guarantees with guarantors whose political connections indicate status and access to resources and less likely to approve those indicating governance problems. b) The relationship in a) is moderated by marketization. c) The association between political ties and guarantee signing is greater for state-owned banks.*

Hypothesis 6: *a) The presence of multiplex ties and previous successful transactions increases the likelihood of a completed guarantee. b) The relationship in a) gets stronger with marketization. c) Because embeddedness depend on the accumulation of trust through a history of positive interactions between two parties, the benefit of embedded relationships will increase with the quality of governance of the guarantor firm.*

4.3 Data

The primary data set for this study is based on a file of over 100,000 guarantee transaction records from the China Stock Market and Accounting Research (CSMAR) databases collected by GTA. These records are based on a machine coding of disclosures made by all firms on the Shanghai and Shenzhen stock exchanges before 2016 in their annual reports or interim announcements. We start our study on January 1, 2008 due to the fact that before the 2006 reform, there is no representative data. After cleaning to remove duplicate transactions and matching the records with other GTA data, this yields a data set of 18,125 guarantee transactions with non-missing data.

Our unit of analysis is a guarantor-guarantee-bank triad within a given year. The dependent variable for our study is a composite measure indicating whether a guarantee contract was signed. This design is possible because the 2006 reform applies to all intended guarantee transactions, and only 50% of the 18125 unique attempts in our data showed evidence of being approved by the bank. We considered guarantees to be signed if they were coded as such in the GTA data or if the GTA data listed an actual, rather than intended, guarantee amount. We verified this decision by comparing over twenty records with their original disclosure documents provided by www.cninfo.com.cn and found that the presence of an actual amount was a consistent indicator of a completed guarantee transaction.

4.3.1 Cleaning

A major challenge of this analysis was identifying unique transactions from the raw disclosure records. Numerous transactions were recorded multiple times as firms disclosed reaching new stages of the process and the different types of disclosures often contain slightly different information. The first step in cleaning these data was to establish a best guess for the date of the attempt or actual signing of a guarantee. We did so by preferring signing date, implementation date, the date of approval by the board of directors or stockholder meeting, and report dates, in that order.

Using this composite date, we grouped the transactions by a standardized version of the names of the guarantor firm, the guarantee recipient, and the bank and used a combination of the available date information, signing status, loan amounts, and other covariates to distinguish between duplicate reports of a transaction among the three firms and a series of multiple transactions occurring within a short time. We deleted all duplicates and combined multiple transaction within the same year by summing the loan amounts.

4.3.2 Merging

The standardized names used to group the guarantee disclosure records were also used to match the organizations in our guarantee data with their appearance in other GTA data sets. These included information on: public firms' directors and interlocking firms; the top ten owners of each public firm; the controlling shareholder and ultimate controllers of each public firm; provincial marketization; the lending history of banks; and various bank and firm covariates

such as size and performance.

Creating the standardized names from both Chinese and Latin characters involved extracting individual firms from fields containing multiple organizations as well as accounting for systematic variations in firm names due to factors like branch location. We also used a handmade data set of name variants that refer to a common firm. This is especially important given that many of the largest banks are frequently referred to by an abbreviated name of as few as two characters.

4.3.3 Predictors

We measure the relationships between the guarantor firm, guarantee recipient, and lender along two dimensions. First, we use a binary indicator of the presence of a direct or indirect ownership connection to measure the financial stake that a bank might have in a guarantor. We measure interlocking directorates using a count of the number of triads linking the two firms. We also measure the presence of a direct ownership or interlock tie between the guarantee recipient and the guarantor.

We measure the embeddedness of the financial transactions between two firms using the log of: the number of guarantees a bank has approved to a guarantor, the number of guarantees that the bank has rejected involving the guarantor, and the number of loans that the guarantor has applied for from the lender. As shown in Table 4.3, the correlations between these variables are quite high, which is to be expected since the number of successes and failures both

	N	Local Ties	Central Ties	Other Ties
UC Private	9598	0.35	0.02	0.34
UC Local	5757	0.64	0.03	0.41
UC SOE	901	0.57	0.05	0.22
UC Central	1959	0.24	0.10	0.23

Table 4.1: Mean political ties by ultimate controller type.

depend directly on the number of attempts.[1]

We assess the evolving role of political connections in securing access to guaranteed loans along three dimensions. First, we code each guarantor firm based on the identity of its ultimate controller, which can be either a private person or firm, a central government organization, a local government body, or an SOE. Firms with a private ultimate controller are considered the reference category. Second, we measure political connections using the log of the number of central, local, and other government officials on the board of each guarantor firm. Finally, we also employ the longitudinal NERI index of provincial marketization to test the contingency of the impact of political capital on the progress of reforms in a given location and time (Fan et al. 2017). Table 4.1 provides a description of the number of transactions and mean political ties by ultimate controller type.

We use three proxies for the corporate governance of each firm. Unlike in advanced market economies where corporate governance is typically evaluated with reference to specific firm policies that regulate the principal-agent conflict between owners and managers (Davis and Greve 1997; Bhagat and Bolton 2008), research in emerging or transitional markets instead focuses on the tension between controlling and minority shareholders that arises from the controller's

[1] An alternative specification would be to use the number of attempts and the percentage of failures or successes. We will evaluate this specification in a future version.

contradictory desire to attract outside capital while maximizing their ability to appropriate firm revenue. Hence, our three measures seek to capture various aspects of this conflict. First, we use the log of the number of shareholders as a proxy for dispersed ownership. Second, we use the percent of the firms shares that may be freely traded as an indicator of the firm's susceptibility to market forces. This measure is especially important given that in 2005 the central government mandated that firms start to reduce their number of such closely-held shares. Finally, we follow numerous studies in the literature on finance and corporate governance in emerging markets by using the gap between an ultimate controller's voting and cash-flow rights in a firm (Boubaker et al. 2014; Gul et al. 2010; Claessens et al. 2000). In other words, we calculate ultimate controller's excess control by subtracting the sum of their direct ownership of a firm and their indirect holdings via pyramid structures from the sum of the voting rights in the target firm that they control directly or through proxies within the pyramid.

We also investigate a relational dimension of corporate governance. Only 5.5% of guarantee transactions in our data take place between an unrelated guarantor and recipient, but these observations allow us to assess banks' reaction to the absence of the conflict of interest created by transactions with related parties. This measure is based on government-mandated disclosures and indicates that the recipient is neither a member of the same business group as the guarantor nor involved in an ongoing alliance or partnership. We also include a measure to differentiate these partners from group members and treat group comembership as the baseline.

Our final set of predictors are based on lender characteristics. We measure

	N Banks	Total	Mean	Std.	Min.	Median	Max.
SOE Bank	5	7344	1468.80	451.71	895	1444	2162
Joint Stock Bank	13	7638	587.54	472.75	10	550	1566
Commercial Bank	81	2790	34.44	58.66	2	17	435
Policy Bank	3	601	200.33	95.85	144	146	311
Other Bank	31	757	24.42	23.57	2	13	62

Table 4.2: Number of transactions by bank type.

bank size in terms of log assets and bank performance using the log of their ratio of performing to non-performing loans. We also include indicators of the type of each bank as either an SOE, a joint-stock bank, a policy bank, commercial bank, or other. Table 4.2 describes the number of transactions and banks by bank type. Table 4.3 provides summary statistics and correlations for the transformed predictor variables.

4.3.4 Controls

In order to control for various forms of heterogeneity across time, location, and industry, we include fixed effects for the year of the transaction, the province of the guarantor, the industry of the guarantor, and the industry of the recipient. This latter measure was not available from GTA, so we instead attempted to parse the name of each recipient firm in order to search for relevant words and phrases. Using this method, we were able to identify an industry for 7511 out of 18215 observations. The remaining 10704 transactions were coded as a residual category rather than excluded from the analysis.

Other controls include the guarantee history of the recipient in terms of the log of the number of successful and failed guarantees, the amount of guarantee

requested in a given year, whether the recipient is a listed firm, the stock exchange on which the guarantor is traded, the type of guarantee (mutual, joint liability, or collateralized); and the age, size, and performance of the guarantor in terms of log years since foundation, log number of employees, log assets, ROE, and Tobin's Q.

4.4 Methods

We use a Bayesian multilevel logistic regression framework in order to estimate the association between the above covariates and the approval of a guarantee transaction. Compared with standard logistic regression models, which attempt to derive point estimates for the coefficients and standard errors of each predictor by maximizing the likelihood of the data, a Bayesian logistic regression combines the likelihood with prior information in order to draw a large number of samples from the resulting posterior distribution. More specifically, we use a pooled design with random intercepts for both the guarantor and bank.[2] This random intercept method is necessary to ensure the conditional independence of observations with a shared guarantor or bank and is a mainstay of the network analysis literature (Van Duijn et al. 1999; 2004). We chose a Bayesian framework primarily because even simple versions of our model failed to converge using standard frequentist software packages.

In addition to allowing for the convenient inclusion of random intercepts for

[2]We did not include a random intercept for the guarantee recipient for three reasons. First, doing so would greatly increase the complexity of our model. Second, most recipients only receive a guarantee from a single firm making it difficult to untangle the two. Third, we use covariates on this guarantee history of each recipient to account for some of the unobserved heterogeneity.

	Mean	Std.	Min	Max	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1. Guarantee Signed	0.50	0.50	0.00	1.00														
2. Market	0.00	1.66	-4.38	2.70	-.08													
3. Prior Guarantees	0.66	0.85	0.00	3.87	-.04	.06												
4. Prior Failures	0.37	0.68	0.00	3.83	-.27	.10	.73											
5. N Loan Attempts	0.76	0.81	0.00	3.50	-.12	.05	.54	.50										
6. N Shareholders	0.00	1.00	-2.49	2.70	.05	-.10	.12	.08	.15									
7. Tradable Shares	0.22	0.24	0.00	0.95	.03	-.04	-.24	-.18	-.26	-.38								
8. UC Excess Control	0.07	0.09	0.00	0.44	.06	-.04	.01	-.04	.05	-.09	-.04							
9. Unrelated Recipient	0.06	0.23	0.00	1.00	.12	-.02	.01	-.06	-.14	.01	-.02	-.01						
10. Bank Interlock	0.21	0.68	0.00	10.00	.02	.02	.01	-.03	-.01	.05	.01	-.00	-.01					
11. Bank Ownership	0.35	0.48	0.00	1.00	.02	.01	.07	.03	.14	.03	-.06	.07	-.01	.02				
12. UC Local	0.32	0.46	0.00	1.00	.04	-.06	.08	.04	.02	.24	-.18	-.26	.04	-.02	.03			
13. UC SOE	0.05	0.22	0.00	1.00	.02	-.01	.06	.08	.08	.05	-.07	.03	.04	-.02	-.02	-.16		
14. UC Cen	0.11	0.31	0.00	1.00	-.00	-.15	.03	.00	.09	.05	-.03	.10	-.06	.09	.02	-.24	-.08	
15. Local Political Ties	0.44	0.56	0.00	2.56	.02	-.04	.02	.02	.00	.26	-.13	-.16	.06	-.03	.02	.24	.05	-.12
16. Central Political Ties	0.03	0.14	0.00	1.10	.03	-.07	.03	-.02	.03	.21	-.04	-.05	.00	.09	-.04	-.01	.03	.18
17. Other Political Ties	0.34	0.52	0.00	2.83	.01	.05	.02	.02	.04	.21	-.13	.06	.03	.03	.06	.08	-.06	-.08

Table 4.3: Summary statistics and correlations for the main predictor variables

each guarantor and bank, this framework also allows for the investigation of how lenders' sensitivities to different covariates might vary by bank. In other words, it provides a compromise between assuming that all banks react similarly to different stimuli and estimating a separate model for each lender by assuming that each bank's coefficient for a particular variable is draw from the same underlying distribution. As a result, banks with few observations will tend to cluster around the mean for the entire sample while the most active banks will be assigned values that reflect their own potentially idiosyncratic criteria for approving a guarantee.

Using this framework, we fit five distinct models. The first includes only guarantor and bank random intercepts. The second includes all covariates and interactions except those involving marketization. The third includes a marketization main effect as well as interactions with all the predictors of interest. The fourth and fifth models allow for varying slopes among the corporate governance and embeddedness and political capital variables, respectively. In doing so, we chose only those variables whose 90% credibility interval does not include zero. We estimated these varying slopes across two separate fits and using only significant predictors because the inclusion of additional varying predictors quadratically increases the number of parameters in the model due to the necessity of estimating a correlation matrix with a row and column for each variable.

We fit five separate pooled logistic regression models as follows:

$$\ln(\frac{P(y_{it}=1)}{P(y_{it}=0)}) = \alpha_k + \delta_l \tag{4.1}$$

$$\ln(\frac{P(y_{it}=1)}{P(y_{it}=0)}) = \boldsymbol{\beta}\boldsymbol{x}_{1it}^{T} + \alpha_k + \delta_l \tag{4.2}$$

$$\ln(\frac{P(y_{it}=1)}{P(y_{it}=0)}) = \boldsymbol{\beta}\boldsymbol{x}_{2it}^{T} + \alpha_k + \delta_l \tag{4.3}$$

$$\ln(\frac{P(y_{it}=1)}{P(y_{it}=0)}) = \boldsymbol{\beta}\boldsymbol{x}_{2it}^{T} + \boldsymbol{\gamma}_l\boldsymbol{z}_{1it}^{T} + \alpha_k + \delta_l \tag{4.4}$$

$$\ln(\frac{P(y_{it}=1)}{P(y_{it}=0)}) = \boldsymbol{\beta}\boldsymbol{x}_{2it}^{T} + \boldsymbol{\gamma}_l\boldsymbol{z}_{2it}^{T} + \alpha_k + \delta_l \tag{4.5}$$

where i refers to a unique combination of public firm k, bank l, and a guarantee recipient, t is the year, α is public-firm intercept, δ is the bank intercept, $\boldsymbol{\beta}$ is a vector of coefficients constrained to be constant across the entire sample, $\boldsymbol{\gamma}_l$ is a vector of coefficients allowed to vary by bank, $\boldsymbol{x}_{1it}^{T}$ is a vector of covariates without the marketization main effect and interactions, $\boldsymbol{x}_{1it}^{T}$ includes these additional predictors, and $\boldsymbol{z}_{1it}^{T}$ is a vector of containing predictors related to financial embeddedness and corporate governance while $\boldsymbol{z}_{2it}^{T}$ contains variables related to political capital.

4.5 Results

This section presents the entire-sample coefficient estimates from Model 3, provides a comparison between the mean coefficients of SOE and Joint Stock banks from Models 4 and 5. and compares the observed data with predictions generated by Model 3. The main effects in Model 3 are similar in magnitude and

significance to those in Model 2 as well as in the sample averages in Models 4 and 5. Table 4.4 gives the 95% credible interval and median value for the coefficient of each financial embeddedness and corporate governance predictor as well as their interactions with each other and marketization. Table 4.4 does the same for each political relationships predictor.

Beginning with financial embeddedness, we find that the main coefficient for number of previous signed guarantees is weakly negative, the number of unsigned is close to zero, and the number of loans attempts is significantly positive. The only significant interaction with marketization is the number of previous unsigned attempts, which is positive. Taken together, these results provide limited support for Hypothesis 6 in that the frequency of other interactions between the firms appears to encourage guarantee relationships. In terms of network ties, the main effect of ownership is close to zero, but firms from more marketized provinces in fact appear *less* likely to receive guarantees from shareholder banks. The density of interlock ties, however, has a weakly positive association that appears to increase for firms in more marketized provinces. This could reflect a contrast between relying on ownership, which might create a conflict of interest, and the density of shared director connections, which might help facilitate trust and communication between firms. This appears to be supported by the positive interaction between ownership and number of failed transactions as well as the negative one between prior failures and the density of interlock ties.

Turning to corporate governance and its interaction with the above variables, we find that the main effects of all three variables have the expected direction are non-zero within a reasonable certainty. The marketization interactions

Coefficient	2.5%	Median	97.5%
Market	-0.33	-0.11	0.12
Prior Guarantees	-0.34	-0.15	0.04
Market:Prior Guarantees	-0.09	-0.02	0.04
Prior Failures	-0.30	-0.07	0.17
Market:Prior Failures	0.06	0.13	0.21
N Loan Attempts	0.07	0.23	0.39
Market:N Loan Attempts	-0.04	0.01	0.06
Bank Ownership	-0.10	0.08	0.26
Market:Bank Ownership	-0.11	-0.05	0.01
Prior Guarantees:Bank Ownership	-0.23	-0.03	0.16
Prior Failures:Bank Ownership	0.06	0.29	0.52
N Loan Attempt:Bank Ownership	-0.24	-0.10	0.05
Bank Interlock Triads	-0.02	0.10	0.22
Market:Bank Interlock Triads	-0.01	0.04	0.09
Prior Guarantees:Bank Interlock Triads	-0.09	0.04	0.19
Prior Failures:Bank Interlock Triads	-0.32	-0.13	0.05
N Loan Attempts:Bank Interlock Triads	-0.18	-0.05	0.07
N Shareholders	0.05	0.20	0.34
Market:N Shareholders	0.08	0.15	0.22
N Shareholders:Prior Guarantees	-0.08	0.05	0.19
N Shareholders:Prior Failures	-0.37	-0.21	-0.05
N Shareholders:N Loan Attempts	-0.07	0.02	0.11
Tradable Shares	-0.05	0.37	0.77
Market:Tradable Shares	0.30	0.48	0.67
Tradable Shares:Prior Guarantees	0.50	1.05	1.60
Tradable Shares:Prior Failures	-2.76	-2.04	-1.30
Tradable Shares:N Loan Attempts	-0.10	0.29	0.70
UC Excess Control	-3.11	-1.63	-0.12
Market:UC Excess Control	0.09	0.80	1.55
UC Excess Control:Prior Guarantees	0.19	1.47	2.71
UC Excess Control:Prior Failures	-1.70	-0.11	1.51
UC Excess Control:N Loan Attempts	-2.08	-1.06	-0.05
Unrelated Recipient	1.14	1.54	1.94
Market:Unrelated Recipient	0.32	0.54	0.76
Unrelated Recipient:Prior Guarantees	0.44	0.99	1.62
Unrelated Recipient:Prior Failures	-1.52	-0.74	0.01
Unrelated Recipient:N Loan Attempts	-0.96	-0.48	0.01

Table 4.4: Median and 95% posterior intervals for entire-sample coefficients related to financial embeddedness and political ties.

are significant with the expected positive sign, suggesting that well governed firms from more marketized provinces are even more likely to receive guarantees while poorly governed firms from the same areas are less strongly penalized. These results provide strong support for Hypothesis 4. Interacted with the financial embeddedness variables, we find that good corporate governance appears to aid the formation of trust and stability through repeated exchanges insofar as that firms with fewer non-tradable shares are more likely to benefit from interacting with the same firm. A similar pattern holds for transactions involving an unrelated recipient.

Table 4.5 shows the main effect and two- and three-way interactions between the types of ultimate controllers (UCs), the log number of political ties, and marketization. Marketization has the same non-significant coefficient, and the associations between political ties and signing status for reference category of private UCs are all close to zero in both their main effects and interactions with marketization. Turning to publicly traded firms with local government UCs, we find that they are no more or less likely than privately controlled firms to receive guarantees, but appear to benefit strongly from the category of residual political ties. This association diminishes with greater marketization. The third UC category, that of other state-owned firms, provides surprising results. This firms appear to be more attractive guarantors when they are located in more marketized provinces, especially when they have connections to local politicians. This suggests that the typical image of SOEs, especially those bound to the conflicting interests of local politicians, as plagued by inefficiency and governance problems perhaps no longer holds in more marketized areas. Finally, high-status firms with central government UCs are more likely to receive guarantees, and marketization only increases the association. Given this strong main

Coefficient	2.5%	Median	97.5%
Market	-0.33	-0.11	0.12
Local Political Ties	-0.40	0.10	0.64
Central Political Ties	-4.32	-1.17	1.89
Other Political Ties	-0.80	-0.20	0.42
Market:Local Political Ties	-0.24	-0.04	0.17
Market:Central Political Ties	-1.53	-0.28	0.94
Market:Other Political Ties	-0.12	0.11	0.34
UC Local	-0.67	-0.09	0.46
Market:UC Local	-0.29	-0.05	0.18
UC Local:Local Political Ties	-1.14	-0.48	0.19
UC Local:Central Political Ties	-3.57	-0.22	3.22
UC Local:Other Political Ties	0.41	1.30	2.22
Market:UC Local:Local Political Ties	-0.30	-0.06	0.19
Market:UC Local:Central Political Ties	-0.92	0.54	2.05
Market:UC Local:Other Political Ties	-0.84	-0.53	-0.20
UC SOE	-0.84	-0.08	0.67
Market:UC SOE	0.06	0.47	0.88
UC SOE:Local Political Ties	-0.41	0.67	1.88
UC SOE:Central Political Ties	-5.65	-0.90	4.12
UC SOE:Other Political Ties	-1.04	0.33	1.71
Market:UC SOE:Local Political Ties	0.37	0.88	1.40
Market:UC SOE:Central Political Ties	-2.61	-0.67	1.19
Market:UC SOE:Other Political Ties	-1.86	-1.12	-0.37
UC Central	0.20	0.96	1.73
Market:UC Central	0.24	0.54	0.82
UC Central:Local Political Ties	-2.40	-1.10	0.10
UC Central:Central Political Ties	-2.58	1.29	5.29
UC Central:Other Political Ties	-0.95	0.46	1.93
Market:UC Central:Local Political Ties	-1.64	-1.07	-0.55
Market:UC Central:Central Political Ties	-0.23	1.57	3.46
Market:UC Central:Other Political Ties	-0.38	0.41	1.24

Table 4.5: Median and 95% posterior intervals for entire-sample coefficients related to political connections.

effect, it is perhaps unsurprising that most of the political ties coefficients have little additional predictive power.

4.5.1 Variation Among Banks

Turning now to between-bank variation in the estimated coefficients in Models 4 and 5, Table 4.6 presents the distributions of the mean difference in coefficients across the draws from the posterior distribution. In general there is little difference between the two groups, but three results are notable. First, joint stock banks react more positively to guarantees to unrelated firms in the context of repeated transactions and more market-oriented institutions. They are also less prone to sign guarantees to firms with local UCs and political ties. They are, however, more likely approve a transaction with a state-owned ultimate controller with local ties from a more marketized province.

4.5.2 Predictive Checks

We carried out two predictive checks to evaluate the fit our model. First, we checked the accuracy our predicted versus observed outcomes and found that using a cutoff of 50% probability our model was correct in 85% of cases. Figure 4.1 shows the frequency of different predicted probabilities, defined as the fraction of posterior draws in which that transaction was more likely to be signed than not, grouped by observed signing status. Furthermore, the clustering of observations at the two extremes of the plot indicates that the model predicted many transactions with a high degree of certainty. Our models make

	5%	Median	95%
N Shareholders	-0.11	-0.02	0.05
Market:N Shareholders	-0.00	0.04	0.10
N Shareholders:Prior Failures	-0.09	0.02	0.14
Tradable Shares	-0.06	0.15	0.54
Market:Tradable Shares	-0.15	0.00	0.15
Tradable Shares:Prior Guarantees	-0.20	0.01	0.30
Tradable Shares:Prior Failures	-0.23	0.15	0.75
UC Excess Control	-0.54	-0.01	0.35
Market:UC Excess Control	-0.21	0.01	0.28
UC Excess Control:Prior Guarantees	-0.23	0.10	0.85
UC Excess Control:N Loan Attempts	-0.18	0.14	0.97
Unrelated Recipient	-0.29	0.03	0.48
Market:Unrelated Recipient	-0.03	0.14	0.49
Unrelated Recipient:Prior Guarantees	-0.03	0.56	1.34
Unrelated Recipient:Prior Failures	-0.26	0.09	1.00
Unrelated Recipient:N Loan Attempts	-0.87	-0.11	0.19
UC Local:Local Political Ties	-0.31	-0.06	0.06
UC Local:Other Political Ties	-0.63	-0.33	-0.04
Market:UC Local:Local Political Ties	-0.16	-0.05	0.04
Market:UC Local:Other Political Ties	-0.08	0.05	0.24
Market:UC SOE	-0.14	0.02	0.29
UC SOE:Local Political Ties	-0.07	0.16	0.85
UC SOE:Other Political Ties	-1.26	-0.15	0.67
Market:UC SOE:Local Political Ties	1.12	2.32	3.91
Market:UC SOE:Other Political Ties	-0.47	0.01	0.47
UC Central	-0.09	0.00	0.14
Market:UC Central	-0.07	-0.00	0.03
UC Central:Local Political Ties	-0.49	-0.02	0.34
UC Central:Central Political Ties	-0.93	0.17	1.92
UC Central:Other Political Ties	-0.42	0.06	0.81
Market:UC Central:Local Political Ties	-0.39	-0.03	0.13
Market:UC Central:Central Political Ties	-2.40	-0.60	0.56
Market:UC Central:Other Political Ties	-0.51	0.04	0.74

Table 4.6: Median and 90% credible intervals for the mean difference between Joint Stock and SOE bank coefficients.

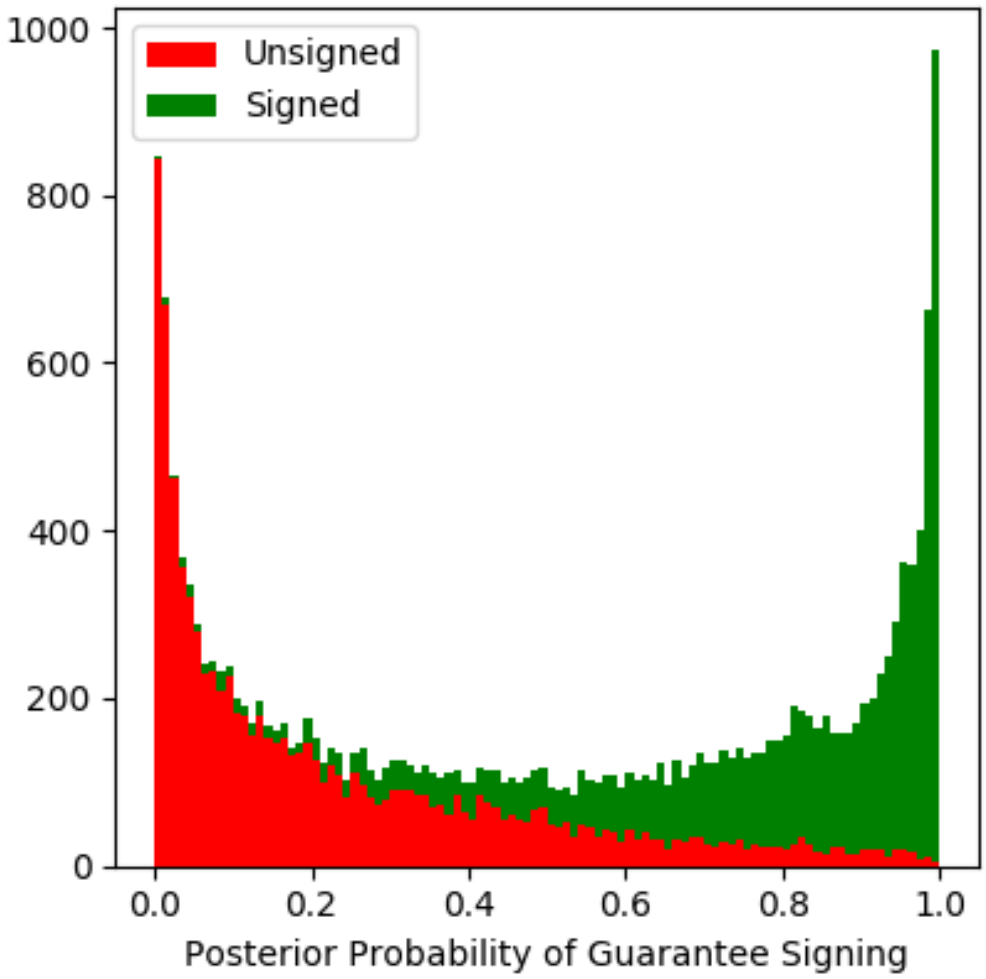

Figure 4.1: Mean predicted outcome across all posterior draws versus actual signing status.

heavy use of firm's institutional environment. However, this value varies both with province and year, suggesting that unobserved factors common to a particular location in a given year might confound these coefficient estimates. Figure 4.2 compares the predicted proportion of signed guarantees in each province in 2008 and 2015 to the corresponding observed success rate and provides the percentile rank of the observed proportion relative to the distribution of predicted values. The model appears to give a reasonable fit for most cases, but

there are enough examples of poor fit to suggest that making better use of over-time variation will improve the fit of the model. We plan to implement a fully longitudinal model in a future revision.

4.6 Discussion

Our study uses transaction-level information to shed light on how banks respond to various measures of firm governance, inter-firm relationships, and political connections. Our perspective assumes that Chinese banks, which operate under close supervision of the state, have internalized the government's desire to reduce the expropriation of minority shareholders through related party transactions and hence will prefer to provide guaranteed loans when they perceive them as less risky and less likely to serve as a means of tunneling. Overall, our findings are consistent with prior work but also include novel results that add nuance to our understanding of the role of various institutional logics in emerging markets in general and China in particular. This section will parse the output of our models in more detail before describing the limitations of our approach and possible improvements.

We find broad support for the importance of corporate governance and each of the main effects has the expected sign. These associations are moderated by marketization in that the coefficients of indicators of relatively minor conflicts of interest between controlling and minority shareholders increase in magnitude in more market-oriented provinces. However, our sole measure of poor governance, UC Excess control, has a positive interaction with marketization perhaps indicating that firms' interactions with market-based institutions can mitigate

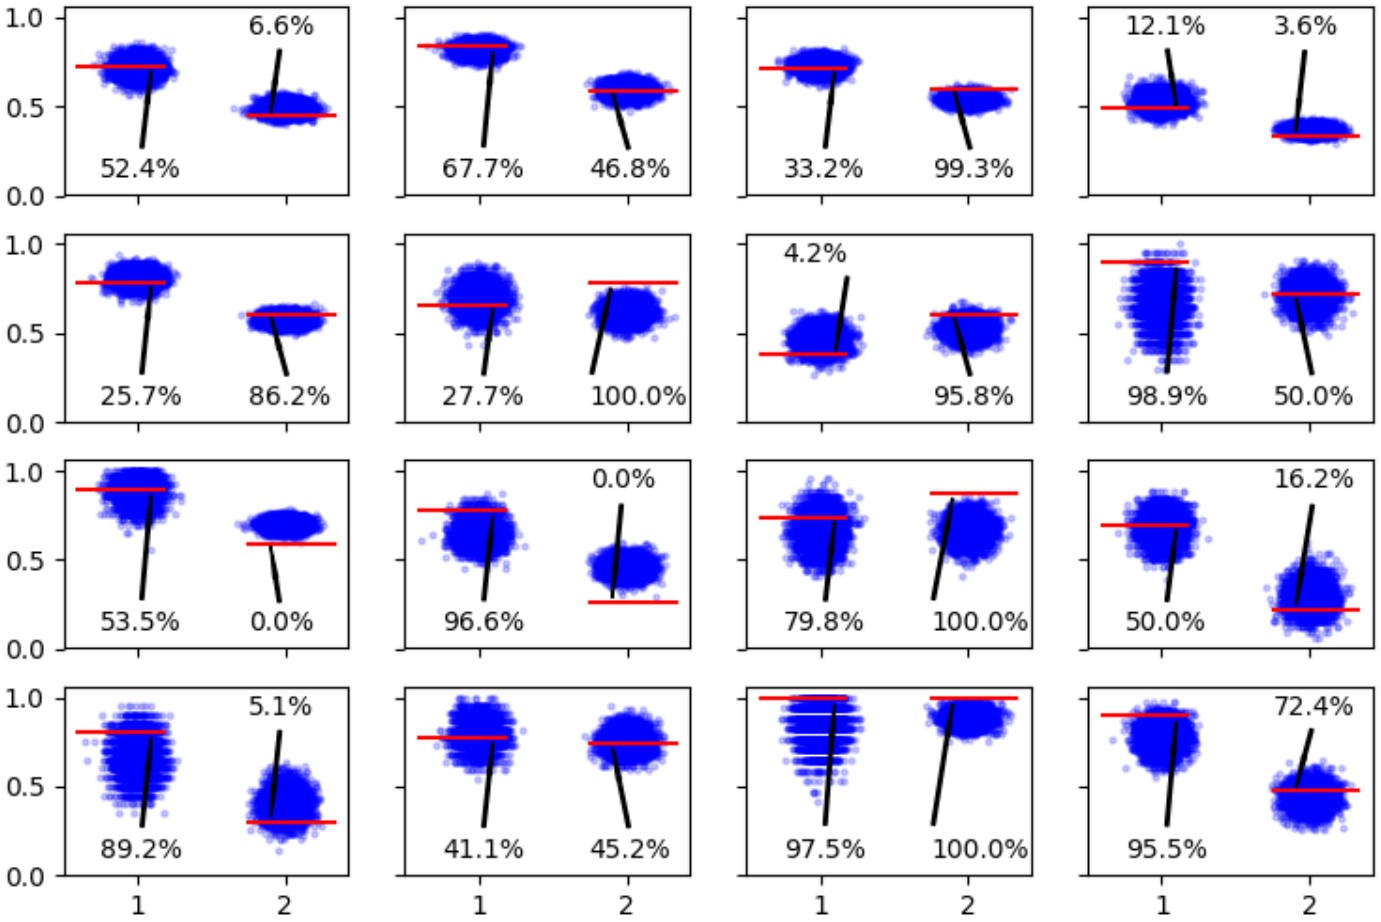

Figure 4.2: Bars indicate the observed proportion of signed guarantees in each province-year. Each dot indicates a predicted proportion derived from a draw of the joint posterior distribution of all coefficients.

the ultimate controllers' incentives to expropriate.

Our analysis of the embeddedness of these transactions in a wider array of multiplex ties paints a more complicated picture. In terms of repeated interactions, the main effects of Prior Guarantees, Prior Failures, and Bank Ownership all close to zero. The number of prior loans sought by the guarantor from the bank is associated with a higher chance of approval, and there is weak evidence that a greater density of interlock connections makes a bank more willing to lend. When interacted with provincial marketization, the coefficients for Prior Guarantees and N Loan Attempts are close to zero. The interaction with ownership is weakly negative, indicating that banks perhaps prefer to avoid their own conflict of interest when dealing with firms from more developed areas. Interlock density, however, appears to become more important in more market-orient areas, which would be consistent with the role of interlock ties as channels of communication and reputation-building. Finally, Prior Failures has a strongly positive interaction with Marketization. This could be due to the fact that banks have a quota of the number of loans they may make to a given firm, and when searching for new loan recipients they appear to prefer not just the firms they have previously rejected, but only those from more developed areas.

Our next set of results concerns the interactions between corporate governance and repeated transactions. Reflecting the idea that trust builds up over the course of repeated positive interactions, we find somewhat mixed evidence that firms less apparent conflicts of interest are more likely to be approved by banks with whom they have a history of interaction. However, there is also a positive interaction with UC Excess Control. The high correlations among Prior Guarantees, Prior Failures, and N Loan Attempts, however, complicate the in-

terpretation of these coefficients since the positive interaction between UC Excess Control and Prior Guarantees would be somewhat offset by the interaction with N Loan Attempts.

The role of political connections in securing guaranteed loans appears even more convoluted. They appear to be completely irrelevant for privately held firms at any marketization level. When considering public firm guarantors with a local government ultimate controller, we find evidence in favor of the central prediction of market transition theory, namely that political ties become less salient with increased marketization. However, we find the opposite trend among transactions involving firms with an SOE ultimate controller. In more marketized areas, not only do such firms appear more likely to have their application approved, but the also benefit more from ties with local politicians. In contrast, their ties with other miscellaneous state entities seem to become a liability in more developed provinces. Finally, firms with ultimate controllers affiliated with the central government are more likely to receive guarantees in general and increasingly so with more marketization.

Finally, our models also provided us with the opportunity to evaluate banks' varying sensitivity to the above factors. Our expectation was that state-owned banks would give more weight to political connections while joint-stock banks, which are more subject to market discipline, would rely more on firm governance and embedded relationships. Our results show that coefficient values do not differ significantly between the two groups with three exceptions. The first two provide support for the above hypothesis insofar as SOE banks give more weight to the political ties of local government firms while joint stock banks are more receptive to guarantees made to unrelated parties. A more surpris-

ing result is that the positive coefficient of the interaction between marketization and local political ties for firms with SOE ultimate controllers is largely due to joint stock banks, not their state-owned counterparts. One explanation for this finding, especially when contrasted with the uniform preferential treatment received by central government firms, is that UC SOE firms have undergone meaningful reforms that make their political capital more of an asset than a liability.

The uniqueness of our data set has allowed us to test the relevance of numerous factors at the transaction level, but our study nevertheless has important limitation, some of which are inherent in the data. First and most seriously, our analysis is largely limited to an investigation of bank and guarantor characteristics despite the fact that the guarantee recipient is arguably the most important actor in the transaction. We have tried to mitigate this issue by inferring the guarantee firm's industry where possible and by including variables on the relationship between the guarantor and the guarantee recipient, but we still lack relevant data on the recipient's performance or governance structure. Second, when cleaning our data, we were able to identify over 60,000 unique transactions. However, only 25,000 of these provided information on the identity of the lender, which was necessary in order to derive our variables on guarantor-bank relationships as well as bank characteristics. This could create a potential selection bias, which we will explore in subsequent work. Finally, posterior predictive checks of our model have found that it does not always accurately predict the proportion of approved guarantees in a given province in a given year, which suggests that a fully longitudinal model might be more appropriate.

4.7 Conclusion

Our results are consistent with the literatures on corporate governance and political ties in China insofar as they demonstrate a growing reliance on market-based governance mechanisms and a complex, evolving role for political connections. We also contribute to the growing literature on the sociology of financial markets in China by testing the role of embeddedness in facilitating guaranteed loan transactions, and found evidence for the relevance of multiplex ties in facilitating exchange, especially for better governed guarantors.

CHAPTER 5

CONCLUSION

This dissertation has sought to provide answers to a series of questions regarding the relationships among firm governance, inter-organizational networks, business groups, and financial markets. In Chapter 2, I find that there is evidence that the results of the community detection procedure for inferring business group membership do indeed correspond to socially meaningful clusters. Chapter 3 bolsters this finding by providing showing that the influence of the inferred business groups can be detected in a separate network of pairwise stock return correlations. Finally, the results from Chapter 4 present a complicated picture of the variety of forces at work in China's financial sector. Firm governance, political ties, and financial embedded all predict the approval of a loan guarantee to varying degrees. There is also limited evidence that different types of banks respond differently to different factors.

APPENDIX A

APPENDIX: COMMUNITY DETECTION

The community-detection procedure used in this paper can be broken down into three steps. First, I adapted the multiplex modularity function described in Equation 2 to fit bipartite and directed networks. Next, I implemented Newman's 2006 leading-eigenvector method as a heuristic for identifying a high quality partition. Finally, I refined this partition using two separate runs of a modified Kerrighan-Lin algorithm and also ensured that the resulting partitions were connected.

In its most general form, the modularity of a given partition and network is proportional to the total weight of observed ties minus the expected value for all potential edges in the network. The key contribution of Mucha et al. (2010) is the principled elaboration of this basic equation to include manually specified connections between the same node in different networks. The specific form of this elaboration does not vary with network type. Thus adapting the modularity function to analyze bipartite and directed networks only requires the correct specification of the expected value of each potential dyad and the appropriate normalization.

The standard modularity equation for undirected networks is

$$Q = \frac{1}{2m} \sum_{ij} \left(A_{ij} - \gamma \frac{k_i k_j}{2m} \right) \delta_{c_{is} c_j} \tag{A.1}$$

where $2m = \sum_{ij} A_{ij}$ and k_i equals the sum of the weights of all of the edges involving node i, δ_{ij} is the Kroenecker delta such that $\delta_{ij} = 1$ if $i = j$ and 0

otherwise, γ is the resolution parameter, and c_i is the index of the community of vertex i (Newman 2006).

In directed networks the expected value of a tie is similar, but rather than use the overall degree of each node, the likelihood of an edge in each direction is proportional to the indegree of the target nodes as well as the outdegree of the source node, as in Equation A.2:

$$Q = \frac{1}{m} \sum_{ij} \left(A_{ij} - \gamma \frac{k_i^{in} k_j^{out}}{m} \right) \delta_{c_i c_j} \tag{A.2}$$

where A_{ij} is the weight of the edge from originating at i and ending at j, $m = \sum_{ij} A_{ij}$, $k_i^{in} = \sum_j A_{ij}$, $k_i^{out} = \sum_j A_{ji}$, and γ, δ_{ij}, and c_i are as defined above (Leicht and Newman 2008). This alternate specification allows the algorithm to distinguish between nodes with many incoming versus outgoing ties; and in more concrete terms it minimizes the impact of ownership edges originating at owners that have holdings in a large number of public firms, such as institutional investors, just as it will maximize the import of an edge ending at a node with a small number of owners. This last point is especially important given the pruning of owners with only one connection to the broader network. Because of this editing, the expected values for each edge are not derived from the overall ownership data for a company, but rather from the what those data tell us about the underlying relational structure of public firms. As a result, a multiplex community-detection algorithm used on the edited data will be less likely to assign a weakly connected node in the ownership network to its own community, and instead focus on the nature of its connection to the broader ownership structure, as well as its position in the director-firm affiliation network. The ability to rely on the director-firm network to help resolve such ambiguous situations

is an important strength of the method.

Finally, bipartite networks are defined by the presence of two disjoint sets of nodes that may only form ties with nodes of the other type, so for example in our director-firm affiliation network directors may not be connected to other directors, but only to firms. Of course, there are certainly important social connections between directors, so our network is bipartite only insofar as it focuses on a specific type of relationship, namely employment as a member of the board of directors. The modularity function for bipartite networks explicitly models this restriction by assigning an expected value of zero to edges between nodes of the same type, as show in Equation A.3:

$$Q = \frac{1}{2\mu} \sum_{ij} \left(A_{ij} - \gamma b_{ij} \frac{k_i k_j}{m} \right) \delta_{c_{is} c_j} \tag{A.3}$$

where A_{ij}, m, k_i, γ, δ_{ij}, and c_i have the same definitions as in the undirected network and $b_{ij} = 1$ if i and j are not the same type and 0 otherwise (Barber 2007). Note that the denominator of the expected edge weight is m instead of $2m$. This reflects the fact that the probability of a tie between two nodes is higher than in an undirected network with a single type of node, conditional on knowing that the two nodes are of different types.

After identifying the appropriate modularity equations, the next step is to choose the γ and ω parameters and combine all three elements to form the overall multiplex modularity function given in Equation 2. As described in Section 2.3.4, I chose to set $\gamma = 1$ and $\omega = 1$. This equation is then used to create the modularity matrix $\boldsymbol{B}$. A final adjustment is to add the modularity matrix to its transpose and divide by 2. This ensures that the modularity matrix is symmet-

ric, but does not change the results of the analysis since it is not possible for i to be in the same partition as j while is j not in the same partition as i.

From this point any one of a number different heuristics can be used to find a high-quality partitioning, and I used Newman's (2006) leading-eigenvector method with two additional refinements. As the name implies, this method uses the eigenvector of $\boldsymbol{B}$ corresponding to the largest eigenvalue as an approximation of the optimal two-way partition of the graph. Thus, nodes with a value less than zero are assigned to one group and the remaining nodes to the other. This approximation provides a good first guess about the best split, and the next step recommended by Newman (2006) is to refine the partition according to the KL-algorithm, which switches the partition assignment of each node exactly once by choosing the unmoved node that would yield the highest increase or smallest decrease in modularity. It then returns the intermediate partition with the highest modularity (Kernighan and Lin 1970). This process is repeated until no improvement is found. Next, I added the step of checking that the resulting partitions are in fact connected across both networks, or in other words that it is possible to reach a node from any other node in the partition using only intra-community ties. If partitions were found to be disconnected, I flipped the partition assignment of each node outside the largest component until no unconnected nodes remained. I do not have a full explanation for why the algorithm would otherwise occasionally return disconnected partitions, but it might be related to the form of the bipartite modularity equation, which does not penalize lumping together unconnected nodes of the same type.

Next, each of the two partitions is further subdivided and refined until doing so does not yield an increase in modularity. This is done using a modified subset

of the original modularity matrix, $\boldsymbol{B}^{(g)}$, such that $B_{ij}^{(g)} = B_{ij} - \delta_{ij} \sum_{k \in g} B_{ik}$ for all $i, j \in g$. After identifying these indivisible partitions, I performed a final set of KL-refinements using a community-level network where a partition is a node and edges are weighted by the number of ties between vertices in each partition. More specifically, I began with the pair of partitions with the greatest number of connections and shuffled nodes between them as described in the previous paragraph. I then moved to the partition dyad with the next highest number of connections and so forth. After exhausting all the community-dyads with more than one shared edge, I repeated the process until no changes were made. I also employed this process to refine the higher-level location based partition.

BIBLIOGRAPHY

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 Nedal Al-Fayoumi. 2009. "Market Value, Book Value and Earnings: Is Bank Efficiency a Missing Link?" *Managerial Finance* 35:156–179.

Aguilera, Ruth V and Gregory Jackson. 2010. "Comparative and International Corporate Governance." *The Academy of Management Annals* 4:485–556.

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-Khouri, Ritab. 2006. "Corporate Governance and Firms Value in Emerging Markets: The Case of Jordan." *Journal of Transnational Management* 12:25–49.

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.

Allen, Franklin, Jun Qian, and Xian Gu. 2017. "An Overview of China's Financial System." *Annual Review of Financial Economics* 9:191–231.

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.

An, Heng and Ting Zhang. 2013. "Stock Price Synchronicity, Crash Risk, and Institutional Investors." *Journal of Corporate Finance* 21:1–15.

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.

Barber, Michael J. 2007. "Modularity and community detection in bipartite networks." *Physical Review E* 76:066102.

Bassett, Danielle S, Mason A Porter, Nicholas F Wymbs, Scott T Grafton, Jean M Carlson, and Peter J Mucha. 2013. "Robust detection of dynamic community structure in networks." *Chaos: An Interdisciplinary Journal of Nonlinear Science* 23:013142.

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.

Berkman, Henk, Rebel Cole, and Jiang Lawrence Fu. 2005. "Agency conflicts, expropriation and firm value: Evidence from securities-market regulation in China." .

Berkman, Henk, Rebel A Cole, and Lawrence J Fu. 2009. "Expropriation through loan guarantees to related parties: Evidence from China." *Journal of Banking and Finance* 33:141–156.

Berkman, Henk, Rebel A Cole, and Lawrence J Fu. 2010. "Political connections and minority-shareholder protection: Evidence from securities-market regulation in China." *Journal of Financial and Quantitative Analysis* 45:1391–1417.

Bertrand, Marianne, Paras Mehta, and Sendhil Mullainathan. 2002. "Ferreting Out Tunneling: An Application to Indian Business Groups." *The Quarterly Journal of Economics* 117:121–148.

Bhagat, Sanjai and Brian Bolton. 2008. "Corporate Fovernance and Firm Performance." *Journal of Corporate Finance* 14:257–273.

Biggart, Nicole W. 1990. "Institutionalized Patrimonialism in Korean Business." *Comparative Social Research* 12:113–133.

Blondel, Vincent D, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. 2008. "Fast unfolding of communities in large networks." *Journal of statistical mechanics: theory and experiment* 2008:P10008.

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.

Calomiris, Charles W, Raymond Fisman, and Yongxiang Wang. 2010. "Profiting from government stakes in a command Economy: Evidence from Chinese asset sales." *Journal of Financial Economics* 96:399–412.

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.

Carney, Michael, Daniel Shapiro, and Yao Tang. 2009. "Business group performance in China: Ownership and temporal considerations." *Management and Organization Review* 5:167–193.

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.

Carruthers, Bruce G and Jeong-Chul Kim. 2011. "The sociology of Finance." *Annual Review of Sociology* 37:239–259.

Chan, Kalok and Allaudeen Hameed. 2006. "Stock Price Synchronicity and Analyst Coverage in Emerging Markets." *Journal of Financial Economics* 80:115–147.

Chekir, Hamouda and Ishac Diwan. 2014. "Crony Capitalism in Egypt." *Journal of Globalization and Development* 5:177–211.

Chen, Xin, Jakob Arnoldi, and Chaohong Na. 2015. "Governance structure and related party loan guarantees: The case of Chinese family business groups." *Management and Organization Review* 11:599–619.

Cheung, Yan-Leung, P Raghavendra Rau, and Aris Stouraitis. 2009. "Helping hand or grabbing hand? Central vs. local government shareholders in Chinese listed firms." *Review of Finance* 14:669–694.

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.

Claessens, Stijn, Simeon Djankov, and Larry HP Lang. 2000. "The Separation of Ownership and Control in East Asian Corporations." *Journal of Financial Economics* 58:81–112.

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. 2010. *Business groups in Turkey*. Oxford University Press.

Colpan, Asli M, Takashi Hikino, and James R Lincoln. 2010. *The Oxford Handbook of Business Groups*. Oxford University Press.

Davis, Gerald F. 2012. "Politics and Financial Markets." *The Oxford Handbook of the Sociology of Finance* p. 33.

Davis, Gerald F and Henrich R Greve. 1997. "Corporate Elite Networks and Governance Changes in the 1980s." *American Journal of Sociology* 103:1–37.

De la Torre, Augusto, Juan Carlos Gozzi, and Sergio L Schmukler. 2007. "Stock market development under globalization: Whither the gains from reforms?" *Journal of Banking and Finance* 31:1731–1754.

De Long, J Bradford, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann. 1990. "Noise Trader Risk in Financial Markets." *Journal of Political Economy* 98:703–738.

Demirgüç-Kunt, Asli and Vojislav Maksimovic. 1998. "Law, Finance, and firm growth." *The Journal of Finance* 53:2107–2137.

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.

Durnev, Art, Randall Morck, and Bernard Yeung. 2004. "Value-enhancing Capital Budgeting and Firm-specific Stock Return Variation." *The Journal of Finance* 59:65–105.

Durnev, Artyom, Randall Morck, Bernard Yeung, and Paul Zarowin. 2003. "DOes Greater Firm-Specific Return Variation Mean More or Less Informed Stock Pricing?" *Journal of Accounting Research* 41:797–836.

El Mehdi, Imen. 2007. "Empirical evidence on corporate governance and corporate performance in Tunisia." *Corporate Governance: An International Review* 15:1429–1441.

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.

Elfenbein, Daniel W. and Todd R. Zenger. 2014. "What Is a Relationship Worth? Repeated Exchange and the Development and Deployment of Relational Capital." *Organization Science* 25:222–244.

Elsayed, Khaled. 2010. "A Multi-theory Perspective of Board Leadership Structure: What Does the Egyptian Corporate Governance Context Tell Us?" *British Journal of Management* 21:80–99.

Faccio, Mara. 2006. "Politically Connected Firms." *The American Economic Review* 96:369–386.

Fama, Eugene F and Michael C Jensen. 1985. "Organizational Forms and Investment Decisions." *Journal of Financial Economics* 14:101–119.

Fan, Gang, Xiaolu Wang, and Jingwen Yu. 2017. *NERI Index of Marketization of Chinas Provinces: 2016 Report*. Social Sciences Academic Press.

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.

Fan, Joseph PH, Li Jin, and Guojian Zheng. 2016. "Revisiting the bright and dark sides of capital flows in business groups." *Journal of business ethics* 134:509–528.

Farooq, Omar, Mohamed Derrabi, and Monir Naciri. 2013. "Corporate governance and liquidity: Pre-and post-crisis analysis from the MENA region." *Review of Middle East Economics and Finance* 8:1–19.

Feld, Scott L. 1981. "The focused organization of social ties." *American journal of sociology* pp. 1015–1035.

Ferreira, Miguel A and Paul A Laux. 2007. "Corporate Governance, Idiosyncratic Risk, and Information Flow." *The Journal of Finance* 62:951–989.

Firth, Michael, Chen Lin, Ping Liu, and Sonia ML Wong. 2009. "Inside the black box: Bank credit allocation in Chinas private sector." *Journal of Banking and Finance* 33:1144–1155.

Fisman, Raymond. 2001. "Estimating the Value of Political Connections." *The American Economic Review* 91:1095–1102.

Fisman, Raymond and Yongxiang Wang. 2010. "Trading favors within Chinese business groups." *American Economic Review* 100:429–33.

Fortunato, Santo. 2010. "Community detection in graphs." *Physics reports* 486:75–174.

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. 2006. *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press.

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.

Haggard, K Stephen, Xiumin Martin, and Raynolde Pereira. 2008. "Does Voluntary Disclosure Improve Stock Price Informativeness?" *Financial Management* 37:747–768.

Hail, Luzi and Christian Leuz. 2006. "International differences in the cost of equity capital: Do legal institutions and securities regulation matter?" *Journal of accounting research* 44:485–531.

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.

Haveman, Heather A, Nan Jia, Jing Shi, and Yongxiang Wang. 2017. "The dynamics of political embeddedness in China." *Administrative Science Quarterly* 62:67–104.

He, Jia, Xinyang Mao, Oliver M Rui, and Xiaolei Zha. 2013. "Business Groups in China." *Journal of Corporate Finance* 22:166–192.

Hearn, Bruce. 2013. "The determinants of director remuneration, executive tenure and individual executive disclosure in North African IPO firms." *Research in International Business and Finance* 27:162–182.

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 groupsThe case of Africa." *Journal of World Business* .

Henry, Peter Blair. 2000. "Do stock market liberalizations cause investment booms?" *Journal of Financial Economics* 58:301–334.

Heydemann, Steven. 2004. *Networks of Privilege in the Middle East: The Politics of Economic Reform Revisited*. Springer.

Hicks, Jacqueline, Vincent A Traag, and Ridho Reinanda. 2015. "Turning digitised newspapers into networks of political elites." *Asian Journal of Social Science* 43:567–587.

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.

Huang, Wei. 2016. "Tunneling through related-party loan guarantees: evidence from a quasi-experiment in China." *Review of Quantitative Finance and Accounting* 47:857–884.

Hutton, Amy P, Alan J Marcus, and Hassan Tehranian. 2009. "Opaque Financial Reports, R^2, and Crash Risk." *Journal of Financial Economics* 94:67–86.

Jia, Nan, Jing Shi, and Yongxiang Wang. 2013. "Coinsurance within business groups: Evidence from related party transactions in an emerging market." *Management Science* 59:2295–2313.

Jiang, Guohua, Charles MC Lee, and Heng Yue. 2010. "Tunneling through intercorporate loans: The China experience." *Journal of Financial Economics* 98:1–20.

Jin, Li and Stewart C Myers. 2006. "R^2 around the World: New Theory and New Tests." *Journal of Financial Economics* 79:257–292.

Johnson, Simon and Todd Mitton. 2003. "Cronyism and Capital Controls: Evidence from Malaysia." *Journal of Financial Economics* 67:351–382.

Keister, Lisa A. 2001. "Exchange structures in transition: Lending and trade relations in Chinese business groups." *American Sociological Review* pp. 336–360.

Kernighan, Brian W and Shen Lin. 1970. "An efficient heuristic procedure for partitioning graphs." *Bell system technical journal* 49:291–307.

Khanna, Tarun and Krishna Palepu. 2000. "Is group affiliation profitable in emerging markets? An analysis of diversified Indian business groups." *The Journal of Finance* 55:867–891.

Khanna, Tarun and Jan W Rivkin. 2001. "Estimating the performance effects of business groups in emerging markets." *Strategic management journal* 22:45–74.

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.

Kogut, Bruce Mitchel. 2012. *The small worlds of corporate governance*. MIT Press.

La Porta, Rafael, Florencio Lopez-de Silanes, and Andrei Shleifer. 1999. "Corporate ownership around the world." *The journal of Finance* 54:471–517.

La Porta, Rafael, Florencio Lopez-de Silanes, and Andrei Shleifer. 2008. "The Economic Consequences of Legal Origins." *Journal of Economic Literature* 46:285–332.

Lagoarde-Segot, Thomas. 2013. "Does stock market development always improve firm-levelf Financing? Evidence from Tunisia." *Research in International Business and Finance* 27:183–208.

Lancichinetti, Andrea and Santo Fortunato. 2012. "Consensus clustering in complex networks." *Scientific reports* 2.

Lancichinetti, Andrea, Filippo Radicchi, José J Ramasco, and Santo Fortunato. 2011. "Finding statistically significant communities in networks." *PloS one* 6:e18961.

Lee, Chang Kil and David Strang. 2006. "The international diffusion of public-sector downsizing: Network emulation and theory-driven learning." *International organization* 60:883–909.

Lee, Keun and Xuehua Jin. 2009. "The origins of business groups in China: An empirical testing of the three paths and the three theories." *Business History* 51:77–99.

Leicht, Elizabeth A and Mark EJ Newman. 2008. "Community structure in directed networks." *Physical review letters* 100:118703.

Leuz, Christian and Felix Oberholzer-Gee. 2006. "Political Relationships, Global Financing, and Corporate Transparency: Evidence from Indonesia." *Journal of Financial Economics* 81:411–439.

Levine, Ross and Sara Zervos. 1998. "Stock markets, banks, and Economic growth." *American Economic Review* pp. 537–558.

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.

Li, Hongbin, Lingsheng Meng, Qian Wang, and Li-An Zhou. 2008. "Political connections, Financing and firm performance: Evidence from Chinese private firms." *Journal of development Economics* 87:283–299.

Liang, Qi, Pisun Xu, and Pornsit Jiraporn. 2013. "Board characteristics and Chinese bank performance." *Journal of Banking and Finance* 37:2953–2968.

Lin, Xiaochi and Yi Zhang. 2009. "Bank ownership reform and bank performance in China." *Journal of Banking and Finance* 33:20–29.

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.

Liu, Bin, Charles Cullinan, Junrui Zhang, and Fangjun Wang. 2016. "Loan guarantees and the cost of debt: evidence from China." *Applied Economics* 48:3626–3643.

Liu, Qiao and Zhou Joe Lu. 2007. "Corporate governance and earnings management in the Chinese listed companies: A tunneling perspective." *Journal of Corporate Finance* 13:881–906.

Liu, Qigui and Gary Tian. 2012. "Controlling shareholder, expropriations and firm's leverage decision: Evidence from Chinese Non-tradable share reform." *Journal of Corporate Finance* 18:782–803.

Luo, Xiaowei and Chi-Nien Chung. 2005. "Keeping It All in the Family: The Role of Particularistic Relationships in Business Group Performance During Institutional Transition." *Administrative Science Quarterly* 50:404–439.

Lynch, Scott M and Bruce Western. 2004. "Bayesian Posterior Predictive Checks for Complex Models." *Sociological Methods and Research* 32:301–335.

Masulis, Ronald W, Peter Kien Pham, and Jason Zein. 2011. "Family business

groups around the world: Financing advantages, control motivations, and organizational choices." *Review of Financial Studies* 24:3556–3600.

McPherson, Miller, Lynn Smith-Lovin, and James M Cook. 2001. "Birds of a feather: Homophily in social networks." *Annual review of sociology* pp. 415–444.

Mizruchi, Mark and Linda Brewster Stearns. 2006. "The conditional nature of embeddedness: a study of borrowing by large US firms, 19731994." *American Sociological Review* 71:310–333.

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 and Bernard Yeung. 2004. "Family control and the rent-seeking society." *Entrepreneurship Theory and Practice* 28:391–409.

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.

Naceur, Samy Ben, Samir Ghazouani, and Mohammed Omran. 2008. "Does stock market liberalization spur Financial and Economic development in the MENA region?" *Journal of Comparative Economics* 36:673–693.

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 Opper. 2010. "Political capital in a market Economy." *Social Forces* 88:2105–2132.

Nee, Victor and Sonja Opper. 2012. *Capitalism from Below: Markets and Institutional Change in China*. Harvard University Press.

Nee, Victor, Sonja Opper, and Sonia Wong. 2007. "Developmental state and corporate governance in China." *Management and Organization Review* 3:19–53.

Newman, Mark EJ. 2006. "Modularity and Community Structure in Networks." *Proceedings of the National Academy of Sciences* 103:8577–8582.

Omran, Mohammed. 2009. "Post-privatization corporate governance and firm performance: The role of private ownership concentration, identity and board composition." *Journal of comparative Economics* 37:658–673.

Orrú, Marco, Gary G Hamilton, and Mariko Suzuki. 1989. "Patterns of interfirm 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.

Pindyck, Robert S and Julio J Rotemberg. 1993. "The Comovement of Stock Prices." *The Quarterly Journal of Economics* 108:1073–1104.

Piotroski, Joseph D and Darren T Roulstone. 2004. "The Influence of Analysts, Institutional Investors, and Insiders on the Incorporation of Market, Industry, and Firm-Specific Information into Stock Prices." *The Accounting Review* 79:1119–1151.

Porta, Rafael, Florencio Lopez-de Silanes, Andrei Shleifer, and Robert Vishny. 2002. "Investor Protection and Corporate Valuation." *The Journal of Finance* 57:1147–1170.

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.

Rangan, Subramanian. 2000. "The Problem of Search and Deliberation in Economic Action: When Social Networks Really Matter." *The Academy of Management Review* 25:813–828.

Roll, Richard. 1988. "R^2." *The Journal of Finance* 43:541–566.

Saadi, Mohamed Saïd. 1989. *Les groupes Financiers au Maroc*. Rabat, Morocco: Ed. Okad.

Shleifer, Andrei and Robert W Vishny. 1997a. "The Limits of Srbitrage." *The Journal of Finance* 52:35–55.

Shleifer, Andrei and Robert W Vishny. 1997b. "A survey of corporate governance." *The journal of Finance* 52:737–783.

Siemon, Michael. 2017. "Identifying and Classifying Business Groups in 12 Arab Countries." Chapter 1.

Sleptsov, Alexander, Jaideep Anand, and Gurneeta Vasudeva. 2013. "Relational configurations with information intermediaries: The effect of firm-investment bank ties on expected acquisition performance." *Strategic management journal* 34:957–977.

Steers, Richard M, Yu-gŭn 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.

Tian, Lihui. 2005. "Bank lending, corporate governance, and government ownership in China." .

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.

Uzzi, Brian. 1999. "Embeddedness in the making of Financial capital: How social relations and networks benefit firms seeking Financing." *American sociological review* pp. 481–505.

Uzzi, Brian and James J Gillespie. 2002. "Knowledge spillover in corporate Financing networks: Embeddedness and the firm's debt performance." *Strategic Management Journal* 23:595–618.

Uzzi, Brian and Ryon Lancaster. 2003. "Relational embeddedness and learning: The case of bank loan managers and their clients." *Management science* 49:383–399.

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.

Wang, Lihong. 2015. "Protection or expropriation: Politically connected independent directors in China." *Journal of Banking and Finance* 55:92–106.

Wang, Qian, Tak-Jun Wong, and Lijun Xia. 2008. "State ownership, the institutional environment, and auditor choice: Evidence from China." *Journal of accounting and Economics* 46:112–134.

Weber, Klaus, Gerald F Davis, and Michael Lounsbury. 2009. "Policy as myth and ceremony? The global spread of stock exchanges, 1980–2005." *Academy of Management Journal* 52:1319–1347.

Wong, Raymond MK, Jeong-Bon Kim, and Agnes WY Lo. 2015. "Are Related-Party Sales Value-Adding or Value-Destroying? Evidence from China." *Journal of International Financial Management and Accounting* 26:1–38.

Wurgler, Jeffrey. 2000. "Financial Markets and the Allocation of Capital." *Journal of Financial Economics* 58:187–214.

Yiu, Daphne, Garry D Bruton, and Yuan Lu. 2005. "Understanding Business Group Performance in an Emerging Economy: Acquiring Resources and Capabilities in Order to Prosper." *Journal of Management Studies* 42:183–206.

Young, Michael N, Mike W Peng, David Ahlstrom, Garry D Bruton, and Yi Jiang. 2008. "Corporate Governance in Emerging Economies: A Review of the Principal–Principal Perspective." *Journal of Management Studies* 45:196–220.

Zhang, Dayong, Jing Cai, David G Dickinson, and Ali M Kutan. 2016. "Non-performing loans, moral hazard and regulation of the Chinese commercial banking system." *Journal of Banking and Finance* 63:48–60.

Zhang, Min, Shenghao Gao, Xinjiao Guan, and Fuxiu Jiang. 2014. "Controlling Shareholder-Manager Collusion and Tunneling: Evidence from C hina." *Corporate Governance: An International Review* 22:440–459.

Zhou, Xueguang, Qiang Li, Wei Zhao, and He Cai. 2003. "Embeddedness and contractual relationships in China's transitional Economy." *American Sociological Review* pp. 75–102.

Zhu, Yanjian and Xiaoneng Zhu. 2012. "Impact of the share structure reform on the role of operating related party transactions in China." *Emerging Markets Finance and Trade* 48:73–94.

Zijlstra, Bonne JH, Marijtje AJ Van Duijn, and Tom AB Snijders. 2006. "The Multilevel p2 Model." *Methodology* 2:42–47.