

17-03 | Sept. 29, 2017

The Complexity of Bank Holding Companies: A New Measurement Approach

Mark D. Flood

Office of Financial Research

mark.flood@ofr.treasury.gov

Dror Y. Kenett

Johns Hopkins University

dkenett1@jhu.edu

Robin L. Lumsdaine

American University, Office of
Financial Research, and National
Bureau of Economic Research

robin.lumsdaine@american.edu

Jonathan K. Simon

University of Iowa

jonathan-simon@uiowa.edu

The Office of Financial Research (OFR) Working Paper Series allows members of the OFR staff and their coauthors to disseminate preliminary research findings in a format intended to generate discussion and critical comments. Papers in the OFR Working Paper Series are works in progress and subject to revision.

Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the OFR or U.S. Department of the Treasury. Comments and suggestions for improvements are welcome and should be directed to the authors. OFR working papers may be quoted without additional permission.

The Complexity of Bank Holding Companies

A New Measurement Approach

Mark D. Flood[◆]
mark.flood@ofr.treasury.gov

Dror Y. Kenett[▶]
dkenett1@jhu.edu

Robin L. Lumsdaine^{▲◆}
robin.lumsdaine@american.edu

Jonathan K. Simon[▼]
jonathan-simon@uiowa.edu

September 29, 2017

- ◆ – Office of Financial Research (OFR), Washington DC
- ▲ – Kogod School of Business, American University; National Bureau of Economic Research (NBER); Center for Financial Stability (CFS)
- ▶ – Johns Hopkins University
- ▼ – Department of Mathematics, University of Iowa

The authors thank Hashim Hamandi, Stathis Tompaidis, Julie Vorman, Marco van der Leij, and participants in the 2017 Vienna–Copenhagen Conference on Financial Econometrics, the 2017 RiskLab/BoF/ESRB Conference on Systemic Risk Analytics in Helsinki, the 10th Annual Society for Financial Econometrics Conference in New York, the 2017 NBER/NSF Time Series Conference in Chicago, the 2017 Heterogeneous Agents and Agent-based Modeling Conference at the Office of Financial Research in Washington, D.C., the Second Conference on Network Models and Stress Testing for Financial Stability at the Banco de Mexico, and research seminars hosted by the Office of Financial Research, the University of Ljubljana, and the Center for Nonlinear Dynamics in Economics and Finance (CeNDEF) at the University of Amsterdam for helpful discussions and suggestions. Any remaining errors are the responsibility of the authors alone.

Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the Office of Financial Research or the U.S. Department of the Treasury.

The Complexity of Bank Holding Companies

A New Measurement Approach

Abstract

Large bank holding companies (BHCs) are structured into intricate ownership hierarchies involving hundreds or even thousands of legal entities. Each subsidiary in these hierarchies has its own legal form, assets, liabilities, managerial goals, and supervisory authorities. In the event of BHC default or insolvency, regulators may need to resolve the BHC and its constituent entities. Each entity individually will require some mix of cash infusion, outside purchase, consolidation with other subsidiaries, legal guarantees, and outright dissolution. The subsidiaries are not resolved in isolation, of course, but in the context of resolving the consolidated BHC at the top of the hierarchy. The number, diversity, and distribution of subsidiaries within the hierarchy can therefore significantly ease or complicate the resolution process. We use graph theory to develop a set of related metrics intended to assess the complexity BHC ownership. These proposed metrics focus on the graph quotient relative to certain well identified partitions on the set of subsidiaries, such as charter type and regulatory jurisdiction. The intended measures are mathematically grounded, intuitively sensible, and easy to implement. We illustrate the process with a case study of one large U.S. BHC.

Keywords: Bank holding company, orderly resolution, complexity, graph quotient

JEL codes: G21, G28, G33, C81

Contents

1	Introduction	2
2	Literature Review	6
2.1	Resolution plans and living wills	6
2.2	Characterizing firms' organizational structure	11
3	Network Structure and Graph Quotients	15
3.1	Basic definitions	17
3.2	Graph quotients and cycles	18
3.3	Edge contraction and label dispersion	26
4	Empirical Application of the Method	32
4.1	Case Study: the Wells Fargo BHC	33
4.2	Before and after the Wachovia acquisition	40
5	Conclusions	48

1 Introduction

In the wake of the Great Depression and the failure of more than 9,000 banks, the Banking Act of 1933 created the Federal Deposit Insurance Corporation (FDIC).¹ Since then, the FDIC has acted as receiver for several thousand failed banks, including 520 since 2008. As Figure 1 indicates, these failures tend to come in waves defined by broader episodes of financial instability.² This means that resolution is not a routine activity that proceeds at a predictable pace, but more typically an urgent and extraordinary supervisory task undertaken in exigent circumstances.




Figure 1: Number of failed U.S. banks, 1935–2016. Three (post-FDIC) failure waves: Great Depression (1935-43), Savings and Loan Crisis (1980-94) and Global Financial Crisis (2009-15).

Source: FDIC

Historically, most of these failed banks were relatively small community banks whose resolution posed little risk to the financial system. Notably, all were considerably smaller and less complex than the most systemically important firms active today. Until the 2008 financial crisis, the largest bank failure in U.S. history was that of Continental Illinois National Bank and Trust Company. At the time of its failure in 1984, it was the seventh largest bank in the U.S., with approximately \$40 billion in assets (approximately \$96 billion in 2016 dollars).³ By comparison, Bank of

¹See Flood (1992). The Banking Act of 1933 is also known as the Glass-Steagall Act.

²Data in Figure 1 are from FDIC (2016b); the majority of Depression-era bank failures preceded start of the FDIC's sample in 1934. For the more recent period, the FDIC (2016a) provides an up-to-date list of bank failures since 2000.

³Haltom (2013) describes the failure of Continental Illinois.

New York Mellon Corp., the seventh largest bank in the U.S. today, has nearly \$300 billion in consolidated assets and the largest bank, JPMorgan Chase & Co., has over \$2 trillion in consolidated assets. Rapid resolution of the largest institutions, if they should fail, is critical to ensure the provision of liquidity to counterparties that are relying on incoming payments from the failing firm and to avoid a cascade of events that can quickly snowball into a systemic panic.

The danger of a disorderly resolution that escalates into a crisis is a channel of systemic risk that is unlikely to correlate with other common measures of systemic risk, especially those that are market-based. This is because such an event involves the intersection of two (ex ante) low-probability events: a large-firm failure and a mismanaged resolution implementation. This does not imply, however, that resolution complexity cannot be measured, or that authorities cannot plan for orderly resolution.

In the aftermath of these challenges, Title II of the Dodd-Frank Consumer Protection and Wall Street Reform Act of 2010 expanded the FDIC's receivership authority to large, complex financial companies, including bank holding companies (BHCs) with \$50 billion or more in assets and nonbanks whose failure regulators have determined could pose a threat to U.S. financial stability.⁴ In addition, to assist with orderly resolution should it become necessary, a regulated firm is now required to submit a plan for supervisory approval, known as a "living will," describing how it should be resolved. As part of this requirement, the Board of Governors of the Federal Reserve System (BoG) and the FDIC (BoG-FDIC) have issued a series of documents that provide guidance regarding the submission and assessment of firms' resolution plans; see BoG-FDIC (2016).⁵ These submissions are then assessed by the agencies,

⁴FDIC (2011) envisions a counterfactual case study of how the Lehman resolution might have proceeded under the new Dodd-Frank authority. For insurance companies, resolution occurs under applicable state law and is conducted by the state regulator; only in the case that the state regulator does not act within 60 days does the FDIC have resolution authority. For a summary of the expanded receivership authority, see FDIC (2010).

⁵Eight domestic bank holding companies participated in the initial wave of submissions. They were: Bank of America Corporation; Bank of New York Mellon Corporation, PLC; Citigroup, Inc.;

which provide feedback identifying deficiencies and required remediation. One of the topics these submission documents are expected to address is the issue of the BHC's complexity. This requirement is in line with the Financial Stability Board's definition of systemically important financial institutions (SIFI) as "financial institutions whose distress or disorderly failure, because of their size, complexity and systemic interconnectedness, would cause significant disruption to the wider financial system and economic activity."

Of these factors, complexity is possibly the hardest to quantify, despite broad agreement that it contributes to the systemic risk posed by a financial institution. In particular, it can be difficult to quantify complexity independently from the size of the institution (one possible approach is proposed by Lumsdaine et al., 2015). Despite the measurement challenge, complexity is a key policy concern. For example, Treasury Secretary Steven Mnuchin stated in his confirmation testimony that banks should be regulated based on "complexity and activity, not just size;" see Tracy and Rubin (2017). In sum, understanding the level of complexity of a given financial institution — separate from its total assets or number of subsidiaries — is crucial for both assessing systemic risk and planning for an orderly resolution in case of a failure event.

The Resolution Plan Assessment Framework, BoG-FDIC (2016), emphasizes a company-specific process that, among other objectives, requires firms to "rationalize their structures." In their April 2016 determinations, the agencies cite a lack of demonstrable progress in remedying deficiencies related to legal entity rationalization (LER) in four of the eight firms, deficiencies that previously had been highlighted in 2014. JPMorgan Chase & Co., State Street Corp., and Wells Fargo & Co., for example, were found to have "inadequate legal entity rationalization criteria," and

Goldman Sachs Group, Inc.; JPMorgan Chase & Co.; Morgan Stanley; State Street Corporation; and Wells Fargo & Company.

are required to establish criteria that “are clear and actionable and promote the best alignment of legal entities and business lines to improve the firm’s resolvability.” The BoG-FDIC guidance (2016, pp. 18–19) for the 2017 submissions notes that “LER criteria should govern the firm’s corporate structure and arrangements between legal entities in a way that facilitates the firm’s resolvability as its activities, technology, business models, or geographic footprint change over time.” It specifies that application of such criteria should “include clean lines of ownership, minimal use of multiple intermediate bank holding companies, and clean funding pathways between the parent and material operating entities,” “adequately protect the subsidiary insured depository institutions from risks arising from the activities of any nonbank subsidiaries of the firm,” and “minimize complexity that could impede an orderly resolution and minimize redundant and dormant entities.” In addition, firms are required to identify discrete operations that can be sold or transferred in resolution.

We use fundamental methods of network analysis and graph theory to propose measures that may help inform resolution plan development and assessment of a firm’s legal entity (organizational) structure and ease of resolvability. Describing the ownership structure of a BHC as a network, we are able to consider topological properties in relation to complexity that may ultimately provide inputs for orderly resolution. In this paper, we focus on the use of the topological quotient. Intuitively, the topological quotient collapses (or “glues together”) certain points of the topology that share a common property. In our case, we are fusing vertices (subsidiaries) of the BHC ownership network. Specifically, we use the topological quotient to partition the holding company graph into clusters of similar subsidiaries, based on charter type (e.g., bank, broker-dealer, insurer, depository, etc.), regulatory or geographical jurisdiction, or activity. The cycle rank of the quotient graph (defined in Section 3.2) reveals how many links of ownership and control must be broken to extract all subsidiaries of a given type from the holding company. Because each such link

(ignoring self-loops) represents a connection between two subsidiaries of different entity types or geographic jurisdictions, it highlights a potential coordination problem during resolution. Therefore, we argue that the number of such links should correlate positively with the complexity of resolution. This is an empirical question, which we do not resolve conclusively here, but which remains an area of ongoing research. The topological quotient also may be useful in assessing ways to wind-down some portions of a firm (e.g., a certain business line) while still ensuring that the firm be no more complex than it is currently.

The rest of this paper is organized as follows: Section 2 briefly reviews the literature linking organizational structure and network analysis. Section 3 describes the network structure and the topological (graph) quotient. Section 4 demonstrates the feasibility of our approach by analyzing the Wells Fargo holding company's legal entity structure as a case study; we consider Wells Fargo's complexity both before and after its 2009 acquisition of Wachovia Corp. Section 5 concludes.

2 Literature Review

In this section, we discuss two main strands of literature relevant to this paper: (1) bank resolution and living wills, and (2) organizational structures and their complexity.

2.1 Resolution plans and living wills

The near-collapse of Bear Stearns Companies, Inc. in March 2008 without an effective resolution mechanism resulted in a hasty acquisition by JPMorgan Chase. Yet as 2008 unfolded, it became apparent that the financial stability risks associated with the potential failure of a large financial institution are neither an idle concern nor a fully-solved problem (Herring, 2014). The sequence of events in September of that

year highlighted the challenges associated with resolution. Both Fannie Mae and Freddie Mac were quickly placed into conservatorship as subprime mortgage defaults escalated. Emergency weekend negotiations to save Lehman Brothers Holdings Inc. terminated unsuccessfully; the firm filed for bankruptcy in the early morning hours of Monday, September 15 and collapsed without an effective resolution mechanism. The next day (Tuesday, September 16), American International Group, Inc. (AIG), which had sold credit protection on large swaths of mortgage securitizations, was on the brink of collapse and was rescued via an \$85 billion cash injection from the Federal Reserve; see FCIC (2011), Valukas (2010). The problem quickly spiraled into a systemic crisis. At the time, none of these firms' parent companies was regulated by the U.S. banking authorities, although some had banking subsidiaries that were, and none of the agencies had resolution authority over them. To avoid complete collapse of the financial system, regulators were left scrambling to untangle the complex web of these firms' exposures and to identify potential suitors in a very short time-frame. Particularly in the case of Lehman Brothers, where a suitor could not be identified in time, the lack of an effective resolution mechanism threatened global financial stability; see Fleming and Sarkar (2014), Carmassi and Herring (2013).

In a recent OFR Research Brief, Bright et al. (2016) study the public portions of the 2014 and 2015 living wills of the eight global systemically important banks (G-SIBs) that participated in the initial wave of resolution plan submissions. They argue that organizational complexity is a big issue for the largest banks. The simpler the structure of a G-SIB, the easier it is to resolve. They find the public portions of living wills include relatively little data, but based on the data available, U.S. G-SIBs have done little to streamline their core businesses since 2013. They conclude that the public information is insufficiently detailed to determine if a failing bank could be resolved without government intervention. Specifically, the living wills' public sections offer only a rough idea of how these banks would manage this complexity in a

failure. They compare information about those structures in the living wills with the information in another regulatory database, the Federal Financial Institution Examination Council's (FFIEC) National Information Center (NIC) database and conclude that data on corporate structures could be more transparent. Finally, Bright et al. (2016) emphasize that cross-border operations present challenges for resolution and that more public information on G-SIBs' cross-border operations would be helpful. For example, a G-SIB's home country and host country may have competing interests. Host countries may protect local interests by shielding local subsidiaries from foreign claims, thus making resolution harder.

Although Bright et al. (2016) argue that complexity is an important dimension for evaluating a living will, they largely treat complexity as a counting exercise — for example, the number of core business lines, critical operations, or material legal entities. Furthermore, they emphasize that assessment by category (for example, clearing) does not necessarily highlight whether a firm has multiple critical operations within the same category. Therefore, simple categorical counts will likely understate true complexity and risk. To address this issue, new metrics are needed. To quantify the complexity of the organization as a whole, these metrics must go beyond simple measures of size, to cover the details of individual subsidiaries and their relationships within the overall organization, ideally in a coherent conceptual framework.

In an attempt to address the resolution of G-SIBs, the Financial Stability Board (FSB, 2013) identified two stylized strategies: Single Point of Entry (SPE), and Multiple Point of Entry (MPE). The SPE strategy applies resolution powers at the top holding or parent company level by a single resolution authority, mainly in the jurisdiction responsible for the global consolidated supervision of a group (e.g., a holding company). The assets and operations of particular subsidiaries are preserved on a going-concern basis, avoiding the need to apply resolution at a lower level within the group. The MPE strategy applies resolution powers by two or more resolution au-

thorities to multiple parts of the group (ideally simultaneously), including strategies in which a group is broken up into two or more parts. The FSB does not argue in favor of one strategy over the other, rather it states that a choice depends on the structure of the organization and its operations. It further proposes to combine the two approaches. For example, some MPE strategies may involve applying multiple SPE resolutions to different parts of the firm, e.g., regional blocs that are separable from one another. The FSB states (2013, p. 13):

The choice of the strategy needs to take account of the existing structure and business model of the individual firm and the firm's particular characteristics. For example, an SPE strategy may represent the most effective option if the debt issued at the top of the group is sufficient to absorb the group's losses and ensure the viability of the operating subsidiaries. It may be more suitable to a firm that operates in a highly integrated manner (through, for example, centralized liquidity, trading, hedging and risk management). An MPE strategy may be suitable for firms with a decentralized structure and greater financial, legal and operational separation along national or regional lines, with sub-groups of relatively independent, capitalized and separately funded subsidiaries, particularly if different parts of the group can continue on a standalone basis. In some cases, a group may need to be restructured to make it amenable to one or other of an SPE or MPE resolution strategy.

Thus, the FSB clearly emphasizes the role the organizational structure plays in the ability to implement resolution plans.

In the U.S., the FDIC (2013) has proposed the SPE strategy. While six of the eight G-SIBs that submitted living wills described their living will as an SPE strategy, the Bank of New York Mellon and Wells Fargo plans did not. In Bank of New York Mellon's plan, its depository institution and its affiliates would be resolved through FDIC receivership, while the company's broker-dealer, asset manager, and other entities would be sold or would file for bankruptcy. Wells Fargo's plan calls for the creation of a bridge bank — a temporary national-chartered bank organized by regulators — to operate its core banking business. The parent holding company

would undergo a bankruptcy reorganization, and other units would be sold (Bright et al., 2016). In April 2016, the FDIC and Federal Reserve publicly announced their views and provided feedback on the 2015 plans submitted by eight large U.S. G-SIBs; see BoG-FDIC (2016). Seven of the banks' plans were deemed "not credible" by at least one of the two regulators (Bright et al., 2016). By December 2016, the Wells Fargo living will was the only resubmitted plan that was rejected by the Federal Reserve and the FDIC. As a result of this failure, the regulators prohibited Wells Fargo from establishing new international units or acquiring a subsidiary that is not a bank (Corkery, 2016). In April of 2017, the FDIC and the Federal Reserve Board announced that Wells Fargo had adequately remediated the deficiencies in its 2015 resolution plan, and consequently it will no longer be subject to growth restrictions (Federal Deposit Insurance Corporation and Board of Governors of the Federal Reserve System (FDIC-BoG), 2017).

The SPE and MPE approaches link resolution plans with the organizational structure of the firm. The actions taken against Wells Fargo by the FDIC and the Fed underscore this point. However, the link between resolution strategy and organizational structure is not a simple matter of the size of the firm. Rather, the organizational structure — in particular, its complexity — plays a crucial role in the formation of living wills and resolution plans and strategies. It is plausible that *holding company* complexity is an endogenous byproduct of a repeated game between firms and regulators. More general notions of complexity, however measured, might be a side effect (intentional or unintentional) of the ongoing firm-regulator interplay. In a crisis, the simple existence of complexity will likely matter more than its source.⁶

⁶(Kaufman, 2014) notes that "too complex to fail" is one of many possible interpretations for the more common phrase, "too big to fail" (TBTF). Carmassi and Herring (2014) similarly highlight the interplay between the notions of TBTF, interconnectedness, complexity, and difficulty of resolution. The often casual use of these various terms in the popular press can generate ambiguity and inappropriate conflation of the meaning of the term "complexity," which we hope to clarify by presenting precise, formal, mathematically grounded definitions.

2.2 Characterizing firms' organizational structure

One way of characterizing the organizational structure of a firm is via its control hierarchy (Vitali et al., 2011). The control hierarchy consists of the (parent) company and all of its subsidiaries, considered in a natural hierarchical and networked arrangement. This is effectively a standard representation of the intraconnectedness of a firm, along the lines of Coase (1937) who described a firm as a “system of relationships.” Graph theory provides a natural conceptual framework for analyzing these hierarchies. Here, we represent the subsidiaries of a BHC as vertices of a graph, and the parent-child ownership/control relationships as edges connecting one subsidiary to the next. We apply established mathematical tools — notably, graph quotients and cycle ranks — to define and measure BHC complexity. We focus on the intricacy with which regulatory jurisdictions (defined either by geography or charter type) are intermingled within the BHC, creating possible coordination problems in the event that the firm must be resolved. These coordination problems emerge naturally as cycles in appropriately constructed quotients calculated on BHC control hierarchies.

Starting with a simple case, if a threshold over 50 percent defines “control,” the network structure will necessarily be a *rooted directed tree* — a system of vertices and edges without loops — because every vertex can have at most one parent above this threshold.⁷ In such tree structures, previous literature has found that organizations tend to be siloed, with more communication occurring along connected lines of the rooted directed tree than along unconnected lines (Kleinbaum et al., 2013). Because organizational structure affects communication, Lumsdaine et al. (2015) posit that similar silos exist for the communication of risk. Specifically, to the extent that firms with complex organizational structures are more likely to suffer communication lapses, either within the firm or among its regulators, risk that in principle would be

⁷Our case study in Section 4 shows that without such a threshold, a large BHC can indeed have ownership loops among its many subsidiaries.

contained within the firm may result in more widespread systemic risk the longer it goes undetected. They use the innate network structure of the control hierarchy to develop metrics that characterize the organizational level of complexity. In doing so, they therefore highlight the importance of considering intra-firm complexity in addition to the more commonly studied inter-firm complexity (i.e., the interconnectedness across firms).

Building on this basic intuition, we propose to use the topological properties of the BHC ownership structures as possible metrics for quantifying BHC complexity. Some basic topological properties include number of vertices (subsidiaries), number of edges (ownership relationships), and number of levels from the high holder in the BHC (depth) — amounts that tend to increase with the size of the BHC. For this reason, in general, we find that the simplest topological characteristics of BHCs are not empirically useful in assessing complexity as distinct from size. Simple vertex and edge counts may not correlate with complexity, and the number of layers from the ultimate parent BHC to the most deeply nested subsidiary tends to be small and does not exhibit much variation. However, there are other topological characteristics of the BHC hierarchy, beyond those most immediately apparent, that we believe may capture certain coordination problems that could contribute to the complexity of resolution.

In general, greater complexity (in terms of organizational structure and business activities) of an individual firm makes it harder for a supervisory entity to disentangle and understand the firm's larger systemic interconnectedness and increases the likelihood that some parts of the firm's activities and interrelationships go unnoticed. It also makes it harder for a resolution authority to disentangle and resolve an institution in times of financial distress. In the case of large multinational organizations, quantitative measures related to resolution complexity should naturally account for the burdens posed by necessary coordination across multiple national and regula-

tory environments.⁸ In addition, the identification of metrics that enable comparison across firms with very different organizational structures is of critical importance in assessing the adequacy of submitted resolution plans.

To date, most measures of complexity have been little more than specific dimensions of the size of a firm. For example, the Basel Committee on Bank Supervision highlights three aspects of a firm’s balance sheet for determining complexity, each with a $1/15$ weight (i.e., $6\frac{2}{3}$ percent): (1) notional amount of OTC derivatives, (2) trading and available-for-sale securities, (3) level 3 assets.⁹ Other measures have been proposed, such as the number of subsidiaries (Carmassi and Herring, 2013), the number of operating segments (Grant et al., 2000), or concentration of activities (Cetorelli and Goldberg, 2014) but all are positively related to size.

For this reason, Lumsdaine et al. (2015) proposed a new metric of complexity that reflects the challenges in supervision from the need to coordinate oversight efforts across a variety of jurisdictions, industries and agencies. In their framework, the working hypothesis is that a “perfect” supervisory tree represents the simplest supervisory structure, one in which each supervisor’s entities are purely domestic (country-perfect tree) or within a single industry classification, also known as SIC (SIC-perfect tree). Of course such a structure does not exist in reality. The conceptual intuition, however, is that less breadth of expertise is required to supervise/evaluate an entity that is closer to a perfect supervisory tree, across different national jurisdictions or industries, than in a firm that has a tree structure that is farther from perfect.¹⁰ In addition, the closer an entity’s organizational structure is to a perfect tree, the less

⁸Carmassi and Herring (2013) note that the FSB identified the need to cooperate as one of the five remaining challenges that regulators faced in order for resolution to proceed in an orderly manner.

⁹The Financial Account Standards (FASB) Statement 157 and International Financial Reporting Standard (IFRS) 13 “Level 3” as the least liquid category of assets, for which market price of comparable asset are unavailable. Fair valuation of level 3 assets is often based on a theoretical model.

¹⁰Such an idea is behind the concept of “ring-fencing,” where each national authority has responsibility for the resolution of banks that fall under its jurisdiction. This type of resolution framework is referred to as the “subsidiarization” model in Carmassi and Herring (2013).

likely that difficulty in one country or industry will spill over into other countries or industries in which the firm operates. By comparing the firm's actual organizational structure to the perfect tree ideal, inference can be drawn about the extent of contagion that a firm would experience if one of its subsidiaries deteriorated.

A similar intuition can be applied to resolution. An entity with many subsidiaries that all fall under the same supervisor might be easier to resolve than an entity with fewer subsidiaries that are spread across different country or SIC classifications and therefore require coordination across a number of supervisors. One way these differences can be assessed is via the computation of a firm's *degree distribution*. This fundamental network descriptor describes the histogram of the network's set of degrees, given by the function $d(i)$ that records the fraction of vertices with i children. Just as the degree distribution of a tree describing a firm's reporting lines can characterize the spans of control of its management (Urwick, 1956), the degree distribution of a firm's control hierarchy analogously can describe a supervisor's span of control in resolving a firm. It is not just the number of children emanating from a vertex that is important for resolution; rather, it is the ease by which a resolution authority can access information in order to make its recommendation. The supervisory coordination challenges we envision are likely what motivated the requirement that a firm's resolution plan identify the "supervisory authorities and regulators, including information identifying any foreign agency or authority with significant supervisory authority over material foreign-based subsidiaries or operations;" BoG-FDIC (2011, p. 67330).

Our approach represents a novel extension of the rapidly expanding literature on financial networks.¹¹ Network theory is a natural tool for understanding complex systems, such as the robustness of the system to specific scenarios or the impact of


¹¹This literature is extensive. For an introduction, see (Battiston et al., 2016; Chan-Lau et al., 2009; Cont et al., 2013; Glasserman and Young, 2016; Haldane, 2009; May, 2013; Squartini et al., 2013; Summer, 2013).

policy on system actions. One advantage of the network approach is a clear focus on the details of interactions in the system, rather than the aggregated behavior of the system as a whole. Such details — the specifics of chains of ownership and corporate control within a large financial firm — are central to our analysis of holding company complexity. To our knowledge, few if any papers have proposed network theory in the context of BHC resolution.


3 Network Structure and Graph Quotients

Many aspects of the resolution process are unpredictable and unavoidable. However, one key challenge to resolving a large financial firm may be measurable in advance, using information regulators already possess. This is the intricacy with which similar businesses are distributed within the holding company. The key to applying graph quotients to the problem of resolving a failed bank is to choose a suitable set of categories or descriptors for the subsidiaries in the BHC (i.e., a set of labels corresponding to the vertices in a graph of the BHC’s organizational structure). The examples below categorize subsidiaries according to either their entity type or geographic jurisdiction (U.S. state or foreign country), as identified in the NIC database, but other categorizations are possible. The basic intuition is that parent-subsidiary communication — for both managers and resolution authorities — between two entities of the same type is easier than communication across types.

The quotients of the holding company graph yield straightforward algebraic measures of the number of legal relationships that must be severed to segregate the holding company into distinct components, each consisting only of legal entities of one type. It is possible that the number of such cuts may tell us something important about the complexity of the resolution process for a given firm. We start with a simple motivating example.



(a) Arbitrary arrangement of subsidiaries



(b) Subsidiaries organized by entity type

Figure 2: Two possible configurations of holding company subsidiaries of four entity types (FHD=financial holding company–domestic; BHC=bank holding company; NAT=national bank; and DEO=domestic entity other): (a) arranged arbitrarily, versus (b) organized by charter type.

Source: Authors' analysis

Figure 2 illustrates the concept by comparing two alternative configurations of subsidiaries within a hypothetical BHC. Large financial holding companies subsume a diverse set of subsidiary types; in general, each type can have its own regulator(s). In Panel (a), the four subsidiary types are distributed arbitrarily within the overall holding company, perhaps reflecting the accumulated history of bank mergers or forays into new business lines or geographic areas. DEO is a miscellaneous category that covers a wide range of business lines and strategic goals.¹² Panel (b) shows the same set of subsidiaries, rearranged into clusters of similar firms. In principle, shareholders of the going concern should be indifferent to the two arrangements, because performance is consolidated to the top-level holding company. Regulators resolving

¹²The NIC database has 43 possible values for the “Entity Type” field; BoG (2016).

the failed firm, on the other hand, must quickly disentangle the subsidiaries, some of which will be sold to other firms, others recapitalized, and some wound down. We suggest that this process should be operationally easier when the organizational structure of the holding company aligns cleanly with the boundaries of regulatory jurisdiction. For example, in Panel (a), six cuts are required to separate all the DEO vertices from the holding company; in Panel (b), only one cut is necessary. (On the other hand, for national bank (NAT) vertices, the reduction is only from three to two.) From a resolution perspective, configuration Panel (b) is likely to be easier to resolve than Panel (a).

The remainder of this section presents intended measures of resolution complexity that potentially quantify the extent to which a BHC hierarchy departs from the “easy” structure exemplified by Panel (b) of Figure 2.

3.1 Basic definitions

To explore the properties of quotient graphs and their relationship to the original graph, we start with some basic notation. For any finite graph G , let $v(G)$ denote the number of vertices in G ; let $e(G)$ denote the number of edges; and let $b_0(G)$ denote the number of connected components.¹³ We are seldom interested in $b_0(G)$ directly, but it is a key building block for many results. Although it would be anomalous for a BHC organizational structure to have isolated (unattached) blocs, these may arise in the course of a resolution, or as an artifact of our network analysis; for example, we might choose to ignore (prune) edges with an ownership percentage below some threshold.

Assume now that the vertices of the graph have all been labeled according to

¹³The notation b_0 is standard in the topology literature, starting a sequence b_0, b_1, \dots of *Betti numbers* that describe basic properties of a topological space. The term is an homage to the 19th-century mathematician Enrico Betti. For simple graphs, such as BHC hierarchies, all Betti numbers above b_1 are zero.

some classification scheme, for example by NIC entity type. If an edge of G connects two vertices carrying the same label, call that edge *homogeneous*; if the vertex labels are different, call the edge *heterogeneous*. Call a subgraph H of G *homogeneous* if all of its vertices carry the same label (which implies that all the edges of H are homogeneous). A homogeneous subgraph H is a portion of the holding company hierarchy comprising only of subsidiaries of the same type.

A *maximal homogeneous connected subgraph of G* is a connected homogeneous subgraph H such that H cannot be enlarged without losing either homogeneity or connectedness.¹⁴ The graph G is the union of its maximal homogeneous connected subgraphs, with various of these special homogeneous subgraphs connected to one another by heterogeneous edges. Given a BHC with a fixed set of subsidiaries, the number of these homogenous patches that must be stitched together to compose the BHC is an indication of its complexity with respect to the categories used to define the labels.

3.2 Graph quotients and cycles

In thinking about the ease of resolution of a firm, we will consider two graph theory constructions: a *quotient* and a *contraction*. These are two different techniques for dimensionality reduction of a large graph to a more manageable size (or for proving theorems inductively), but doing so in a way that preserves certain key topological properties of interest. The general idea of a *quotient space* involves:

- partitioning the points of some set into equivalence classes and constructing a new set whose “points” represent those equivalence classes, and

¹⁴A maximal homogeneous connected subgraph can also be defined as the union of all connected homogeneous subgraphs of G containing a given vertex, a . It is the largest subgraph of G in which each vertex is reachable from a without ever traversing a heterogeneous edge. When G is a tree, each connected subgraph of G is also a tree, and we will talk about *maximal homogeneous subtrees* of G .

- giving that new set an appropriate structure such as interconnections, algebraic operations or a sense of distance.

We start with the set consisting of a holding company and its subsidiaries, which we represent as the vertices of a graph, connected to each other by legal relationships, mostly hierarchical, of ownership and control. To calculate a quotient of this graph, we make a new graph by assigning each vertex to an equivalence class based on a classification rule, such as entity type or geographic location. These equivalence classes (for example, national bank, state member bank, etc.) are then the vertices of the quotient graph.


3.2.1 BHC quotient graphs

This section presents four ways to calculate quotients of graphs representing a BHC's organizational structure. Each of the four variations — full versus heterogeneous and condensed versus uncondensed — quotients the graph in a different way along a given dimension to assess the complexity of resolving the BHC along that dimension. We start by representing the BHC as a directed graph G , whose vertices represent the BHC's subsidiaries and whose edges represent relations of ownership and control. The entities in the BHC are labeled according to some categorization such as geographic location or NIC entity type. Partition the vertices of G into equivalence classes V_1, V_2, \dots, V_N , based on the labels: that is, each subset V_i consists of all the vertices of G that have one particular label, and every vertex of G belongs to exactly one subset, V_i . The number of equivalence classes, N , is the number of different labels available.


The *full quotient graph*, Q , is a new graph derived from G based on the partitioning of the vertices according to the elements of the equivalence class (See Figure 3). Q has one vertex for each equivalence class, V_1, V_2, \dots , defined on G . Because each equivalence class contains at least one vertex, Q cannot have more vertices than G :

$v(G) \geq v(Q)$. The edges in Q are defined as follows: for any vertices $v_1 \in V_i$ and $v_2 \in V_j$ that share an edge in G , put an edge in Q joining V_i and V_j . We can formalize this as a mapping from G to Q . That is, define a function $p : G \rightarrow Q$, called the *projection* by $p(v) = [v]$, where v is a vertex of G and $[v]$ denotes the corresponding equivalence class $V_i \in Q$ that contains vertex v . For an edge e of G , define $p(e)$ to be the edge in Q corresponding to e .


Each edge in G yields an edge in Q , so $e(G) = e(Q)$. For a given edge, $e \in G$, if $V_i = V_j$ (i.e., $i = j$), then $p(e) \in Q$ is a self-loop on V_i . The distinction between ordinary edges and self-loops in Q is important to our analysis. Each homogeneous edge in G produces a self-loop in Q and each heterogeneous edge in G produces an edge in Q that joins two different vertices in Q (i.e., an edge that connects two vertices that belong to two different equivalence classes).



(a) Original graph, G



(b) Full quotient, Q



(c) Heterogeneous quotient, $Q_{||}$

Figure 3: Examples of quotient graphs: (a) Hypothetical BHC tree, G , with labeled entities; (b) full quotient, Q ; and (c) heterogeneous (no self-loops) quotient, $Q_{||}$.

Source: Authors' analysis

Figure 3 illustrates a full quotient graph, Q , starting with a hypothetical BHC hierarchy, G , in Panel (a). Entities are labeled as operating in either category “E” (six

subsidiaries) or “A” (four). Suppose for concreteness that these labels indicate the geographic headquarters’ location of each subsidiary. These labels define two equivalence classes. Panel (b) shows the matching full quotient graph, after the vertices have been pulled together into the two equivalence classes. The four heterogeneous edges in G correspond to the four edges between vertices A and E in Q .

The heterogeneous edges highlight communication challenges faced by the firm managers or resolution authorities who must handle the firm. We can think of the quotient graph as the organizational structure that results from grouping all the entities of each category together under the auspices of the single functional regulator appropriate for that category. The benefit of the full quotient graph representation is the ability to group entities according to a specific characteristic. For example, a visualization of Q might help resolution authorities in quickly determining the ease with which some subset of entities and their associated risks can be ring-fenced.

No edges are lost in the mapping from a graph, G , to its full quotient, Q , but the number of distinct connected components can shrink. Formally:

Proposition 3.1. *If G is a graph, V_1, V_2, \dots a partition of the vertices of G , and Q is the associated full quotient graph, then*

$$e(Q) = e(G) , \text{ and} \tag{1}$$

$$b_0(Q) \leq b_0(G) \tag{2}$$

Proof. By definition, each edge (recall we are including loops) in Q is created from an edge of G , so the numbers of edges are identical.

Two vertices v, w of G are in the same connected component when there is (ignoring orientations) a sequence of edges in G connecting v and w . That is, we

have a sequence $v = v_0, v_1, \dots, v_k = w$ of vertices of G and a sequence of (virtually) oriented edges e_1 from $v = v_0$ to v_1 , e_2 from v_1 to v_2 , etc. Their projections $p(e_1), p(e_2), \dots, p(e_k)$ in Q then form a path (which is a connected subset) in Q that runs through the equivalence classes, $[v] = [v_0]$ to $[v_1]$ to $[v_2]$ etc. to $[v_k] = [w]$. So the number of connected components cannot increase when we pass from G to Q . \square

One way to think of the full quotient, Q , is as a reduction in the dimensionality of the BHC graph. Three related reductions of the full quotient emphasize different aspects of the BHC organizational structure. One important variation on the full quotient, Q is what we call the *heterogeneous quotient*, denoted $Q_{||}$. We derive $Q_{||}$ from Q by deleting any self-loops, effectively ignoring the homogeneous edges in the original graph, G . The notation emphasizes that $Q_{||}$ admits parallel edges (from heterogeneous edges in G , e.g., the four edges connecting A and E in Panels (b) and (c) of Figure 3) but not the loops (from homogeneous edges in G). Removing loops from Q facilitates visualization of the amount of regulatory *coordination* required for resolution — by focusing attention on the heterogeneous edges connecting legal entities from different categories.

It is sometimes convenient to work with a simplified quotient graph that condenses multi-edges (i.e., multiple parallel edges, including multiple self-loops on the same vertex) into a single edge between equivalence classes. We denote such a *condensed quotient graph* with a bar, e.g., \bar{Q} or $\bar{Q}_{||}$. These are the graphs obtained from Q or $Q_{||}$, respectively, by replacing any multi-edges between V_i and V_j with a single edge (where $i = j$ indicates a self-loop). For example, \bar{Q} is a concise representation of the full quotient, where one might include edge weights on \bar{Q} to record the number of parallel edges in Q . \bar{Q} visually identifies the presence of both intra- and inter-supervisory coordination required for resolution of a firm, G , with the given organizational structure. In contrast, $\bar{Q}_{||}$ identifies the presence only of needed inter-

supervisory coordination.

3.2.2 Cycles

Even if a graph G is a tree, its full quotient graph is unlikely to be a tree because Q will typically have a number of loops and longer cycles. A multi-edge cycle can be thought of as a sequence of edges that starts and ends at the same vertex, or equivalently as two different edge paths connecting the same pair of vertices. Passing from a graph to a quotient graph can both create cycles and reduce the number of connected components.

The act of identifying two vertices of any graph, G , into a single vertex in Q must either create a cycle (perhaps as a loop or parallel edges) or reduce the number of connected components. If the two vertices lie in different connected components of G , then the identification reduces the number of components but does not create cycles. If G itself has some cycles, then passing to Q will generally create more cycles. We make this precise in Theorem 3.2 below. In our holding-company context, self-loops in Q indicate subsidiaries that were contiguous with other subsidiaries of the same type in G ; other (non-loop) cycles indicate subsidiaries of a given type that were dispersed non-contiguously.

BHC graphs, for example as extracted from the NIC data, typically are connected. Any connected graph G can be viewed as a tree containing all the vertices (known as a *spanning tree*), together with some additional edges. Each new edge added to a spanning tree creates a cycle. Reversing perspective, any connected graph can be reduced to a spanning tree (covering all its vertices but containing no loops) by removing some number of edges. That number of edges that must be cut to get to a spanning tree is called the *cycle rank* of G , and is measured by the first Betti number of the graph, $b_1(G)$, as in equation (3).¹⁵

¹⁵On cycle ranks, see Harary (1969, p. 192-194). The cycle rank can also be defined in terms

Although the number of spanning trees for a given graph can be large¹⁶, the cycle rank is unique. The basic relationship between the cycle rank and the number of edges, vertices, and connected components of a graph is given by the *Euler characteristic* equation: For a graph G having $v(G)$ vertices, $e(G)$ edges, and $b_0(G)$ connected components,

$$\text{cycle rank of } G = b_1(G) = e(G) - v(G) + b_0(G) . \quad (3)$$

As is typical for a BHC, for any connected graph G , the whole graph consists of a single component and $b_0 = 1$. Combining this observation with equation (3) we note:

$$\text{If } G \text{ is any connected graph then } b_1(G) = e(G) - v(G) + 1 . \quad (4)$$

Separately, all trees are acyclic, so the cycle rank of a tree is $b_1 = 0$, so:

$$\text{If } G \text{ is a tree, then } e(G) = v(G) - 1 . \quad (5)$$

The Euler characteristic is a foundation for our proposed measures of the complexity of holding company resolution. For example, Theorem 3.2 (equation(6)) shows how the cycle rank of the full quotient, $b_1(Q)$, begins to incorporate information about the assignments of categories to the entities in G . Theorem 3.3 (equation(8)), shows how the cycle rank of the heterogeneous quotient, $b_1(Q_{||})$, reflects the dispersion vs. concentration of homogeneous subgraphs in G .

Theorem 3.2. *Let G be a connected graph with vertices partitioned into N equivalence classes of an algebraic structure called the cycle space, which is an instance of homology theory (Massey, 1991, Section VIII.3).*

¹⁶For example, the Borchardt-Cayley theorem says the number of spanning trees for a complete graph on n vertices is n^{n-2} ; see the discussion at (Stanley, 2013, p. 184).

lence classes V_1, V_2, \dots, V_N , and let Q denote the full quotient graph. Then

$$b_1(Q) = b_1(G) + v(G) - N . \quad (6)$$

Proof. From the Euler characteristic equation, and the assumption that G (hence also Q) is connected, we have

$$b_1(Q) = e(Q) - v(Q) + 1 , \text{ and}$$

$$b_1(G) = e(G) - v(G) + 1 .$$

Substitute $e(Q) = e(G) = b_1(G) + v(G) - 1$ and $v(Q) = N$ into the first equation to get

$$b_1(Q) = b_1(G) + v(G) - 1 - v(Q) + 1 = b_1(G) + v(G) - N .$$

□

Remark 1. Note that because $v(Q) = N$ and $e(Q) = e(G)$, it also follows that $b_1(Q) = e(G) - N + 1$.

Remark 2. If G is not connected, then $b_0(Q)$ depends not only on $b_0(G)$ but also on how the various vertex equivalence classes intersect the various connected components of G . The expression for $b_1(Q)$ then needs to be corrected by the difference between $b_0(G)$ and $b_0(Q)$.

Ideally, we would like to relate the cycle rank of the heterogeneous quotient, $b_1(Q_{||})$, to the category assignments, along the lines of equation (6). But for general graphs, unfortunately, when G itself has cycles, the interpretation of $b_1(Q_{||})$ is more complicated. The problem is nonetheless tractable in an important case, covered in Corollary 3.4 to Theorem 3.3, below. As a prerequisite to presenting those results, however, we introduce another graph theory construction, *edge contraction*.

3.3 Edge contraction and label dispersion

Edge contraction is an alternative way to collapse a BHC graph while preserving certain key complexities. For an edge $e \in G$, form a new graph G_e by “shrinking e to a point” — that is, remove the edge e and replace its two endpoints by one new vertex whose neighbors are the neighbors of both endpoints. If the edge $e = \{v, x\}$ is a leaf of G , where vertex v has degree 1, then contraction of edge e is equivalent to removing edge e and vertex v , leaving vertex x as it was. Let us call the process of contracting a graph on one edge *an elementary edge contraction*.

If the starting graph is a tree, then edge contraction always results in another tree. In contrast, a full quotient graph of a tree is not a tree, except in the trivial case where each equivalence class consists of a single point. If working with a tree, we can iterate elementary edge contractions successively to contract all the edges to one vertex. If the tree is a subgraph of a larger graph, then shrinking the tree to a point gives us a new graph that retains key properties of the original larger graph. Figure 4 illustrates the effect of one elementary edge contraction, contraction of an entire homogeneous subtree, and finally contraction of each of the maximal homogeneous subtrees. Let Γ denote the fully contracted graph (i.e., performing all possible elementary contractions of homogeneous edges in G). Note that the resulting Γ is independent of the order in which edge contractions occur.

Iterated edge contractions can yield a collapsed graph that, in a sense, lies between a given labeled graph G and its heterogeneous quotient $Q_{||}$. That intermediate graph is our route to proving Theorem 3.3. Theorem 3.3, together with Corollary 3.4, are the main results of this Section, and the theoretical basis for our intended complexity measurements. For a fixed number of subsidiaries in the holding company, the number of homogeneous patches that must be stitched together to compose G is an indication of its complexity. Of course, the number of different labels is a lower bound on this




Figure 4: Example of quotients via edge contractions: (a) Original tree, (b) contraction of one edge $\{2,5\}$, (c) contraction of the maximal homogeneous subtree on vertices $\{2,4,5,8,10\}$, and (d) contraction of all homogeneous subtrees.

Source: Authors' analysis

kind of complexity, and the cycle rank of the quotient graph Q_{\parallel} exactly captures the excess over that number.

Theorem 3.3. *Assume the BHC graph G is a tree, with labeled vertices, and V_1, V_2, \dots, V_N the partition of vertices of G into equivalence classes according to the vertex labels. Suppose G has exactly M maximal homogeneous subtrees according to this labeling. Then*

$$e(Q_{\parallel}) = M - 1 \tag{7}$$

and

$$b_1(Q_{\parallel}) = M - N . \tag{8}$$

Proof of Theorem 3.3. Let $\mathcal{T} = \{T_1, T_2, \dots, T_M\}$ denote the maximal homogeneous subtrees of G . Use successive edge contraction to shrink each subtree T_j to a single point to produce a quotient space, $\Gamma_{\mathcal{T}}$. The graph $\Gamma_{\mathcal{T}}$ has M vertices. The edges of $\Gamma_{\mathcal{T}}$ come from heterogeneous edges in G since each homogeneous edge in G has been contracted in the construction of $\Gamma_{\mathcal{T}}$.

$\Gamma_{\mathcal{T}}$ has no cycles. This follows from general topological theorems about shrinking contractible sets to points, but we can see it here in terms of edge-paths. It is convenient to think of the edges composing a path as oriented, having starting and ending points. An edge path in $\Gamma_{\mathcal{T}}$ is a sequence of edges e_1, e_2, \dots where each e_i corresponds to a heterogeneous edge \hat{e}_i in G . The edge \hat{e}_i starts at a vertex v_i in some maximal homogeneous subtree T_i and ends at another vertex w_i in a different T_j . The next edge e_{i+1} comes from an edge \hat{e}_{i+1} that starts at some vertex v_{i+1} that is contained in the same maximal homogeneous tree T_j as w_i . Since each maximal homogeneous tree is connected, we can find an edge path in T_j between w_i and v_{i+1} . If the quotient

$\Gamma_{\mathcal{T}}$ had a cycle, then we could construct a cycle in G consisting of heterogeneous edges in G corresponding to the edges in $\Gamma_{\mathcal{T}}$, together with homogeneous edge paths within various of the T_i . Because G is a tree, we conclude $\Gamma_{\mathcal{T}}$ has no cycles and so also is a tree.

Because $\Gamma_{\mathcal{T}}$ is a tree with M vertices, we know from equation (5) (the Euler characteristic), that $\Gamma_{\mathcal{T}}$ has $M - 1$ edges. But each heterogeneous edge in G appears exactly once in $\Gamma_{\mathcal{T}}$, so $M - 1$ is precisely the number of heterogeneous edges in G . On the other hand, each heterogeneous edge of G appears exactly once among the non-loop edges of Q , so $Q_{||}$ also has $(M - 1)$ edges. Because $Q_{||}$ has N vertices and is connected (G is a tree and therefore connected), we conclude:

$$b_1(Q_{||}) = e(Q_{||}) - v(Q_{||}) + 1 = M - N .$$

□

If G itself has cycles, as can occur when ownership is defined using a less than 50 percent threshold, then the above calculations become more complicated, because homogeneous cycles in G would disappear in $Q_{||}$ but heterogeneous cycles in G would still be represented in $Q_{||}$. We make this precise in the following corollary, which introduces the subgraph H (illustrated in Figure 5), which retains only the homogeneous edges of G .

Corollary 3.4. *Let G be a connected graph with labeled vertices and V_1, V_2, \dots, V_N the partition of vertices of G into equivalence classes according to the vertex labels. Suppose G has exactly M maximal homogeneous connected subgraphs according to this labeling, and let H denote the subgraph of G composed of these M disjoint homogeneous subgraphs. Then*

$$e(Q_{||}) = M - 1 + b_1(G), \text{ and} \tag{9}$$




Figure 5: Example of quotients via edge contractions: (a) Original tree, G , from Figure 4, and (b) the subgraph, $H \subset G$, composed of the seven disjoint maximal homogeneous subgraphs of G .

Source: Authors' analysis

$$b_1(Q_{||}) = M - N + b_1(G) - b_1(H). \quad (10)$$

Proof of Corollary 3.4. We know we can remove (cut) a number of edges of G to get down to a spanning tree of G , in fact $b_1(G)$ is the number of edges to cut. Start by partitioning the set of cycles in G into the set C_1 of cycles that include at least one heterogeneous edge, and the set C_0 of those that do not. Note that each cycle in C_0 must lie within one of the disjoint homogeneous subgraphs in H , while each cycle in C_1 must involve at least two of the subgraphs in H . To obtain a spanning tree T_G , we can cut $b_1(H)$ homogeneous edges to eliminate the cycles in C_0 , plus an additional $b_1(G) - b_1(H)$ heterogeneous edges to eliminate the remaining cycles in C_1 . Using in T_G the same vertex labeling as for G , apply Theorem 3.3 to T_G . The heterogeneous quotient $Q_{||}$ for G can be recovered from the heterogeneous quotient for T_G by adding back the edges in the quotient corresponding to the heterogeneous edges that were cut. So the number of edges in $Q_{||}$ is exactly $b_1(G) - b_1(H)$ more than the edges in the heterogeneous quotient for T_G , and so (from the Euler characteristic equation)

likewise for the cycle ranks. □

Remark 3. One can derive the subgraph H by deleting all heterogeneous edges from G .

Remark 4. If each simple cycle in G includes at least one heterogeneous edge, then the formula for the cycle rank of $Q_{||}$ simplifies to: $b_1(Q_{||}) = M - N + b_1(G)$.

3.3.1 Measuring complexity

The preceding results help measure the fragmentation of vertex labels throughout a BHC graph. We start by observing that a fundamental attribute of subsidiaries, such as charter type or regulatory jurisdiction, might be concentrated in one part of the BHC network, or might be dispersed throughout. The degree of fragmentation versus concentration is key in assessing the potential complexity of resolution, beyond simple measures of BHC size.

Suppose we have a BHC graph, with all the entities labeled in some category, e.g., geographical location or NIC entity type. As discussed in the previous section, the graph might be considered as simple as possible — or “perfect,” in the nomenclature of Lumsdaine et al. (2015) — if all the entities with each particular label were grouped in their own homogeneous connected subgraph. In such a perfect hierarchy, each label exists in exactly one maximal connected homogeneous subgraph, so the number of maximal connected homogeneous subgraphs in excess of the number of labels is a quantitative measure of the fragmentation of the labeling scheme throughout the network. For example, in Panel (b) of Figure 2, the number of maximal homogeneous connected subgraphs is five (including two separate NAT singletons) while in Panel (a) of Figure 2 it is eight.

Definition 3.1. For a given BHC graph, G , whose vertices are each classified (for example, by geographical location or NIC entity type) with one of N possible la-

bels, resulting in a partition of the vertices into M maximal homogeneous connected subgraphs, the *fragmentation* or *dispersion complexity* of this labeling is given by:

$$\text{dispersion complexity} = M - N. \tag{11}$$

When G is a tree, this is precisely the cycle rank $b_1(Q_{||}) = M - N$ calculated in equation (8) of Theorem 3.3. When G itself has cycles, but no purely homogeneous cycles, we can incorporate $b_1(G)$ as in Corollary 3.4

When the BHC network is “perfect” (i.e., each label’s equivalence class is contiguous), dispersion complexity by this definition is zero. Cycles in the quotient graphs combine two kinds of complexity. First, they embody cycles in the original BHC hierarchy, possibly complicated relationships such as joint control over a subsidiary. Alternatively, cycles in the quotient may capture regulatory, communication, and resolution complexities due to dispersion of entity types.

4 Empirical Application of the Method

We emphasize that our approach is not a purely theoretical exercise. Having defined the formal notation and results in Section 3, we return to the central question of how we might gauge the resolution complexity of a BHC. As a proof of concept, this section presents a case study that applies our proposed quotient metrics to a specific BHC. The purpose is to demonstrate the empirical feasibility of the approach, and to glean a preliminary sense of the orders of magnitude involved. Additionally, the purpose here is not to undertake a systematic empirical study of any BHC or set of BHCs. Done properly, that is a significant undertaking for future research, which we are pursuing separately.

We propose to measure the difficulty of resolution through the cycle rank of the

quotient graph, as set out in equations (6) or (10). In calculating the quotient, the choice of categories is critical. These should correspond to groupings within which communication and coordination is relatively easy, and between which communication is relatively difficult. Obvious choices are categories that correlate with regulatory jurisdictions, because these will correlate with existing trust relationships developed through past supervision, and with shared line-of-business expertise. In the case study in Section 4, we consider both NIC entity type and geographic location (countries and U.S. states) as the quotienting dimension.

4.1 Case Study: the Wells Fargo BHC

Our case study considers the structure of a single BHC over time. We extracted holding-company attributes and ownership relationships from the FFIEC’s public NIC database.¹⁷ The NIC database provides quarterly snapshots of all distinct subsidiaries (legal entities) in U.S. BHCs, together with a cumulative record of relationships among them. Timestamps on the relationships permit the construction of ownership networks as of any given date.

Of the eight original firms that submitted resolution plans, we selected Wells Fargo for our case study. We consider BHC complexity as it evolves over time, but focus on two dates, year-end 2006 and 2010, which are roughly equidistant from year-end 2008, when Wells Fargo’s acquisition of Wachovia was completed, as well as the time immediately surrounding the acquisition, namely the final two quarters of 2008 and the first quarter of 2009. Notably, neither firm was formally resolved (Wells Fargo is still a going concern), so although our measures capture certain aspects of complexity, we do not connect them directly to resolution in this example. At the

¹⁷See BoG (2016). The database is also available for public download; see FFIEC (2016). Many relationships in the NIC database lack information on the percentage of ownership, because there is no regulatory requirement to report this. We include all subsidiaries (i.e., we do not “prune” the graph), because all must be handled by the resolution process.

time of the acquisition, Wells Fargo was the fifth largest commercial bank in the U.S. and Wachovia was the fourth. Both firms appear in the NIC database separately prior to the acquisition; the combined firm also remains in the database afterward. By examining Wells Fargo, we are able to observe the evolution of the organizational structure following the acquisition. The other seven firms are not as appealing as a case-study subject for the following reasons:¹⁸

- Goldman Sachs Group, Inc. and Morgan Stanley are not available in the NIC database in the pre-crisis period.
- Like Wells Fargo, both JPMorgan Chase and Bank of America acquired large institutions during the crisis, Bear Stearns and Merrill Lynch, respectively. But because the acquired firms were investment banks and not commercial banks, there is no information corresponding to them in the NIC database in the pre-crisis period. In addition, the Federal Reserve was the primary regulator (at the holding company level) of both Wells Fargo and Wachovia, so any network structure effects that might be caused by an associated change in supervisor should not be present for the Wells Fargo-Wachovia example.
- Citigroup did not make any major acquisitions in the aftermath of the crisis.
- Bank of New York and State Street, respectively the eighth and 14th largest holding companies by total assets, are much smaller than the other top six holding companies by asset size. Both are less than half of the assets of Morgan Stanley (the sixth largest) and less than 16 percent of the assets of JPMorgan Chase (the largest). In addition, being monoline firms, they only have a small number of branches/entities and hence are not as interesting as networks.

The NIC database classifies the Wells Fargo top-level holding company as a

¹⁸Importantly, beyond these considerations, there was no other reason for choosing Wells Fargo from the set of large BHCs available as possible case studies. The selection of Wells Fargo as our case study does not imply a statement of its stability or instability relative to any other BHCs in the database.

“Financial Holding Company – Domestic.” The history of Wells Fargo within the database provides a glimpse into the nature of the data and BHC measurement. The NIC database emphasizes the identity of the formal legal charter in determining the surviving firm following a merger. As a result, the history as constructed from the NIC database will, in general, not be identical to the history as a firm might describe themselves; compare Wells Fargo’s (2017) own timeline with Table 1. Regardless, the Wells Fargo brand has a long history. The most familiar subsidiary — Wells Fargo Bank, National Association (RSSD #451965) — began as a state bank in San Francisco in 1852 to provide financial services in the wake of the California Gold Rush and did not become a national bank until 1968. The “high-holder” Wells Fargo holding company (RSSD #1120754), on the other hand, followed a very different path.

Table 1 (top panel), shows that Northwest Bancorporation was founded as a regional banking cooperative in Minneapolis in early 1929, shortly before the stock market crash and banking crisis of the early 1930s. Northwest Bancorporation formally reorganized as a BHC after the passage of the Bank Holding Company Act in the 1950s. Finally, in the late 1990s, the renamed Norwest Corporation merged with the California-based Wells Fargo BHC. The combined organization retained the Northwest charter, but renamed again to take advantage of the better-established Wells Fargo brand. In short, the NIC data on large financial firms embeds a wealth of historical and organizational detail that network analysis may help to disentangle.

Wachovia similarly has a long history. The original Wachovia Bank and Trust Co. (RSSD #392620) was chartered as a national bank in Winston-Salem, N.C. in 1879. In 2001, Wachovia Corp. (the parent financial holding company, RSSD #1136157, which itself was established in 1985), merged with a larger, but faltering, First Union Corp. (RSSD #1073551) of Charlotte, N.C. Legally, the combined entity retained the larger firm’s charter (and headquarters location), but renamed itself as Wachovia Corp. to exploit the more valuable brand. This “new” Wachovia Corp. (RSSD

#1073551) faltered again during the crisis, and was acquired by Wells Fargo & Co. (RSSD #1120754) in 2008.

Table 1: Histories of Wells Fargo & Co. and Wachovia Corp. in the NIC database

Date	Event
Wells Fargo & Co. (RSSD #1120754)	
1929-01-08	Northwest Bancorporation located at Minneapolis, MN was established as a Domestic Entity Other.
1956-05-09	Northwest Bancorporation changed from Domestic Entity Other to Bank Holding Company.
1983-05-01	Northwest Bancorporation was renamed to Norwest Corporation.
1984-12-31	Norwest Corporation moved to Sixth & Marquette Street Minneapolis, MN.
1998-11-03	Norwest Corporation was renamed to Wells Fargo & Company and moved to 420 Montgomery Street San Francisco, CA.
2000-03-13	Wells Fargo & Company changed from Bank Holding Company to Financial Holding Company – Domestic.
Wachovia Corporation (RSSD #1073551)	
1968-05-04	Cameron Financial Corporation located at Charlotte, NC was established as a Domestic Entity Other.
1970-12-31	Cameron Financial Corporation changed from Domestic Entity Other to Bank Holding Company.
1974-12-31	Cameron Financial Corporation was renamed to First Union Corporation.
1992-12-03	First Union Corporation moved to One First Union Center Charlotte, NC.
2000-03-13	First Union Corporation changed from Bank Holding Company to Financial Holding Company – Domestic.
2001-09-01	First Union Corporation was renamed to Wachovia Corporation and moved to One Wachovia Center Charlotte, NC.
2008-12-31	Wachovia Corporation was acquired by Wells Fargo & Company.
<small>Source: Federal Financial Institutions Examination Council</small>	

When Wells Fargo acquired Wachovia in 2008, the latter firm was struggling, due in significant part to its own 2007 acquisition of World Savings Bank, F.S.B. (RSSD #1157433). World Savings was part of Golden West Financial Corp., a large California-based mortgage lender that Wachovia had purchased at the peak of the

subprime mortgage lending boom in 2006. Although experiencing difficulties, Wachovia did not fail, and Wells Fargo absorbed it as a going concern as of year-end 2008. Figure 6 presents the separate ownership hierarchies of Wells Fargo and Wachovia in 2006, and the merged hierarchy as of 2010. Figure 7 depicts the diaspora of the former Wachovia subsidiaries in the combined BHC, seven years after the merger. In many cases (blue edges in Figure 7), former Wachovia subsidiaries now own non-Wachovia subsidiaries. Figure 7 suggests that a partitioning (i.e., quotienting) scheme based on the provenance of subsidiaries may itself be of interest.

Given that cycles in the quotient graph are our key intended metric of resolution complexity (i.e., a higher cycle rank indicates greater complexity), it is important to note that a BHC ownership hierarchy will not, in general, form a tree. Cycles will typically exist in a large BHC graph, even before quotienting. For example, Figure 8 depicts some of the cycles hidden in the BHC in panel (c) of Figure 6. These involve some of the most important entities in the BHC, including the high-holder financial holding company and the flagship national bank.¹⁹ The top-level holding company has direct ownership in the national bank, but this stake is not 100 percent. Instead, Figure 8 shows that there are several lower-level BHCs and other entities owned by the Wells Fargo high holder, and which maintain partial interest in the national bank. It is possible, given the exact anatomy of the subgraph in Figure 8, together with various summary statistics describing the rest of the BHC, to calculate the total number of maximal connected subgraphs in the overall BHC cycle graph, along the lines of Theorem 3.3 and its corollary. Figure 9 illustrates our various quotients, defined in section 3.2 — with/without self-loops and with/without condensing — applied to the cyclic subgraph from Figure 8.

¹⁹Figure 8 shows only a subset of key cycles, in the interest of presenting a legible visualization. These cycles involve only those edges for which ownership percentages are recorded in the NIC database. As Table 2 indicates, the cycle rank of the Wells Fargo hierarchy in 2010 exceeded 450 independent cycles.




Figure 6: Merger of two BHC hierarchies. (a) Wells Fargo & Co. in 2006; (b) Wachovia Corp. in 2006; and (c) the merged firm in 2010, after Wells Fargo's 2008 acquisition of Wachovia. In each panel, vertices represent the top-level BHC and all subsidiary entities; edges represent ownership relationships. Red edges are heterogeneous by entity type.

Source: Federal Financial Institutions Examination Council; authors' analysis




Figure 7: Wells Fargo & Co. at year-end 2015, highlighting former Wachovia subsidiaries. A blue edge indicates a former Wachovia subsidiary owning a non-Wachovia subsidiary (i.e., a Wells Fargo subsidiary that was not present in the Wachovia hierarchy at the 2008 acquisition). A red edge indicates a non-Wachovia subsidiary owning a former Wachovia subsidiary.

Source: Federal Financial Institutions Examination Council; authors' analysis




Figure 8: Ownership cycles in Wells Fargo BHC hierarchy, 2010, before quotienting.

Source: Federal Financial Institutions Examination Council; authors' analysis

4.2 Before and after the Wachovia acquisition

A large acquisition or merger is a major undertaking, and we expect it to create a significant dispersion of entities of various types within the post-merger holding company. The primary issue for the post-merger firm should be ease of coordinated management among the conglomerated activities. Similarly, the primary concern of a regulator tasked with approving such a merger might be ensuring that the combined holding company does not pose greater risk to financial stability, for example if the merger creates greater challenges in coordinating risk management or coordinating supervision. Given consolidation planning and accounting at the BHC level, the details of ownership complexity may be a formality of secondary consideration for management of the going concern. Even presuming a strong business motivation to coordinate the location of similar business lines, such as banking or insurance sub-




Figure 9: Four quotients by entity type, for the Wells Fargo cycles of Figure 8: (a) Full quotient, Q , (b) Heterogeneous quotient, $Q_{||}$, (c) Condensed quotient, \bar{Q} , and (d) Condensed heterogeneous quotient, $\bar{Q}_{||}$.

Source: Federal Financial Institutions Examination Council; authors' analysis

sidiaries, as close neighbors within the merged BHC ownership hierarchy, however, it could take months or years to iron out the legal details to achieve this structure. We might therefore expect an initial jump in complexity when comparing the pre- and post-acquisition BHC. Over time, if the acquired firm has been successfully integrated into the acquiring firm, the organizational complexity should wane and perhaps return to pre-acquisition levels. Wells Fargo's extensive operations are represented by its BHC hierarchy that contains at least one instance of almost every entity type identified in the NIC schema, both before and after the crisis and Wachovia acquisition.

Table 2 presents size and complexity statistics based on our intended metrics for the WFC and Wachovia BHCs at four specific dates: Wells and Wachovia separately at Q4 2006 and Q3 2008 (immediately prior to the merger), and for the combined firm in Q4 2008 and Q4 2010. The top panel lists two size-based measures (the number of edges and the number of vertices) and the cycle rank of the overall firm. The middle and bottom panels of Table 2 present various intended complexity metrics for the BHC snapshots after quotienting by the entity type of each vertex, and by the geographic location (U.S. state or foreign country) of the entity headquarters.

Table 2: Holding Company Complexities for Wells Fargo & Co. and Wachovia Corp.

Intended Measure	Wells F. Q4 2006	Wachovia Q4 2006	Wells F. Q3 2008	Wachovia Q3 2008	Wells F. Q4 2008	Wells F. Q4 2010
Statistics on unquotiented BHC graphs						
Vertex count, $v(G)$	609	1388	612	1037	1710	1544
Edge count, $e(G)$	648	1688	644	1442	2158	2025
Cycle rank, $b_1(G)$	40	301	33	406	489	482
Cycle ranks (b_1) and component counts (M, N), quotienting by entity type						
Full quotient, $b_1(Q)$	636	1674	631	1429	2143	2011
Heterogeneous quotient, $b_1(Q_{ })$	296	457	301	678	1062	1054
Condensed full quotient, $b_1(\bar{Q})$	18	19	18	18	30	23
Condensed heterogeneous quotient, $b_1(\bar{Q}_{ })$	14	15	14	14	25	19
Edge-contracted graph, $b_1(\Gamma)$	20	8	14	10	22	16
Disjoint homogeneous subgraph, $b_1(H)$	7	79	6	12	17	8
Number of disjoint homogeneous components, M	276	249	288	298	646	595
Number of labels, N	13	14	14	14	16	15
Cycle ranks (b_1) and component counts (M, N), quotienting by geographic jurisdiction						
Full quotient, $b_1(Q)$	568	1616	564	1373	2051	1922
Heterogeneous quotient, $b_1(Q_{ })$	352	955	341	684	1191	1176
Condensed full quotient, $b_1(\bar{Q})$	85	58	87	67	194	172
Condensed heterogeneous quotient, $b_1(\bar{Q}_{ })$	65	38	69	38	165	146
Edge-contracted graph, $b_1(\Gamma)$	21	13	18	21	42	52
Disjoint homogeneous subgraph, $b_1(H)$	6	22	6	8	11	5
Number of disjoint homogeneous components, M	399	748	395	356	861	803
Number of labels, N	81	72	81	70	108	104

Source: Federal Financial Institutions Examination Council; authors' calculations

The statistics include cycle ranks for all four quotient graphs for each period and classification.

Table 2 provides some preliminary evidence on the effectiveness of the quotienting approach. First, the results demonstrate that the basic calculations are feasible, using actual holding company data. Second, the quotient cycle rank magnitudes give a sense of the magnitude of complexity involved across two different classifications, entity-type and jurisdiction. Of particular interest are the heterogeneous quotient cycle ranks, $b_1(Q_{||})$, and their condensed counterparts, $b_1(\bar{Q}_{||})$. Each independent cycle in these indicates the presence of a heterogeneous relationship that crosses entity-type or geographical boundaries. Unwinding each such relationship would typically involve a conversation with at least two separate primary regulators.

Comparing all of the intended metrics of Wells Fargo in Q4 2006 with those in Q3 2008 (columns 2 and 4), it appears that Wells Fargo had remarkably stable levels of complexity leading up to the merger, both in terms of entity-type and jurisdiction type. In contrast (compare columns 3 and 5), while Wachovia reduced its size (i.e., the number of vertices and edges), the cycle rank measures, b_1 , tell a more nuanced story. In terms of entity quotients, while the cycle rank of its full quotient, $b_1(Q)$, declined, the cycle rank of the heterogeneous quotient, $b_1(Q_{||})$, increased by almost 50 percent (678/457). In addition, the homogeneous cycle rank, $b_1(H)$, dropped dramatically; was six-fold higher in 2006 by entity quotient and nearly three times higher by jurisdiction. The homogeneous subgraph count (M) increased by 20 percent (298/249). There was little change in the number of labels, N , however, so that the condensed quotient cycle ranks remained quite stable. The edge contraction cycle rank, $b_1(\Gamma)$, increased by both entity type (25 percent) and geographic jurisdiction (60 percent). In contrast, Wachovia's jurisdictional complexity appears to have declined in the run-up to the merger, although the condensed quotient cycle rank, $b_1(\bar{Q})$, increased. That the condensed quotient cycle increased but the heterogeneous condensed quotient cy-

cle rank, $b_1(\bar{Q}_{||})$, decreased indicates an increase in within-supervisor coordination and a reduction in cross-supervisor coordination.

In terms of entity complexity, at the point of the merger Wells Fargo and Wachovia were quite similar in many ways (compare columns 4 and 5). Although Wachovia may appear complex based on the proposed unquotiented metrics and even the full quotient and heterogeneous quotient, the other proposed quotients are identical. The edge contractions tell a different story, however, with Wachovia having a lower contraction cycle rank and double the homogeneous cycle rank. In terms of jurisdictional complexity, Wells Fargo actually looks more complex along a number of dimensions — larger cycle ranks in the two condensed quotient graphs and in terms of the number of connected components, M and the number of labels.

The table clearly shows that these proposed metrics are not additive, as none of the intended measures of the merged firm (column 6) is the sum of the intended individual measures (columns 4 and 5). More specifically, for most intended metrics, the merged firm has more complexity than the sum of the separate (pre-merger) complexities. Notably, the cycle ranks increase after the Wachovia acquisition, even beyond just the sum of the cycle ranks of the two separate firms. There is a roughly four percent increase in the entity full-quotient cycle rank, and a nearly six percent increase in geography full-quotient cycle rank. This increase mirrors the 11 percent increase in the overall firm's cycle rank.

Consistent with our intuition, complexity as defined by our intended measures declines over time. In almost every case, the numbers in 2010 are lower than in Q4 2008. The only exception is the jurisdictional edge contraction.




Figure 10: Evolution of the Wells Fargo ownership hierarchy, quarterly from 1986 to 2016. Top panel: Cycle rank of the heterogeneous quotient, $b_1(Q_{||})$, for both the entity-type (blue) and the jurisdiction (gold) vertex partitions. Bottom panel: Counts of edges (red dashes) and vertices (red dots) for the full (unquotiented) graph; and number of maximal homogeneous subgraphs, M , for both the entity-type (blue) and the jurisdiction (gold) vertex partitions.

Source: Federal Financial Institutions Examination Council; authors' analysis




Figure 11: Evolution of the Wachovia ownership hierarchy, quarterly from 1986 until its acquisition in 2008. Top panel: Cycle rank of the heterogeneous quotient, $b_1(Q_{||})$, for both the entity-type (blue) and the jurisdiction (gold) vertex partitions. Bottom panel: Counts of edges (red dashes) and vertices (red dots) for the full (unquoted) graph; and number of maximal homogeneous subgraphs, M , for both the entity-type (blue) and the jurisdiction (gold) vertex partitions.

Source: Federal Financial Institutions Examination Council; authors' analysis

Finally, we investigate the evolution of the Wells Fargo ownership hierarchy from 1996 to 2016, on a quarterly basis. For each quarter, we study the topological properties of the ownership structure, showing the number of vertices and edges, as measures of the size of Wells Fargo. In addition, we compute the cycle rank, b_1 , for two different quotient procedures, the entity-type quotient and the jurisdiction-based quotient. Figures 10 and 11 present the results. All four intended measures of the evolution of the Wells Fargo ownership hierarchy clearly show the impact of the Wachovia acquisition. The numbers of vertices and edges are highly correlated. They exhibit slow and gradual changes until the Wachovia acquisition, at which point they exhibit a significant jump. Interestingly, following the acquisition, there is a nearly 20 percent increase in the cycle rank, revealed in the difference in the number of edges versus the number of vertices. This suggests that following the acquisition, there are significantly more cross-ownership relationships.

As the merger episode evolved post-crisis, the size of the combined firm (measured by vertex and edge counts) remained relatively stable; on the other hand, from a resolution perspective, the firm became more complex, due to the increase in the number of cycles. Studying the two quotient-derived cycle ranks, we find a few key insights into the topology of Wells Fargo. First, it is clear that the cycle ranks vary substantially more than the size-based metrics. As such, they potentially provide a quantitative measure for complexity that is not simply a function of the institution's size. Second, there is greater variation in the cycle rank based on the organizational structure of entity types, versus the cycle rank based on the organizational structure of jurisdiction types. Third, there is a clear spike in both cycle ranks as a result of the Wachovia acquisition. This is consistent with the interpretation that organizational complexity increased dramatically as a result of the acquisition. Following this event, both cycle ranks have been decreasing on a quarter-to-quarter basis. This is consistent with the interpretation that post-merger, the firm is returning to a reduced

level of complexity as the Wachovia entities are fully subsumed into the Wells Fargo organizational structure. Finally, it is possible to observe that during the merger wave in the early years of the dataset, the entity-based quotient cycle rank increased from 1996 to 2001, while the size related measures were mostly constant.

5 Conclusions

Large bank holding companies (BHCs) are structured into intricate ownership hierarchies involving hundreds of legal entities. If a firm's financial distress deteriorates to the point that resolution is required, regulators will be called in to help. The difficulty — indeed, feasibility — of the task depends on the complexity of the struggling organization. It is therefore useful to understand, *ex ante*, the difficulty of the resolution challenge.

We emphasize that complexity should not be judged out of context as undesirable in some absolute sense. Certain beneficial activities naturally require complex organizations. The *relative* costs and benefits of organizational complexity are outside of the scope of this paper, but remain important topics for future research. Nonetheless, complexity measurement is critical for evaluating the efficacy of the current organizational structure of a firm and anticipating possible resolution challenges associated with coordination across different resolution authorities. For example, in our case study in Section 4, we observe that large holding company mergers induce a pronounced transient spike in complexity as defined by our proposed metrics, which gradually dissipates over subsequent quarters as the firm rationalizes its post-merger structure. Complexity measurement is also crucial for assessing the difficulty of post-failure resolution, when there are no going-concern benefits of complexity left to balance the costs. In the absence of regulation, post-failure resolution would not be an important consideration for current management, but is a prime consideration for

resolution authorities.

We have proposed a set of metrics, based on the cycle rank of the quotient graph of the BHC hierarchy, to help assess this complexity at any point. The proposed measures are mathematically grounded, intuitively sensible, and easy to implement. We illustrate the process with a simple case study of the Wells Fargo & Co. holding company. The geographical and entity type complexity of the Wells Fargo holding company increases after the acquisition of Wachovia during the financial crisis. Although this example demonstrates the potential of our proposed measures to show the evolution of a single BHC's complexity over time, our proposed measures provide no direct evidence of the complexity of its resolution, as Wells Fargo is still a going concern. A comprehensive econometric analysis of resolution complexity using our measures is an important topic for future research.

The feasibility and plausibility of the approach raise a number of possibilities for future research. First, in quotienting by entity type, we encountered large numbers of miscellaneous entities (which the NIC dataset lists as “domestic entities – other”). Many of these may be inconsequential or trivial to resolve, but it is difficult to make broad assertions, because the range of business lines is so broad. This points to a need for more granular resolution of entity types in the holding company data. Second, longer-range and higher-frequency analysis would help reveal whether the observed jump in complexity with the Wachovia merger was unusual, or something that falls within the ordinary range of variation for large BHCs. Lastly, we hope to make cross-sectional comparisons, to assess how Wells Fargo's experience relates to that of other large BHCs.

References

- S. Battiston, J. D. Farmer, A. Flache, D. Garlaschelli, A. G. Haldane, H. Heesterbeek, C. Hommes, C. Jaeger, R. May, and M. Scheffer. Complexity theory and financial regulation. *Science*, 351(6275):818–819, 2016.
- Board of Governors of the Federal Reserve System and Federal Deposit Insurance Corporation (BoG-FDIC). Resolution plans required. *Federal Register*, 76:67323–67340, November 2011.
- Board of Governors of the Federal Reserve System and Federal Deposit Insurance Corporation (BoG-FDIC). Resolution plan assessment framework and firm determinations (2016). Technical report, Federal Reserve, April 2016. URL <https://www.federalreserve.gov/newsevents/press/bcreg/bcreg20160413a2.pdf>.
- Board of Governors of the Federal Reserve System (BoG). Bulk data download: Data dictionary and reference guide, version 1.1. Technical report, Federal Reserve System, 2016. URL <https://www.ffiec.gov/nicpubweb/content/DataDownload/NPW%20Data%20Dictionary.pdf>.
- S. Bright, P. Glasserman, C. Gregg, and H. Hamandi. What can we learn from publicly available data in banks’ living wills? *OFR Brief*, (Series No. 16-05), 2016.
- J. Carmassi and R. Herring. Living wills and cross-border resolution of systemically important banks. *Journal of Financial Economic Policy*, 5(4):361–387, 2013.
- J. Carmassi and R. Herring. Corporate structures, transparency and resolvability of global systemically important banks. Technical report, U. of Pennsylvania, August 2014. URL http://www.systemicriskcouncil.org/wp-content/uploads/2015/01/Carmassi-Herring_Corporate-Structures-Transparency-and-Resolvability-of-G-SIBs.pdf.
- N. Cetorelli and L. S. Goldberg. Measures of global bank complexity. *Economic Policy Review*, 20:107–126, 2014.
- J. Chan-Lau, M. Espinosa, K. Giesecke, and J. Solé. Assessing the systemic implications of financial linkages. *IMF Global Financial Stability Report*, 2, 2009.
- R. Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937.
- R. Cont, A. Moussa, and E. B. Santos. Network structure and systemic risk in banking systems. In J.-P. Fouque and J. A. Langsam, editors, *Handbook on Systemic Risk*, chapter 13, pages 327–368. Cambridge University Press, 2013.

M. Corkery. Wells Fargo's 'living will' plan is rejected again by regulators. *New York Times*, December 2016.

Federal Deposit Insurance Corporation. FDIC's role and authorities under the new financial reform law. Technical report, FDIC, November 2010. URL <https://www.fdic.gov/regulations/reform/role.html>.

Federal Deposit Insurance Corporation. The orderly liquidation of Lehman Brothers Holdings Inc. under the Dodd-Frank Act. *FDIC Quarterly*, 5(2), 2011.

Federal Deposit Insurance Corporation. The resolution of systemically important financial institutions: The single point of entry strategy. *Federal Register*, 78(243): 76614–76624, December 2013.

Federal Deposit Insurance Corporation. Failed bank list. Internet site, 2016a. URL <https://www.fdic.gov/bank/individual/failed/banklist.html>.

Federal Deposit Insurance Corporation. Failures and assistance transactions: Number of institutions: US and other areas: 1934-2016. Internet site, 2016b. URL <https://www5.fdic.gov/hsob/>.

Federal Deposit Insurance Corporation and Board of Governors of the Federal Reserve System (FDIC-BoG). Guidance for 2017: 165(d) Annual resolution plan submissions by domestic covered companies that submitted resolution plans in July 2015. Technical report, FDIC-FRB, 2016. URL <https://www.fdic.gov/news/news/press/2016/pr16031b.pdf>.

Federal Deposit Insurance Corporation and Board of Governors of the Federal Reserve System (FDIC-BoG). Agencies announce wells fargo has remediated resolution plan deficiencies. FDIC press release PR-32-2017, FDIC-FRB, 2017. URL <https://www.fdic.gov/news/news/press/2017/pr17032.html>.

Federal Financial Institutions Examination Council (FFIEC). National Information Center. Internet site, 2016. URL <https://www.ffiec.gov/nicpubweb/nicweb/NicHome.aspx>.

Financial Crisis Inquiry Commission. *The Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States, Official Government Edition*. U.S. Government Printing Office, January 2011.

Financial Stability Board. Recovery and resolution planning for systemically important financial institutions: Guidance on developing effective resolution strategies. Technical report, FSB, 2013. URL <http://www.fsb.org/wp-content/uploads/>

r_130716b.pdf?page_moved=1.

- M. J. Fleming and A. Sarkar. The failure resolution of Lehman Brothers. *FRBNY Economic Policy Review*, pages 175–206, December 2014.
- M. Flood. The great deposit insurance debate. *Federal Reserve Bank of St. Louis Review*, 74(4):51–77, July/August 1992.
- P. Glasserman and P. Young. Contagion in financial networks. *Journal of Economic Literature*, 54(3):779–831, September 2016.
- J. Grant, R. J. Bricker, T. J. Fogarty, and G. Previts. Complexity and analyst following of multinational firms. Technical report, Case Western Reserve University, January 2000. URL <https://ssrn.com/abstract=220421>.
- A. G. Haldane. Rethinking the financial network. Speech delivered at the Financial Student Association, Amsterdam, April, Bank of England, 2009.
- R. Haltom. Failure of continental illinois, may 1984. Technical report, Federal Reserve Bank of Richmond, 2013. URL <http://www.federalreservehistory.org/Events/DetailView/47>.
- F. Harary. *Graph Theory*. Addison-Wesley Series in Mathematics. Addison Wesley, 1969.
- R. Herring. The challenge of resolving cross-border financial institutions. *Yale Journal on Regulation*, 31(3):853–881, 2014.
- G. G. Kaufman. Too big to fail in banking: What does it mean? *Journal of Financial Stability*, 13:214–223, August 2014.
- A. Kleinbaum, T. Stuart, and M. Tushman. Discretion within constraint: homophily and structure in a formal organization. *Organization Science*, 24(5):1316–1336, 2013.
- R. Lumsdaine, D. Rockmore, N. Foti, G. Leibon, and J. D. Farmer. The intrafirm complexity of systemically important financial institutions. Technical report, American University, 2015. URL <http://ssrn.com/abstract=2604166>.
- W. Massey. *A basic course in algebraic topology*. Springer-Verlag, 1991.
- R. M. May. Networks and webs in ecosystems and financial systems. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1987):20120376, 2013.

- T. Squartini, I. van Lelyveld, and D. Garlaschelli. Early-warning signals of topological collapse in interbank networks. *Nature*, 3(3357), 2013.
- R. P. Stanley. Topics in algebraic combinatorics. Technical report, Massachusetts Institute of Technology, February 2013. URL <http://www-math.mit.edu/~rstan/algcomb/algcomb.pdf>.
- M. Summer. Financial contagion and network analysis. *Annual Review of Financial Economics*, 5(1):277–297, 2013.
- R. Tracy and R. Rubin. Treasury nominee Steven Mnuchin says bank regulation should be tailored to activity. *Wall Street Journal*, January 2017.
- L. Urwick. The manager’s span of control. *Harvard Business Review*, pages 39–47, May-June 1956.
- A. R. Valukas. Report of Anton R. Valukas, examiner: Sections I & II: Introduction, executive summary & procedural background. Technical report, U.S. Bankruptcy Court, Southern District of New York, 2010.
- S. Vitali, G. J.B., and S. Battiston. The network of global corporate control. *PLoS ONE*, 6(10):e25995, 2011.
- Wells Fargo & Co. Wells Fargo history: Timeline. Online report, Wells Fargo, 2017. URL <https://www.wellsfargohistory.com/timeline/>.