DATA SCIENCE IN SCIENCE
2023, VOL. 2, NO. 1, 2151949
https://doi.org/10.1080/26941899.2022.2151949

Taylor & Francis
Taylor & Francis Group

OPEN ACCESS Check for updates

Learning from Lending in the Interbank Network

Paul Laux[a], Wei Qian[b], and Haici Zhang[c]

[a]Department of Finance, Institute for Financial Services Analytics, University of Delaware, Newark, Delaware, USA; [b]Department of Applied Economics and Statistics, Institute for Financial Services Analytics, University of Delaware, Newark, Delaware, USA; [c]Institute for Financial Services Analytics, University of Delaware, Newark, Delaware, USA

ABSTRACT

Empirical analysis of a major overnight-funding network of European banks shows that, when liquidity disruptions occur in a part of the network, lending banks in other parts of the network broaden their cohorts of borrowers in the part of the network that is subject to the disruptions. Measures of this broadening are useful new statistics for the amount of information conveyed from one part of the network to another. In our setting, we call this broadening "counterparty sampling," and present evidence that it improves the network's stock of information about future interest rates. By comparing to linkages forecast by an LSTM deep learning model for counterparty linkages, we find that the extent of surprising new linkages predicts lower future rates. Our evidence supports the idea that interbank funding networks provide benefits of learning and information aggregation, and our measures suggest new ways of looking at sparse networks with stable structures but dynamically-changing linkages.

ARTICLE HISTORY
Received 3 August 2022
Revised 16 November 2022
Accepted 21 November 2022

KEYWORDS
Entropy; information; network statistics; interbank network; LIBOR; overnight funds

1. Introduction

Different from a social network where linkages tend to be persistent, other important network settings, such as traffic, weather, and financial trading, have stable node structures but their linkages or linkage intensities are often sparse and fleeting. With a focus on modern finance studies, we investigate one such financial network, the overnight-funding network of commercial banks. In this network, banks loan each other deposits for short periods of time to satisfy the liquidity (cash) needs of the borrowing banks and earn fees for the lending banks. This specific network is empirically important in itself and has been a topic of interest in modern finance due to economy-wide implications.[1] Our work intends to suggest new network theory inspired data metrics for interbank networks and employ data science tools in econometrics and deep learning to systematically investigate whether banks are learning macro information *via* broadening their lending connections and whether banks' re-configurations of lending links in reaction to routine shocks to the liquidity needs of other banks is predictive of future interest rates

A recent important event in the fed funds market illustrates the need for our investigation. On September 17, 2019, the benchmark overnight unsecured fed funds rate, a measure of the fee that lending banks receive from borrowing banks, spiked to 5.25%, after trading in a range of 2.09–2.20% over the previous month. The mean daily absolute deviation that week was 129 basis points, after averaging 3 basis points over the year before that point. The precipitating event was an unexpected shortage of liquid reserves, the portion of deposits retained by banks in cash, not lent out (Kaminska 2019). An analysis by Federal Reserve Board staff comments: "Lenders did not appear to step into the market to take advantage of the higher rates … Borrowers generally trade with the same group of lenders at similar volumes every day and even market volatility as substantial as what we saw in mid-September does not seem to change that" (Anbil et al. 2020). Apparently, in this case, banks did not reconfigure their lending relationships in response to the underlying liquidity shock. The Federal Reserve staff report suggests that interbank lending failed to help the U.S. fed funds market to learn its new equilibrium rate.

By studying a time series of exogenous liquidity stresses with regular periodicity in the European interbank market, we provide an inference that this failure may be more the exception than the rule.[2] Consistent with the observations above regarding September 2019, the topology of the European interbank network changes little with the onset of stress and uncertainty. Under the surface, however, we find that distant banks, i.e. those in less-stressed parts of the network, re-establish older lending ties with banks in the most-stressed part, and also broaden lending ties to include more borrowers. The distant lending banks include the large money-center banks that comprise the rate-setting panel for 3-month LIBOR, a key interest rate for bank loans. On a time-series basis, the breadth of distant banks' lending cliques—groups of borrowing banks from the liquidity-stressed

CONTACT Paul Laux laux@udel.edu Department of Finance, Institute for Financial Services Analytics, University of Delaware, Newark, DE, USA

regions—predicts 3-month LIBOR changes a week later. This predictability is most dependable for the day of the week on which the monetary authority routinely makes its policy announcements, suggesting that information relevant to upcoming macroeconomic policy determinations is conveyed from the stressed groups of banks to the groups that lend to them.

We choose to study the market for overnight funds in Europe because a rare opportunity existed for researchers to observe interbank deposits trading directly in an electronic market known as "e-MID" during the early 2000s. In its heyday, most European banks were members of the e-MID system.[3] The typical e-MID transaction occurs when a borrower bank quotes an interest rate at which it would like to borrow overnight (with the quote and bank identifier visible to all member banks), and where a lending bank then accepts the quote to lend at that rate—technically sells a deposit (Angelini et al. 2011; Beaupain and Durré 2008). Settlement is electronically automatic upon the quote acceptance, so any off-line information gathering or conversations between banks must take place first.

During the period of its operation, e-MID is headquartered in Italy, and the largest *number* of banks actively trading on it are Italian banks. In contrast, the *largest* banks on the network are non-Italian banks. Moreover, Italian banks experience a repeated exogenous uncertainty-increasing liquidity event on a roughly monthly basis during the period of our data: Italian banks are required to transmit their customers' ongoing tax payments to the government on a specific schedule each month. The tax payment amounts are systemically important and difficult to predict: the European Central Bank states that it considers the time-varying bank system transfers to the Italian government to be a significant "autonomous factor" influencing the varying liquidity needs of the banking system overall (see https://www.ecb.europa.eu/mopo/liq/html/treas.en.html).[4] Uncertainty is high during the period just before the tax transmittal period when the tax collection process is underway, and then is resolved (i.e. the liquidity in the system becomes known) as the payments are transmitted to the government. We, therefore, compare network activity differences during the more-uncertain "pre-tax-transmittal periods" (PTTPs) when information about the balance of reserves is relatively scarce, especially from the point of view of non-Italian banks, as compared to later adjacent "tax-transmittal periods" (TTPs), periods when uncertainty is resolved. For our investigation, we designate TTPs and PTTPs to be 5 business days centered around the tax-transmittal date and the same day of the previous week, respectively.

The non-Italian banks on the e-MID network include all the large money-center banks that are part of the rate setting-panels for LIBOR rates during this period. A central idea for our investigation is that these large non-Italian banks that are important for establishing key rates, such as 3-month LIBOR face an information deficit about an important component of system-wide liquidity during pre-tax-transmittal periods. If they choose, they can lend to Italian banks overnight at competitive rates *via* e-MID, gaining credit information and perhaps other insight into the Italian liquidity situation. If micro and macro learning are complimentary (Hirshleifer and Sheng 2022), this information might help them to better understand the developing equilibrium for rates important to their loan business.

The idea that network connections serve as a conduit for information flow has intuitive appeal and has been much considered in the literature on social structure and economics (Granovetter 2005). The idea that ties bridging to distant parts of a network are important for information originated in sociology (Granovetter 1973, 1983). A node's "strong" locally-connected nodes are apt to also be connected to each other or at least to be similar in important ways, and thus might not conduct much new information. In contrast, Burt (1992) explicates a key reason for "the strength of weak ties": though no particular weak tie (i.e. nodes that are distantly but not locally connected) is of much importance, it is weak ties overall that disproportionately bridge different parts of a network and convey distant information. Our investigation can be characterized as measuring the extent to which weak ties are important for linking different bank cohorts and enabling information to flow between them.

Pfaltz (2012) considers these ideas in the context of the mathematical network theory. Separated parts of networks, defined as parts with no direct bridges, can only be joined by adding edges (i.e. links or connections). By linking previously distant, tightly-interconnected clusters (such as triads), such additional edges move the network toward being fully connected. Pflatz notes that adding such bridging edges can lead to a meaningful increase in connectedness while implying only small changes in standard summary measures of network topology.[5] Pfalz suggests that one natural measure of change in this direction would then be the increase in entropy, an index of disorder. We will use the change in entropy as one statistical indicator of possible bridge-building when systems become stressed. We also suggest, and show the usefulness of, some intuitive measures of changes in banks' lending cliques, such as the number of trading partners that are brought back from previous spates of system stress and the cross-sectional variation in banks' sets of trading partners.

Having described the setting and network reasoning, we can now state our findings more specifically. Using vector autoregressions (VARs), we establish that recent changes in the nature of distant banks' lending cliques Granger-causes interest rates (i.e. Granger non-causality is rejected). Additionally, we find that changes in rate volatility do not Granger-cause future rates. That is, the lending-clique related actions of banks are the key, not the volatility that may have led to them. Further, the evidence for Granger causality is strongest when we focus on changes that occur on the day of the week when the European Central Bank (ECB) Governing Council makes it key monetary policy decisions for the Eurozone, suggesting that changes in the nature of lending cliques may help distant banks learn about how recent liquidity pressures feed into ECB decisions. This core evidence is consistent with the idea that the distant banks that set 3-month LIBOR learn about the implications of the Italian liquidity stress for future rates *via* the broader lending cliques that they establish. We call this behavior "counterparty sampling" because it has the effect of adding to the quantity of information about future interest rates that is present in the network.

Even so, it is possible that counterparty sampling is a reaction to information banks already have, rather than a way of gathering new information. This possibility is highlighted because counterparty sampling increases are directionally associated with increases in rates. Banks might broaden their lending cliques to better profit from upcoming higher rates. To investigate, we measure the surprises in counterparty sampling by comparing link outcomes to the link predictions of a standard long short-term memory (LSTM) deep learning model and then investigate whether the extent of surprising new linkages is also predictive of future interest rates. The LSTM has appeal as a prediction mechanism because it allows for non-linear connections (which the VARs do not), because it allows linkage patterns from further in the past to revive in importance (in a way that typical time-series models do not), and because it deals well with the relatively sparse linkage matrices that characterize interbank networks (in which banks tend to have lending links to only a few other banks). We find that the extent of surprises is predictive of *lower* future rates. The positive association with linkages overall thus emanates from the predictable part of changes in linkages, which could reflect banks' reacting to add borrowers based on expecting better rates. But such a reaction is not a plausible explanation for the negative association with surprises.

Overall, we establish the existence of counterparty sampling in interbank lending relationships as a response to routine spates of uncertainty and provide evidence that suggests counterparty sampling is predictive of future rates of direct concern to banks. The evidence is supportive of the idea that banks learn in the lending process.

Our paper contributes to the literature in several ways. Regarding networks generally, we first propose several data metrics inspired by network topology theory for assessing the extent of dynamic changes in interbank lending linkages, and show that even when network topology does not change in response to shocks, changes in specific linkages can have strong effects on the functioning of the network. Additionally, we show that these change measures map the movement of economic information from one part of the network to another, and the breadth of nodes' linkage sets can be effective indexes relating to information transmission to distant nodes. By further extending our empirical analysis to a dynamic deep learning model, we show that the application of an LSTM can improve inference in empirical settings beyond what typical time series models can provide. In particular, we substantiate that it is the innovations (surprises) in linkages, not just the changes, that conveys information.

In addition, we make specific contributions in understanding banks' systemic risk and interbank network structure effects. Regarding systemic risk, Caballero and Simsek (2013) have emphasized that banks' knowledge of other banks' financial health is limited by the extent of their counterparty relationships. They reason that during times of systemic stress, when the health of counterparties' counterparties (and counterparties of those counterparties, etc.) are of increased concern, banks may retrench into a liquidity conservation mode. Our evidence suggests that banks sometimes engage in more extensive counterparty sampling during times of stress, which may broaden their information set and could partly offset the need for retrenchment. Regarding interbank network structure, our evidence suggests that interbank lending relationships matter for informational reasons, not only because of domino-theory concerns about interlinked liabilities. Our evidence complements the recent model of Denbee et al. (2021), which is premised on the idea that the liabilities topology of the interbank network is less important than banks' interpretations of their peers' actions. Denbee et al. (2021) show that the key to the systemic demand for liquidity is whether banks consider each other's reserves to be substitutes (i.e. reservoirs of readily-borrowable liquidity) or compliments (i.e. indications of the value of retaining liquidity on their own accounts). We show that changes in banks' links to other banks, which can bring in new information on which to base their views, are predictive of rates in the markets for liquidity. Thus, we show a price effect of linkages that complements their quantity effect.

The next section of our paper overviews aspects of the literature on interbank networks as necessary context. The subsequent two sections then describe our data and empirical methods, respectively. A fifth section presents our results on network changes in uncertain times, and a sixth section reports on the predictive power of the counterparty sampling changes for interest rates during spates of uncertainty. A seventh section considers and rejects some non-learning explanations for the importance of counterparty sampling. A final section concludes.

2. Interbank networks

The importance of bank interconnectedness for systemic risk, i.e. the risk of a collapse of the banking system, is a foundational concept relating to financial stability and financial crises (Allen and Gale 2000; Diamond and Dybvig 1983; Eisenberg and Noe 2001). Banking systems generally include some mechanisms, such as overnight funds trading networks, intended to help the system self-insure against localized liquidity shortfalls that could otherwise cause instability or force "fire sales" of bank assets to raise funds (Greenwood et al. 2015; Shleifer and Vishny 2011). The actual cross-bank liabilities can be highly variable, depending on which banks need to call on the insurance, even though the structure of the system itself may be stable. Glasserman and Young (2016) provide a comprehensive survey from this interbank-liability perspective.

Interbank networks have also been usefully analyzed from a trading perspective. A relevant example is Brunetti et al. (2019), who empirically build the network of interbank cross-liabilities from bids, asks, and trades for the same e-

MID market that we study. They find that this “physical network” deteriorated and fragmented as the Global Financial Crisis proceeded. Afonso et al. (2014) also study an interbank funding network during the Global Financial Crisis, finding that borrower characteristics became more important for lending decisions. Ashcraft and Duffie (2007) find that the probability of two banks trading is increasing in the strength of their prior relationship. Hatzopoulos et al. (2015) document preferential trading partners over time in the same market we investigate. Though a focus on actual or potential crises clearly pervades these literatures, they also provide insights into the functioning of interbank networks at times of stress and uncertainty that fall short of being cataclysmic.

Finally, interbank networks have been studied from a network topology perspective. Researchers often document a core-periphery or tiered network topology.[6] Craig and von Peter (2014) point out that any non-random structure to interbank networks suggests a purpose in addition to merely diversifying liquidity shocks. Blasques et al. (2018) point to easing the costs of credit risk monitoring as one reason bank networks might not be random, and suggest a model in which counterparty monitoring helps generate a core-periphery topology. Also in this vein, Erol and Ordoñez (2017) emphasize that banks make relationships to assure themselves access to reserves to support profitable lending on client projects. When circumstances limit the profitability of such lending, the benefits of relationships are likewise limited. Beyond a tipping point where any bank finds that many other banks have decided to withdraw, the network can become very sparse. Their model suggests that banks will form isolated “cliques” to obtain the benefits of support for reserves without taking more-than-necessary risk of being harmed by the shocks of other banks. Erol and Ordoñez (2017) thus provide a basis for examining the possibility of cliques and other persistent counterparty groupings, and for considering the possibility of quick changes in their compositions. We emphasize the importance of cliques from an informational perspective.

Taking together these three research perspectives on interbank networks—interlinked liabilities, the trading, and the network topology—we draw several points to motivate and guide our work. First, network linkages can be highly dynamic, even if the network structure is stable. Second, linkages, structure, or both might change in times of stress or crisis. Third, important topological aspects of interbank networks driven by banks’ underlying economics, such as lending cliques, may deserve a special focus.

3. Data

Our data detail trading on the e-MID platform (formally, Mercato Interbancario dei Depositi), and have been obtained directly from e-MID. e-MID was a quote-driven screen-based platform for interbank deposit trading by European banks. Quotes (i.e. binding bids and offers for bank deposits) were visible to all participating banks, and trades against those quotes could be executed electronically. Trading was available in several currencies, though trading in euros was the most common. Deposits were lent and borrowed on an unsecured basis. Various loan maturities up to a year were available, but the great majority of trades were for overnight deposits. We focus on euro-denominated overnight borrowing and lending.

We begin with all trades where a lending bank accepts a borrowing bank’s bid for an overnight deposit. The data identify the lending and borrowing banks with unique reference numbers that anonymize each bank’s identity except for its country. Key fields include the timestamp and size of the loan in euros. From this trade record, we first construct a daily dataset with observations on the total of loans each day from each lender to each borrower and then identify bank-pair linkages for each week according to whether one bank lent to the other at least once.

e-MID was headquartered in Italy, and Italian banks represented the largest cohort of banks in the system. However, many banks from Germany, France, and other European countries were also active traders. Given our purposes as described in the introduction, we key on Italian borrowers and non-Italian lenders. These non-Italian lenders include the large money-center banks that sat on panels to determine LIBOR interest rates, which were the benchmark rates for most bank loans to customers.

During the early 2000s, interbank trading on the e-MID platform was very active and was regarded as typical of deposit trading by European banks overall (Finger et al. 2013). Aciero et al. (2016) report that e-MID trading represented 20% of all interbank transactions in Europe as late as 2007. Monticini and Ravazzolo (2011) state that e-MID was still representative of Euro area interbank trading overall even in 2008. After 2008, trading on e-MID declined and became more dominated by Italian banks. The e-MID market eventually closed after 2011.

Before 2004, the ECB’s operational framework for providing liquidity was different than afterward in several salient ways that would affect comparability.[7] After 2006, the Italian Ministry of Finance undertook a new practice of concurrent money market operations to improve the predictability of the government account, with the result that tax-payment-associated systemic liquidity uncertainty that is at the heart of our empirical design became somewhat curtailed. Further, the Global Financial Crisis, which coincided with the various LIBOR rate setting scandals that eventually led to LIBOR’s demise, began to affect Europe beginning in summer 2007.

We study data from 2005 to 2006. All things considered, 2005 and 2006 is the right time period for our investigation. We focus on data after 2004 to avoid any lack of comparability within our samples due to the changes in the ECB operations starting in 2004. We focus on data before 2007 so that we have an active market rather than eMID in decline, when it may be less representative of the full market for overnight funds, and so that we have a market during normal times rather than a crisis-time market.

Our investigation is facilitated by the temporal pattern of ECB system-wide liquidity management in the early 2000s. Banks' reserves were assessed for regulatory purposes over rolling "reserve maintenance periods." A reserve maintenance period is the timespan over which banks must hit targeted reserve levels. The ECB's calendar of reserve maintenance periods, tied to the schedule of ECB Governing Council meetings, is published well in advance. During the period of our data, reserve maintenance periods ranged in length from 20 to 43 days, with most periods being about a month long. The rough correspondence between maintenance periods and calendar months allows us to conveniently refer to specific maintenance periods according to the calendar month within which they begin.

To have comparisons that are robust to the normal influence of ECB liquidity/regulatory architecture on banks' reserves-management incentives, we organize our sample as a set of consecutive maintenance periods, each with a tax-transmittal period and a pre-tax-transmittal period included. Beaupain and Durré (2008) document platform-wide trading regularities corresponding to the start of a period (light activity and rate volatility), end of a period (heavy trading and rate volatility), as well as to the roughly-weekly occurrences of ECB MROs.

In the introduction, we have already described the rationale and institutional setup of the Italian tax-transmittal sub-period (TTP) and a pre-tax-transmittal sub-period (PTTP). For our investigation, we designate TTPs and PTTPs to be 5 business days centered around the tax-transmittal date and the same day of the previous week, respectively. Each maintenance reserve period contains one PTTP and one TTP, with the TTP coming close to the end of the maintenance reserve period. Given that the Italian tax calendar is an autonomous factor driving liquidity risk in the eurozone according to the ECB, we consider PTTPs as likely higher uncertainty times, and TTPs as times when uncertainty is resolved. Later we provide evidence this is indeed the case.

4. Empirical Design

Our empirical design has two main components. In the first component, we investigate the effects on network topology and lending relationship patterns when a tax-transfer-period spate of uncertainty occurs. In the second component, we examine the time-series relations of topology and lending relationships with the future realizations of key interest rates and spreads of concern to banks. We next describe, in turn, the set-up of each of these two components of our empirical plan.

4.1. First Empirical Design Component: Effects of Uncertainty-Increasing Cases on Network Topology and Lending Relationships

Our first empirical design component corresponds to our purpose to characterize the reaction of lenders in the interbank network to increased uncertainty and liquidity stress. We identify cases of repeated increased uncertainty with a predictable timing, with corresponding control-group cases for each that are otherwise similar but removed from the uncertainty. The control cases are the key to holding constant the effects of the ECB maintenance period and refinance operation regime, as well as financial and economic conditions more generally.

The uncertainty *vs.* control case split is based on the Italian tax collection cycle. As discussed above, the monthly process of transferring ongoing tax payments in the government's accounts in Italy drains a significant but highly uncertain amount of liquidity from banks. The ECB notes this as having eurozone-wide significance, not only for Italy. The transfers occur around the 23rd of each month. The uncertainty is resolved as the transfers occur and the amount of liquidity drained becomes known to system participants. Our control period is thus defined as the 5 business-day period beginning on the tax transfer date (the 23rd or the first Italian business day after the 23rd), and the uncertainty period is defined as the 5 business-day period before the tax transfer date. For nomenclature, we call these the uncertainty case as the "pre-tax-transfer sub-period" (PTTP for short) and the control case as the "tax-transfer sub-period" (TTP for short), respectively.[8] There is one PTTP and one TTP within each monthly ECB reserve maintenance period. Therefore, our study effectively focuses on the 24 ECB reserve maintenance periods during 2005 and 2006.

This uncertainty case *vs.* control case split can also be motivated statistically. The mini-table below shows the rolling mean levels of 5-day variance for several money market interest rates and spreads with strong relevance to banks during each of these case periods, with t-statistics for the difference in parentheses.[9] The sample for the mini-table includes all the tax-transfer periods and pre-tax-transfer periods in 2005 and 2006. The variance in these overnight rates is sharply greater during the PTTP uncertainty cases than during the TTP control/uncertainty-resolution cases. At a minimum, the multiple is about four times.

We measure the effects of uncertainty increases on two aspects of the interbank lending network. The first is network topology, using several standard metrics. The other is the tendency of lenders to broaden or narrow their group of borrowers, using several different metrics. For nomenclature, we refer to a lender's group of borrowers as its "lending clique." These measures are described below.

Tax-transfer-based cases			
Weekly variance in	TTP control case		PTTP uncertainty case
EONIA rate	0.0055		0.0228
t-statistic for difference		(3.63)	
Euro LIBOR overnight rate	0.0048		0.0182
t-statistic for difference		(3.49)	
Euro LIBOR 3-month spread	0.0053		0.0200
t-statistic for difference		(3.46)	

To assess network topology, we measure network density, clustering/transitivity, in-degree centralization, and out-degree centralization. These are standard descriptive measures of network topology and have specific economic importance in our context. Brunetti et al. (2019) emphasize the importance of network topology. Among other things, they consider especially density and transitivity/clustering as being important. Denbee et al. (2021), in their theoretical model of interbank liquidity linkages, propose that a bank's centrality, i.e. the extent to which it is linked to the network of other banks by either borrowing (for the in-degree case) or lending (for the out-degree case) is salient for its contribution to systemic risk. Centralization measures, as network-wide aggregates of bank-specific centrality measures, provide a summary of the tendency of the network to have banks that are systemically important. Below we provide some detail on the calculation of these specific network topology measures.

- Bank network density is computed as the number of banks with a lending link between them relative to the maximum possible number of links for a network with the same number of nodes.
- A bank's transitivity (also known as clustering) is computed as ratio of the count of closed three-link lending triads in the network (i.e. where A lends to B, B lends to C, and C lends to A) relative to the number of two-link lending relationships possible in the network (A lends to B, B lends to C). It is the ratio of triads that are closed relative to the number that could be closed. Since the most connected network is one where every triad is fully linked, transitivity is a simple proportional measure as to how far the actual network rises toward this potential level of connectedness. In terms of the theoretical discussion of networking above, a network with high transitivity would have a large number of strong ties.
- A bank's out-degree centrality is computed as $C_{out}(i) = \Sigma_{j=1}^{N} y_{ij}$ where y_{ij} is an indicator variable equal to one if bank i lends to bank j and zero otherwise. Similarly, its in-degree centrality is computed as $C_{in}(i) = \Sigma_{j=1}^{N} y_{ji}$ where y_{ji} is an indicator variable equal to one when bank j lends to bank i and zero otherwise. Then the (Freeman) degree centralization measures are computed as the average deviation of each node's degree centrality from the most-central node's centrality (i.e. a measure of the variation in centrality across the network), standardized as a proportion of the summed deviations for a hypothetical network with the largest possible summed deviations for a network with the same number of nodes. The theoretically most centralized network (a star network with each node connected only to a single central node) has the largest average variation in centrality because it has only one highly linked node connecting all others, so the relative average variation is then a proportional measure as to how far the actual network rises toward this potential level of centralization.

As Pfaltz (2012) notes, a network that adds links that bridge between disparate groups of network banks may experience only very minor changes in the standard network topology statistics, such as those laid out above. Therefore, to specifically assess the tendency of lenders to broaden or narrow their lending cliques, we use several specific metrics as follows.

- First, we assess the proportion of network banks' lending relationships that are with banks they lent to during the previous tax-transfer period (TTP) or pre-tax-transfer period (PTTP), respectively. If banks tend to bring back relationships from a month earlier similar-uncertainty situation, it is an indication of restoring lending relationships that might convey information useful during such times. This contrasts with cases where banks might simply continue to trade with the counterparties they have been using most recently. We call this metric the proportion of restorations of previous trading relationships.
- Second, we count the number of borrowers in each network bank's lending clique during a particular TTP or PTTP. Then, for the network as a whole, we compute the mean size of lending cliques and the cross-sectional standard deviation of lending cliques. The mean is a measure of the extent of an average lending banks' relationships. The standard deviation impounds this information also, in that the minimum size of a clique is zero. Additionally, the standard deviation impounds information about the variety of clique sizes in the system.
- Finally, we use Shannon's entropy, a summary statistic from network science that measures the overall level of disorder in the network. For a network with N banks, entropy is defined as $H = -\sum_{i=1}^{N-1} p_i \ln(p_i)$, where $p_i = k_i \times \frac{1}{N-1}$ with k_i being the number of banks that bank i lends to. Networks with larger entropy are more disorderly, i.e. have a stronger tendency for any two banks to be connected. When banks spread their lending across larger cliques, entropy is apt to increase. One reason for including entropy is that it has prominence in information theory as a summary measure of the information in a system.

Each of these three measures is a different summary of the overall breadth of relationships in the interbank lending network. We will therefore refer to them as network breadth measures. Foreshadowing our finding that broader networks convey information in the manner of drawing larger data samples, we also call them "counterparty sampling" measures.

To use the standard topology and the counterparty sampling measures for our purpose, we assemble interbank lending data for a particular time period (such as the TTP or PTTP sub-period within a reserve maintenance period) into a set of directional "links" corresponding to lender-borrower relationships between bank pairs. We begin at a daily level, totaling up the lending from each bank i to each other bank j from the e-MID transaction record. To focus on lending decisions, we consider only cases where a lending bank lifts a bid for an overnight deposit from a borrowing bank, as discussed above. Next we aggregate the pair-specific lending over the entire time period. Whenever the aggregate lending from i to j during the period is positive, we record a link in the interbank network "adjacency matrix" showing connections of all possible bank pairs i, j. Adjacency matrices then become the input data for computing the network topology and counterparty sampling measures above.

In addition to computing aggregate lending considering all loans from i to j, we sometimes separately compute aggregate lending using only "small" loans, where the dividing line between large and small is set on a lender specific basis. That is, for each lending bank i we examine its full distribution of loans during the maintenance period and designate as "small" all loans below the 25th percentile in size by euros lent. The point of measuring small loan networks is to focus on connections for which credit risk considerations are less important.

In addition to connectedness, the size of the interbank network should be expected to be economically important. Therefore, as control variables for our study we use some standard measures as follow:

- *Total euros lent*, defined as the number of euros lent by all banks in aggregate on a day,
- *Number of distinct lenders*, defined as the number of banks that lend on a day, and
- *Number of distinct borrowers*, defined as the number of banks that borrow on a day.

The various measures above are applied to the full network of lenders, the network of non-Italian banks lending to Italian banks, and, for comparison, the network of Italian banks lending to non-Italian banks.

4.2. Second Empirical Design Component: Predictive Relationships between Changes in Lending Relationships and Interest Rates

We also undertake a second analysis to assess the predictive relation between network patterns and the 3-month LIBOR interest rate. This rate is of key importance a benchmark interest rate for bank loans, it is determined by a daily survey of large non-Italian banks in the e-MID network, and it is not directly determined by trading in the e-MID network. For all these reasons, 3-month LIBOR is a good focus to test if the extent and variety of e-MID lending relationships provide useful information to lenders about the future state of the banking system and the economy.

We use Granger causality tests to assess the predictive power of our counterparty sampling measures to predict these rates on a day-ahead basis. In general, Granger causality tests are χ^2 tests of predictive effects based on estimates of linear vector autoregression systems (VARs). We estimate 5-equation VARs for the system of the change in 3-month LIBOR, one of the network connectedness measures (i.e. either the proportion of restorations, the standard deviation of clique size, or the entropy), the number of distinct borrowers, the number of distinct lenders, and the number of euros lent in aggregate. We also include day-of-the-week indicator variables to allow for weekly seasonal effects. Importantly for our purpose, we often allow for an additional equation for the interaction of a network connectedness measure and an indicator variable for the PTTP uncertainty-case period.

Our basic implementation of this VAR system uses five day lags of all the endogenous variables since our counterparty sampling measures are computed on a weekly basis and so do not change from day to day. In the empirical section later, these results are shown in Table 4.[10] To check the results' robustness, we also use a one-week-lag setup, measuring weekly changes in LIBOR. We show the results for weekly changes in Table 6.

With an included lag at only one horizon, Granger-causality can be assessed *via* the t-statistic on the relevant coefficient. The most important equation in our VARs is the one for changes in 3-month LIBOR, and our reporting focuses on this one. We measure daily changes in 3-month LIBOR in order to avoid overlapping data. As noted, we also estimate VARs where the change in 3-month LIBOR is measured as a weekly change, but in that case use only data for a particular day of the week. One important case regards Thursday data, for this is the day when the ECB Governing Council holds its monetary policy meetings and press conferences.

As with the first design element, we carry out these investigations based on, alternatively, the full network of lending banks and the sub-network of non-Italian banks that lend to Italian banks. Our hypothesis about bridging links is reflected most directly in measures regarding the non-Italian bank sub-network, but if the effects are pervasive we might also expect to see similar evidence for the full network.

5. Findings regarding the Effect of Uncertainty on Network Topology and Lending Relationships

5.1. Descriptive Example of the e-MID Interbank Network

To begin, it is helpful to have a sense of the general arrangement of the interbank network and to visualize some of the points we are making about counterparty sampling during lower-uncertainty *vs.* higher-uncertainty times.

Figure 1 presents a graphical depiction of the interbank network during the February and March 2006 reserve maintenance periods. The figure is based on an adjacency matrix constructed separately for each maintenance period that includes all banks that lend by lifting borrowers' quotes.[11] The top panel of the figure shows the network during TTP control-case periods, and the bottom panel shows the network during PTTP uncertainty-case periods.

Network nodes represent individual banks in the network, and are arrayed around a circle or clockface for visual convenience. The darker green colored nodes, from about 11 o'clock on the clockface to about 4 o'clock, represent non-Italian banks from across Europe. The lighter colored nodes represent Italian banks.

Link-lines connecting nodes represent that at least one loan was made between that pair of banks during the maintenance period. Red links represent relationships that are present only during February and green links represent

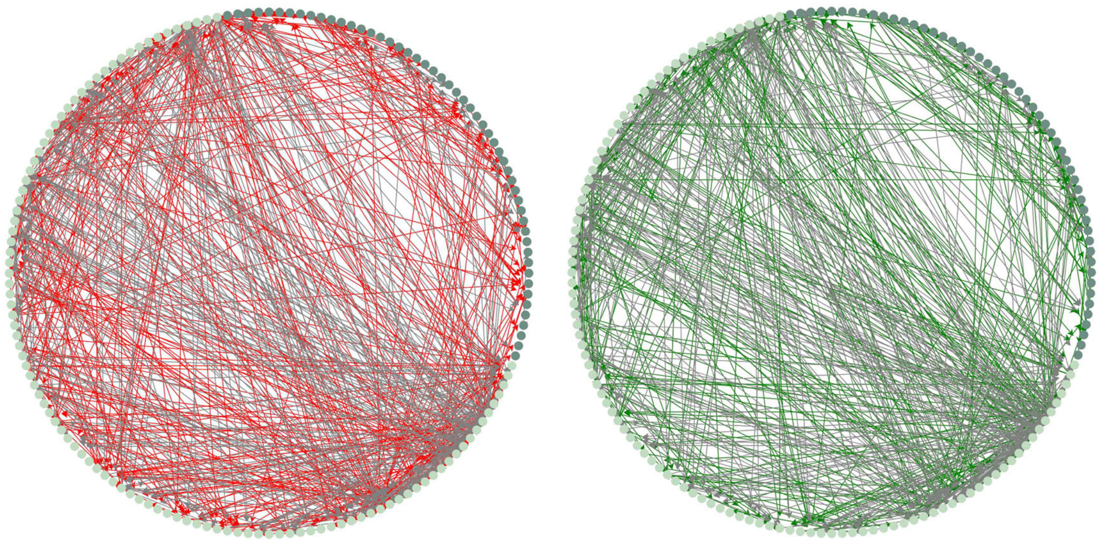

(a) February-March 2006 TTP network breakout, with restoration (gray), stopping (red) and starting (green) links.

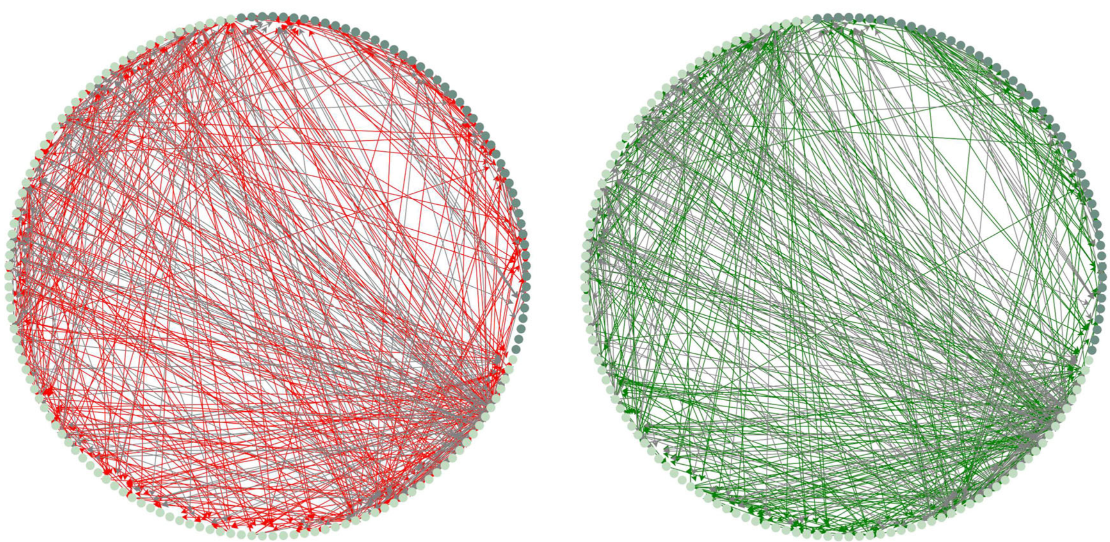

(b) February-March 2006 PTTP network breakout, with restoration (gray), stopping (red) and starting (green) links.

Figure 1. February–March 2006 tax-transfer period (TTP) and February–March 2006 pre-tax-transfer period (PTTP) interbank lending networks, with restoration, stopping, and starting links. Networks are based on all loans made when lending banks accept e-MID borrowing request quotes. Gray edges indicate links that exist in the TTP or PTTP case, respectively, of the February reserve maintenance period and continue to exist in the March reserve maintenance period's TTP or PTTP case, respectively. Red edges indicate lending links that exist in the respective February case but do not exist in the respective March case ("stopping links"). Green edges indicate lending links that begin in the respective March case but did not exist in the respective February case ("starting links"). Edges begin at the lender node and are terminated with a directional arrow short of reaching the borrower node so that lending cliques can be visualized. Node shading indicates bank nationality, with Italian banks shaded lighter, and other nationalities' banks shaded darker.

relationships that are present only during March. Grey links represent relationships that are present during both months.

It is evident in the figure that the interbank networks are similarly dense for both maintenance periods (i.e. comparing left and right in the figure). The network is also highly interlinked during both TTP and PTTP periods (i.e. comparing top to bottom in the figure), though it appears a little less-dense in the PTTP periods each month. Lending links involving both Italian and non-Italian banks are commonplace, though lending links within geographic blocks are more common, as indicated by the less-dense more-uncolored swaths toward the top right of each clockface.

The overall pattern is suggestive of both "strong" links that densely interconnect local nodes and "weak" links that bridge to distant parts of the network as in Granovetter (1973). The pattern is also a first visual indication that entropy may be greater within the non-Italian and Italian clusters, respectively, than in the network overall, in that the full network has additional structure (which would reduce entropy by definition).

Next, we discuss implications suggested by the coloring of the links, i.e. ones that are unique to one month or are present across the months. From that perspective, repeat links connecting a particular type of sub-period (PTTP or TTP) across time for adjacent maintenance periods are "restoration links," i.e. evidence of relationships that have been restored from a past time of either higher (PTTP) or lower (TTP) uncertainty that have a time gap between them. Lending banks could be expected to return to borrowers that they have found useful given the type of market conditions. We want to investigate the possibility that the usefulness includes information flow about the uncertainty situation. With this in mind, the count of restoration links is one measure we use to assess the extent to which banks keep past condition-specific information flow possibilities open.

Taking the March network as an example, while there is a slightly larger proportion of restoration links in the full network of all lenders in the PTTP sub-period than in the TTP sub-period (28 *vs.* 27%), restoration links are less common than starting links or stopping links in both sub-periods. The comparison of PTTP restoration links to TTP restoration links is much more lopsided when the network of non-Italian lenders to Italian banks is considered (15 *vs.* 7%). We are not arguing that these differences are necessarily representative of the network in general over time, but only explicating the counterparty-sampling measurement approach and showing with an example how counterparty patterns can vary across the higher- and lower-uncertainty cases. Later, we will consider whether they differ on average and whether it makes a difference for future interest rates how they differ.

Another perspective on counterparty sampling can be taken by considering the breadth, variation, and changes in banks' lending cliques. Table 1 provides a numerical accounting of lending relationships for example of the February 2006 network. The table shows the distributions of the number of unique counterparties lent to by lending banks as a numerical histogram where buckets are formed according to the number of unique borrowers for each given lending bank. Panel A shows the control-group TTP case, Panel B shows the uncertainty-increased PTTP case, and Panel C shows the differences. During the PTTP sub-period, more banks choose to do some lending, fewer banks choose to lend to only a small number of counterparties, and more banks choose to lend to a medium size number of counterparties. About the same proportion of banks choose to lend to a large number of counterparties. The indication is that the weight of the distribution has shifted toward larger lending cliques for the PTTP sub-period case, which would be consistent with more counterparty sampling in this example case. In later sections of the paper when we need summary indications of the extent of such tendencies we will use the median size of lending cliques and the cross-sectional standard deviation. Given the lower bound of zero on lending clique size, we lean toward using the cross-sectional standard deviation because it is sensitive to both the tendency for large cliques and the tendency for different banks to have different clique choices.

The link-count and clique-size examples discussed above focus on how the patterns of counterparty relationships change differently over time for higher *vs.* lower uncertainty cases. With this in mind, Table 2 reports some standard topology characteristics of the TTP-PTTP comparison using the February 2006 maintenance period as an example. Panel A is constructed for the full network using all loans where a lender lifts a quote from a borrower. Panel B uses only lender-specific small loans, i.e. those in the lower quartile of loan size for the specific lender. The purpose of including Panel B is to move away from situations where credit-risk considerations might be a key driver of interbank lending choices toward measuring the pattern of lending relationships where a possibly-informational connection exists but a counterparty default would not be devastating to the lender. The panels give similar impressions overall, suggesting that the link structure of the network is not driven only by credit risk considerations. For this example of a PTTP uncertainty

Table 1. Percentage distribution of lending clique sizes.

Number of unique banks lent by a lending bank	(1) % Of all network banks
Panel A: TTP	
0 "Non-lenders"	49.1
1–5 "Small clique lenders"	20.8
6–20 "Medium clique lenders"	13.8
21–90 "Large clique lenders"	16.4
Panel B: PTTP	
0 "Non-lenders"	35.0
1–5 "Small clique lenders"	22.3
6–20 "Medium clique lenders"	27.3
21–90 "Large clique lenders"	15.3
Panel C: Differences	
0 "Non-lenders"	−14.1
1–5 "Small clique lenders"	1.50
6–20 "Medium clique lenders"	13.5
21–90 "Large clique lenders"	−1.1

A lending clique is defined as the set of borrowing banks that a particular lending bank provides with overnight deposits during a particular time period by lifting a quotes requesting a deposit on the e-MID system during a particular time period. Panel A shows the distribution of lending clique sizes during the tax-transfer period (TTP) in the February 2006 ECB reserve maintenance period. Panel B shows the distribution during the pre-tax-transfer (PTTP) period of the same maintenance period. Panel C shows the differences between those two distributions. For this table, the network for all banks refers to the set of banks that lend at least once during the February or March 2006 reserve maintenance periods that have been the subject of our introductory examples.

Table 2. Comparison of lending network descriptive characteristics and their changes across tax-transfer period (TTP) and pre-tax-transfer period (PTTP) for February 2006.

Characteristics	(1) TTP	(2) PTTP	(3) Difference
Panel A: Full network			
Number of network banks	159	157	−2
Number of lending relationships	880	644	−236
Density	0.035	0.026	−0.009
In-degree centralization	0.175	0.154	−0.021
Out-degree centralization	0.353	0.238	−0.115
Transitivity	0.225	0.124	−0.101
Panel B: Small loan network			
Number of network banks	155	151	−4
Number of lending relationships	406	263	−143
Density	0.017	0.012	−0.005
In-degree centralization	0.087	0.055	−0.032
Out-degree centralization	0.238	0.170	−0.068
Transitivity	0.148	0.088	−0.060

Panel A describes the lending network for all loans made when lending banks accept e-MID borrowing request quotes. Panel B describes the lending network for normal-size loans, and Panel B describes the network for small loans, defined as those in the lowest quartile of loan size for each specific lending bank.

Table 3. Selected characteristics of interbank lending networks across months for tax-transfer sub-periods (TTP) and pre-tax-transfer sub-periods (PTTP).

	(1) TTP	(2) PTTP	(3) Difference	(4) *z*-statistic
Panel A: Standard network topography statistics				
Sub-panel A.1: Network of non-Italian lenders to Italian borrowers				
Out-degree centralization	0.103	0.132	0.029	1.660
In-degree centralization	0.203	0.215	0.002	0.426
Sub-panel A.2: Network of Italian lenders to non-Italian borrowers				
Out-degree centralization	0.194	0.190	−0.004	0.254
In-degree centralization	0.173	0.156	−0.017	1.077
Sub-panel A.3: Full network of all lenders to all borrowers				
Out-degree centralization	0.307	0.316	0.009	0.447
In-degree centralization	0.169	0.170	0.001	0.180
Transitivity	0.236	0.231	0.006	0.274
Panel B: Information-flow-oriented statistics				
Sub-panel B.1: Network of Italian lenders to non-Italian borrowers				
Average clique size	4.984	3.756	−1.228	−4.491
Standard deviation of clique size	3.161	2.451	0.710	−2.322
Proportion of restorations	0.208	0.164	−0.044	−3.662
Proportion of small restorations	0.180	0.128	−0.052	−3.156
Proportion of small restorations	0.180	0.128	−0.052	−3.156
Entropy	1.065	0.901	−0.164	−2.753
Sub-panel B.2: Network of non-Italian lenders to Italian borrowers				
Average clique size	3.876	5.369	1.493	2.570
Standard deviation of clique size	1.968	3.044	1.076	2.627
Proportion of restorations	0.075	0.134	0.059	3.383
Proportion of small restorations	0.069	0.107	0.308	1.961
Entropy	0.574	0.840	0.266	2.821
Sub-panel B.3: Full network of all lenders to all borrowers				
Average clique size	11.919	11.799	−0.120	0.407
Standard deviation of clique size	7.612	7.414	−0.198	−0.809
Proportion of restorations	0.258	0.267	0.009	1.190
Proportion of small restorations	0.171	0.167	−0.004	−0.529
Entropy	1.300	1.334	0.034	1.266

The table shows averages of select network characteristics over all ECB reserve maintenance periods in 2005 and 2006 according to TTP and PTTP designations. Panel A shows standard network topography statistics. Panel B shows information-flow-oriented statistics as developed in the test. In each panel, statistics are shown for the network of non-Italian lender banks to Italian borrower banks, for the network of Italian lender banks to non-Italian borrower banks, and for the full network, respectively.

case as compared to a TTP control case, the number of links, density, in- and out-degree centralization, and transitivity are all reduced. Based on these topological measures, the network appears to move somewhat toward more disorder, i.e. a state more like a random network. Especially meaningful may be the drop in transitivity, indicating relatively fewer closed triads and thus possibly more bridging links that might be important for information flow.

We now proceed to a systematic analysis of data for PTTP and TTP effects across all the 2005 and 2006 reserve maintenance periods.

5.2. Results regarding Effects of Uncertainty Cases on Network Topology and Relationships

In this subsection, we extend the investigation of network topology and relationships to provide more systematic inference, following the empirical design developed above. Standard network topology measures may be at least somewhat sensitive to the issues we are investigating and contribute to understanding the e-MID network in the context of the literature. And the counterparty sampling measures should focus even more on the issues at hand.

With this in mind, Table 3 provides both topology and counterparty sampling statistics summarized separately for the TTP and PTTP sub-period cases during the 2005–2006 reserve maintenance periods. The table is constructed by first measuring each respective statistic (whether a standard topology statistic or a counterparty sampling statistic) for the TTP and the PTTP within each of the 24 reserve maintenance periods in our sample, and then constructing averages, differences in averages, and test statistics across all 24 observations. Within each panel are sub-panels for the sub-network of non-Italian banks lending to Italian banks, the sub-network of Italian banks lending to non-Italian banks, and the full network. Panel A presents standard topology statistics, and Panel B presents counterparty-sampling statistics.

Examination of the standard network topology statistics in Panel A shows no statistically-significant differences in out-degree centralization, in-degree centralization, or transitivity. This is true for the two sub-networks and the full network, i.e. in all three sub-panels. Evidently, the topology of the e-MID interbank network is similar on average across uncertainty *vs.* control cases. In keeping with the points made in Pfaltz (2012), any change in the networks tendency to link in general or to link across geographical regions is not reflected in the overall standard topology measures. If there is a change, more focused measures are needed to illuminate it.

Panel B presents the more information-flow-oriented counterparty sampling statistics that we have developed, for the two sub-networks and the full network. During PTTP sub-periods as compared to TTP sub-periods the network of non-Italian lenders making loans to Italian banks (Sub-panel B.1) shifts on average to larger cliques, more variation in cliques, more restorations, more restorations considering only small loans, and more entropy. All differences are statistically significant and seem large enough to be economically meaningful. Overall, more connections, more variety in connections, and more disorder are evident.

In contrast, the network of Italian banks lending to non-Italian banks (Sub-panel B.2) exhibits changes in exactly the opposite direction comparing across TTP *vs.* PTTP cases. All of the counterparty sampling statistics are smaller on average in the PTTP uncertainty cases (Sub-panel B.3). All the differences are statistically significant, though not as economically large as the ones in Sub-panel B.1. The spates of uncertainty represented by the PTTP cases are rooted in the Italian bank sub-network. The statistics may therefore indicate that Italian lenders have no need to gather information by establishing broader, more disorderly sets of relationships across the gap to continental Europe. Establishing such sets of relationships might be more costly during a time of stress, and so are avoided absent the need.

For the full network (Sub-panel B.3), there are no statistically significant differences comparing across TTP and PTTP cases, neither in mean clique size, standard deviation of clique size, proportion of restorations, proportion of restorations for lender-specific small loans, or entropy. This is evidence that the changes discussed above are rooted specifically in connections across the regions.

The counterparty sampling statistics each measure a different mechanical aspect of the network, but they all give the same statistically-significant message. Non-Italian banks are reaching more broadly to lend in Italy. It is as if the banks are sampling relationships in Italy, perhaps hoping to learn. This fits with the idea of bridging ties being more strongly present during the high-uncertainty PTTP cases: non-Italian lenders are reaching across the gap into the set of Italian borrowers with lending relationships. If these relationships open the door to information flow, we might expect to see the tracks in further analysis.

In this section, we have assessed the on-average differences in the interbank network across the lower- and higher-uncertainty cases. The extent of the differences does vary from one reserve maintenance period to another, which opens the door to analysis based on the extent of the variation. In the next section, we provide such an analysis with a focus on how the network variations may relate to future interest rates.

6. Results regarding Predictive Relationships between Changes in Lending Relationships and Interest Rates

The previous section shows that, when volatility spikes due to liquidity shocks building in one part of the banking sector, there are some on-average reactions in the linkages to distant sections of the interbank market that are reflected in counterparty sampling measures. To see if these reactions contain information that is relevant for the future path of interest rates, we next exploit the variation in those reactions across the reserve maintenance periods in our sample. In this section, we report the results of time series tests that connect future interest rates with interbank network trading relationship data over all days during 2005–2006, relying on the empirical design outlined earlier.

6.1. Predicting Week-Ahead Daily Changes

Table 4 reports on a single equation, the one for the 3-month LIBOR changes equation, from vector-autoregression systems (VAR) involving for the sub-network of non-Italian banks that lend to Italian banks. We focus on changes in 3-month LIBOR changes because an augmented Dickey-Fuller test (not reported) indicates that the raw 3-month LIBOR time series may be non-stationary.

The VARs reported in columns (1), (3), and (5) are constructed exactly as in Equation (1), and each includes a single counterparty sampling measure in the system: either the proportion of restorations, the standard deviation of clique size, or entropy. In every case, the coefficient linking the lagged counterparty sampling measure to the change in 3-month LIBOR is statistically insignificant. There is no evidence of Granger causality. Counterparty sampling apparently does not influence 3-month LIBOR changes unconditionally.

The VARs reported in columns (2), (4), and (6) are respectively analogous to those in columns (1), (3), and (5), and also report on the 3-month LIBOR change equation with each of the three different counterparty sampling effects. These VARs are also constructed as in Equation (1) except that they additionally include an interaction term for the counterparty sampling measure and a PTTP indicator variable, as discussed in the Empirical Design section above. The estimated coefficients on the interaction term are uniformly positive and statistically-significant, providing positive evidence of Granger causality due to counterparty sampling changes during these periods. The indication is that measures relating to more counterparty sampling and network disorder in the bridging links between non-Italian bank lenders and Italian bank borrowers are positively related to upcoming changes in LIBOR a week ahead, but only during PTTP uncertainty-case time periods.

Each of the counterparty sampling measures is different in the way it captures network disorder and broader linkages with banks in the part of the network that is affected by the liquidity shocks. Yet they are moderately correlated (not reported in the table), indicating that they may capture similar underlying economic regularities. For this reason, as well as due to the limited number of daily observations in our 2 year time period, if we use all three measures in the same VAR system, the interaction terms are statistically insignificant. Even so, an F-test (not reported in the table) rejects

Table 4. Non-Italian lender bank network vector-autoregressions showing the influence of information-flow-oriented network statistics on one-week-ahead daily changes in 3-month euro LIBOR.

Regressors	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Proportion of restorations	−0.001 (−0.04)	−0.002 (−0.45)						
Proportion of restorations × PTTP dummy		0.013** (2.28)						
Standard deviation of clique size			−0.001* (−1.77)	−0.001** (−2.13)				
Standard deviation of clique size × PTTP dummy				0.001* (1.85)				
Entropy					−0.001 (−1.42)	−0.002* (−1.84)		
Entropy × PTTP dummy						0.002** (2.02)		
Summary counterparty sampling measure							−0.001 (−1.35)	−0.001 (−1.61)
Summary counterparty sampling measure × PTTP dummy								0.001* (1.91)
Current change in 3-month LIBOR	0.082* (1.84)	0.084* (1.89)	0.073 (1.63)	0.074* (1.67)	0.078* (1.75)	0.077* (1.75)	0.076* (1.71)	0.77* (1.73)
Number of distinct lenders	−0.001*** (−2.42)	−0.001*** (−2.35)	−0.001*** (−2.49)	−0.001*** (−2.47)	−0.001*** (−2.40)	−0.001*** (−2.36)	−0.001*** (−2.40)	−0.001*** (−2.37)
Number of distinct borrowers	0.001*** (2.59)	0.001*** (2.55)	0.001*** (2.69)	0.001*** (2.73)	0.001*** (2.60)	0.001*** (2.62)	0.001*** (2.63)	0.001*** (2.67)
Total euros lent	−0.000 (0.28)	−0.000 (−0.32)	−0.000 (−0.13)	−0.000 (−0.22)	0.000 (−0.10)	−0.000 (−0.13)	−0.000 (−0.14)	−0.000 (−0.18)
Day of week dummies	Y	Y	Y	Y	Y	Y	Y	Y
Observations	499	499	499	499	499	499	499	499
R-square	0.052	0.062	0.057	0.064	0.056	0.063	0.053	0.062

The table reports selected statistics from vector autoregressions (VARs) based on the 5th daily lag of 5 endogenous variables: a summary measure of the breadth of the lending network, the change in 3-month LIBOR, the number of active lending banks, the number of active borrowing banks, and the total volume of lending. For some VARs, an interaction term involving the summary network breadth statistic and a dummy indicating the pre-tax-transfer period is additionally included as an endogenous variable with its own VAR equation. The VAR systems reported here thus have either 5 or 6 equations, depending on whether an interaction is included. Weekly changes are measured, alternatively, on each day of the week and reported separately. The table reports only on the equation having the 3-month LIBOR change as the dependent variable. The significance levels of 10% (*), 5% (**), and 1% (***) are used in the analysis.

the null hypothesis that the coefficients are jointly zero with a p-value of about 0.09.

In consideration of this, Columns (7) and (8) provide a parsimonious way of incorporating the information of all three counterparty sampling measures in the same VAR system. We compute a summary counterparty sampling measure as the first principle component score of the three measures.[12] The takeaway from VARs using this summary measure is the same as from the single-measure VARs: the measure for the sub-network that bridges from non-Italian bank lenders to Italian bank borrowers is positively related to upcoming 3-month LIBOR changes in a statistically-significant manner (t-statistic $= 2.19$), but only during high-uncertainty PTTP times.

We next re-estimate all these same VAR systems, but now using counterparty sampling measures based on the full network of lenders. Our purpose in doing so is to check if the influence of variation in counterparty sampling at a time of uncertainty might be more pervasive than only across geographical divides. That is, even though the most obvious information-flow reason for any importance of counterparty sampling related to bridging between the Italian bank sector and the non-Italian bank sector, conditions and information also vary within each of those sectors. Broader links within the Italian bank sector might eventually bridge to the non-Italian bank sector, e.g. in the sense of Granovetter (1973). And broader links within the non-Italian bank section may be a path for information to spread within that sector. Table 5 reports the results for the 3 month LIBOR equation. Individual measures of counterparty breadth and network disorder are in fact positively related to upcoming 3-month LIBOR changes for uncertainty-case PTTP times, as is the summary counterparty sampling measure. Comparing the size and statistical significance of coefficients with those in the previous table, the evidence for Granger causality is at least as strong. *R-square* statistics from Table 5 are as large as those in the previous table as well. The evidence is that counterparty sampling in a broad sense, i.e. in the full network, is important for the path of upcoming LIBOR rates.

6.2. Predicting Week-Ahead Weekly Changes

Our VARs suggest that network re-configurations in terms of information-flow-oriented counterparty sampling characteristics do in fact help information in the uncertainty-stressed part of the network (Italian banks) to flow across geography to the rest of the interbank network. The configuration changes are predictive interest rate changes a week ahead, but the interest rate changes we have examined so far are one-day changes. We next examine if longer-period changes can also be predicted. To do so, we apply our VAR setup with a focus on weekly changes in one-week ahead 3-month LIBOR.

To avoid overlapping data in our VARs, we measure the weekly changes in 3-month LIBOR on the same particular day of the week. We are especially interested in Thursday-to-Thursday changes because they coincide with the day of the week on which ECB Governing Council holds its

Table 5. Full lender bank network vector-autoregressions showing the influence of information-flow-oriented network statistics on one-week-ahead daily changes in 3-month euro LIBOR.

Regressors	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Proportion of restorations	−0.016 (−1.36)	−0.015 (−1.26)						
Proportion of restorations × PTTP dummy		0.006** (2.15)						
Standard deviation of clique size			−0.001 (−1.51)	−0.001 (−1.47)				
Standard deviation of clique size × PTTP dummy				0.001** (1.99)				
Entropy					−0.004 (−0.83)	−0.005 (−1.11)		
Entropy × PTTP dummy						0.001*** (2.33)		
Summary counterparty sampling measure							−0.001 (−1.20)	−0.001 (−1.26)
Summary counterparty sampling measure × PTTP dummy								0.001** (2.19)
Current change in 3-month LIBOR	0.085* (1.90)	0.084* (1.90)	0.080* (1.80)	0.081* (1.83)	0.082* (1.84)	0.081* (1.84)	0.083* (1.87)	0.083* (1.89)
Number of distinct lenders	−0.001** (−2.21)	−0.001** (−2.22)	−0.001*** (−2.74)	−0.001*** (−2.66)	−0.001*** (−2.49)	−0.001*** (−2.48)	−0.001*** (−2.34)	−0.001** (−2.30)
Number of distinct borrowers	0.001*** (2.51)	0.001*** (2.59)	0.001*** (3.00)	0.001*** (2.99)	0.001*** (2.65)	0.001*** (2.74)	0.001*** (2.54)	0.001*** (2.58)
Total euros lent	−0.000 (−0.045)	−0.000 (−0.25)	−0.000 (−0.46)	−0.000 (−0.29)	−0.000 (−0.13)	−0.000 (−0.35)	−0.000 (−0.42)	−0.000 (−0.50)
Day of week dummies	Y	Y	Y	Y	Y	Y	Y	Y
Observations	499	499	499	499	499	499	499	499
R-square	0.055	0.064	0.056	0.064	0.053	0.063	0.055	0.064

The table reports selected statistics from vector autoregressions (VARs) based on the 5th daily lag of five endogenous variables: a network characteristics, the change in 3-month LIBOR, the number of active lending banks, the number of active borrowing banks, and the total volume of lending. The network characteristic is either the proportion of restorations, the standard deviation of clique size, the network entropy, or a summary measure based on all three, respectively. For some VARs, interaction terms involving the network statistic and a dummy indicating the pre-tax-transfer period is also included as an additional endogenous variable with its own VAR equation. The VAR systems reported here thus have either 5 or 6 equations, depending on whether an interaction is included. Day-of-the-week dummy variables are also included as exogenous variables in every VAR. The table reports only on the equation having the 3-month LIBOR change as the dependent variable. The significance levels of 10% (*), 5% (**), and 1% (***) are used in the analysis.

meetings to determine monetary policy. If the network measures predict Thursday-to-Thursday LIBOR changes better than weekly changes ending on other days of the week, it would suggest that some of the information flow is useful for reasons beyond the Italian tax-transmission events, but also link to the information that the ECB uses in setting its policy.

Formally, we are applying a variation on Equation (1) where the periodicity of the data is weekly, the lag-length is one and day-of-the-week dummy effects are zeroed-out. The number of observations drops to about 100, matching to the number of business weeks in our two-year sample period.

Table 6 shows the key results of this investigation. Because we have found that counterparty sampling in the full network of all lenders is important, we use the full network.[13] In the table, we report only the 3-month LIBOR estimating equation as usual and, to save space, we show only the coefficients relating to the network summary counterparty sampling measure. We have confirmed that the findings are consistent with those for VARs based on the specific counterparty sampling measures.

VAR systems with the counterparty sampling measure but no interaction term show only a statistically-insignificant effect of the measure on LIBOR for every day-of-the-week case. VAR systems with an interaction term show significant effects of the PTTP interaction term for LIBOR changes ending on Wednesday and Thursday. The network changes during uncertainty-case PTTP periods are predictive of weekly LIBOR changes that coincide with the day of the week of the ECB meetings or the day before. The meetings occur every other week, so the correspondence is strong.

7. Exploring Possible Non-learning Explanations

We have shown that changes in interbank network lending linkages that broaden banks' counterparty cliques in more varied and disorganized ways, a pattern we call "counterparty sampling," is predictive of changes in the 3-month LIBOR interest rate during disequilibrating spates of uncertainty. The evidence is consistent with the idea that banks gain information *via* counterparty sampling that is useful for understanding the upcoming path of an important interest rate.

In this section, we consider and rule out two alternative explanations. First, we consider the possibility that the statistical predictive power of counterparty sampling that we have documented is only a stand-in for the spates of volatility that are the occasion of the counterparty sampling. In that case, counterparty sampling is simply a coincident occurrence around the time of interest rate shifts, and not causal in any sense.

Second, we consider the possibility that lending banks do not learn about the interest rate path by counterparty sampling, but rather lenders already have private information about the likely path of rates. In that case, the lending that we have termed counterparty sampling would simply be a business use of the information. In particular, lenders who

Table 6. Vector-autoregressions showing the influence of the information-flow-oriented network summary statistic on one-week-ahead weekly changes in 3-month euro LIBOR.

Weeks ending on	Monday		Tuesday		Wednesday		Thursday		Friday	
Regressors										
Summary counterparty sampling measure	−0.014 (−0.14)	−0.017 (−0.16)	0.001 (0.31)	0.001 (0.30)	−0.001 (−0.26)	−0.001 (−0.35)	−0.002 (−0.92)	−0.002 (−0.88)	−0.002 (−1.03)	−0.002 (−1.12)
Summary counterparty sampling measure × PTTP dummy		−0.005* (−1.84)		0.001 (0.33)		0.003*** (2.98)		0.003*** (2.70)		0.001 (1.16)
Current change in 3-month LIBOR	−0.014 (−0.14)	−0.017 (−0.16)	0.197** (2.09)	0.198** (2.10)	0.310*** (3.39)	0.329*** (3.74)	0.174* (1.85)	0.208** (2.26)	−0.050 (−0.50)	−0.042 (−0.43)
All other endogenous variables included as regressors	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	91	91	101	101	103	103	102	102	98	98
R-square	0.075	0.082	0.224	0.225	0.179	0.242	0.211	0.263	0.132	0.144

The table reports selected statistics from vector autoregressions (VARs) based on a one-week lag of four variables: the network summary breadth characteristic, the change in 3-month LIBOR, the number of active lending banks, the number of active borrowing banks, and the total volume of lending. The VAR systems reported here thus have either 5 or 6 equations, depending on whether an interaction is included. The table reports only on the equation having the 3-month LIBOR change as the dependent variable. Separate VARs are estimated for weekly changes ending with each respective day of the week. The significance levels of 10% (*), 5% (**), and 1% (***) are used in the analysis.

expect rates to rise might want to establish broader lending relationships to profit.

7.1. Is It Network Changes or It Is Heightened Stress That Is Predictive?

Our story is that changes in banks' counterparty sampling is predictive of changes in an important interest rate during spates of disequilibrium. The evidence suggests that the counterparty changes not only help the system re-equilibrate but also perhaps help banks that are remote from the disequilibrating event to learn about the situation as it matters for future rates. Since the counterparty changes and the stress periods are aligned in time, it is important to confirm that it is not the liquidity stress itself that is predictive of future interest rates.

In Table 7 we present evidence that it is the network changes that matter for future rates, not the spates of volatility that are the temporal occasion of the network changes. The table presents the results of our now-familiar VAR systems with the addition of equations for the recent (5-business-day) variance of overnight euro LIBOR rates and the recent variance of the euro LIBOR spread. The euro LIBOR spread is defined as the difference between 3-month LIBOR and overnight LIBOR. We allow for the time-series connection between these rates and the future 3-month LIBOR rate to be different in the PTTP uncertainty cases *vs.* the TTP control cases by including interaction terms, analogous to the ones in our VAR systems above.

Table 7 presents findings for the VAR equation with future 3-month LIBOR as the dependent variable. Column (1) presents a VAR that does not include the lagged counterparty sampling regressors, but only the variance measures, while Columns (2), (3), and (4) present VARs that each include one of the counterparty sampling measures as well. The variances have uniformly tiny coefficients and *t*-statistics in all four VARs. The counterparty sampling measures' interaction terms have statistically-significant positive coefficients, just as in the VARs presented above. The bottom line is that the counterparty sampling measures Granger-cause future 3-month LIBOR, not the elevated variances.

We have now checked whether it is the increases in uncertainty themselves that correspond to the predictability. It is not. Network changes in response to spates of uncertainty are what predict future rates.

7.2. Are Lending Banks Exploiting Information They Already Have?

It seems reasonable that lending banks might broaden their lending cliques when they expect interest rates to rise, consistent with our evidence so far. New or restored lending relationships would become more profitable when rates are higher. In this section, we pursue two additional analyses that, in the end, instead shore up the learning interpretation we have proposed.

Table 7. Vector-autoregressions showing the influence of variance summaries on one-week-ahead daily changes in 3-month LIBOR in the full network of all lenders.

	(1)	(2)	(3)	(4)
Regressors				
Recent LIBOR sum of squares	−0.001 (−0.03)	−0.009 (−0.14)	−0.020 (−0.32)	−0.004 (−0.08)
Interaction with PTTP	0.062 (0.41)	0.056 (0.37)	0.078 (0.52)	0.054 (0.36)
Recent spread sum of squares	0.006 (0.09)	0.020 (0.30)	0.031 (0.46)	0.018 (0.28)
Interaction with PTTP	−0.045 (−0.27)	−0.057 (0.35)	−0.080 (−0.49)	−0.054 (−0.33)
Proportion of restorations		−0.015 (−1.28)		
Interaction with PTTP		0.006* (1.88)		
Std. Dev of clique size			−0.001 (−1.54)	
Interaction with PTTP			0.001* (1.68)	
Entropy				0.084* (1.91)
Interaction with PTTP				0.001** (2.01)
Endogenous variables included	Y	Y	Y	Y
Day of week dummies included	Y	Y	Y	Y
Observations	499	499	499	499
R-square	0.055	0.065	0.065	0.065

The table reports selected statistics from vector autoregressions (VARs) based on a 5-business-day lag of several variables: the sums of squares of recent changes in overnight LIBOR and LIBOR spread, the change in 3-month LIBOR, a specific counterparty sampling characteristic (alternatively, either no counterparty sampling characteristic included, or the proportion of restorations, or the standard deviation of clique size, or the entropy). The VAR systems also include counterparty sampling characteristic interaction terms relating to the pre-tax-transfer uncertainty PTTP cases. All VARs also include other the endogenous variables as in Equation (1): the number of active lending banks, the number of active borrowing banks, and the total volume of lending, with coefficients not reported to save space. All VARs include day-of-the-week dummy variable regressors. The table reports only the system equation having the 3-month LIBOR change as the dependent variable. Column (1) reports for the VAR including no breath characteristic. Column (2) reports for the VAR including also the proportion of restorations, column (3) for the VAR including the standard deviation of clique size, and column (4) for the VAR including no breath characteristic, respectively. The significance levels of 10% (*) and 1% (***) are used in the analysis.

7.2.1. Counterparty Sampling Does Not Predict the Rate Directly Traded in the e-MID Market

Our evidence on rate predictability so far is based on week-ahead 3-month LIBOR rates. One reason to think this is learning-driven is that lending banks would not profit more by lending a week before interest rate increases. But, even so, they might be established in advance some lending relationships and patterns that would have good profit potential a week later. Another reason working against such a private information explanation is that the major trading rate in e-MID is not a 3-month rate, but rather an overnight rate. Private information on the 3-month rate cannot be directly exploited by trading a 3-month rate. But, even so, 3-month LIBOR and overnight rates are closely correlated, so perhaps overnight trading could link to 3-month rate information.

To check, we estimate some sets of VAR systems that include the overnight LIBOR rate, which is essentially the same as the rate traded in our e-MID data. We ask if the counterparty sampling measures predict changes in overnight LIBOR on a week-ahead basis and/or on an overnight basis, either during normal times or during the pre-tax-transfer period spates of uncertainty. In a nutshell, we find that they do not. Counterparty sampling is predictive of only the 3-month rate—a rate that relates most closely to bank's loan businesses as a major benchmark but is not traded in our data and is not a major part of e-MID value outside our data either. The evidence is therefore more consistent with banks learning than it is with banks exploiting private information in their trading.

The tabulated evidence is in Table 8. The mechanical structure of the VAR systems we report in that table is analogous to those in earlier tables, so we tabulate only the key estimates focusing on week-ahead changes in overnight LIBOR (i.e. a timing exactly the same as in Equation (1). We report coefficients on our four counterparty sampling measures as a possible Granger-cause of week-ahead changes in overnight LIBOR (i.e. the three specific measures pertaining to restorations, variation in clique size, and entropy, and the principal-component summary measure of counterparty sampling). As usual, there is a separate VAR system for each measure. Each VAR system has the same full set of endogenous variables as in previous tables, except that now overnight LIBOR is included instead of 3-month LIBOR. The counterparty sampling measures used here are based on the lending network of non-Italian banks that lend by lifting the quotes of Italian banks, to maximize the chance of prior private information being available to the lenders.

The estimated coefficients are directly comparable to those in Table 5 for 3-month LIBOR. In contrast to the coefficients in that table, the ones in Panel A are uniformly small and statistically insignificant in every case. There is no counterparty sampling measure that shows any sign of being predictive for week-ahead overnight LIBOR, whether in normal times or during spates of pre-tax-transmittal period uncertainty.

Table 8 has been constructed using the non-Italian lender network. We note that we obtain qualitatively identical results if we use the full network of lenders.

Our takeaway is that the learning story is better supported than the prior private information story. Banks that are choosing their e-MID lending counterparties based on private information about upcoming overnight LIBOR would most profitably use it to trade on overnight LIBOR, but we do not see the tracks of such trading.

7.2.2. Surprises in Counterparty Sampling Are Negatively Related to Future Rates

Our main evidence on predicting 3-month LIBOR with counterparty sampling measures shows positive

Table 8. Non-Italian bank network vector-autoregressions showing the influence of information-flow-oriented network statistics one-week-ahead daily changes in overnight euro LIBOR.

Regressors	(1)	(2)	(3)	(4)
Proportion of restorations	0.027 (0.78)			
Proportion of restorations × PTTP dummy	−0.032 (−0.81)			
Standard deviation of clique size		0.001 (0.02)		
Standard deviation of clique size × PTTP dummy		−0.001 (−0.82)		
Entropy			−0.008 (−1.05)	
Entropy × PTTP dummy			−0.004 (0.62)	
Summary counterparty sampling measure				−0.001 (−0.24)
Summary counterparty sampling measure × PTTP dummy				−0.001 (−0.74)
Other endogenous variables included?	Y	Y	Y	Y
Day of week dummies included?	Y	Y	Y	Y
Observations	499	499	499	499
R-square	0.025	0.024	0.027	0.024

The table reports selected statistics from vector autoregressions (VARs) based on lags of 5 variables: a network characteristics, the change in overnight LIBOR, the number of active lending banks, the number of active borrowing banks, and the total volume of lending. The network characteristic is either the proportion of restorations, the standard deviation of clique size, the network entropy, or a summary measure based on all three, respectively. The VARs additionally include interaction terms involving the network statistic. The VAR systems reported here thus have 6 equations. Day-of-the-week dummy variables are also included as exogenous variables in every VAR. The table reports only on the equation having the overnight LIBOR change as the dependent variable.

relationships. The counterparty sampling measures combine expected and surprise components. If banks are adjusting their lending cliques based on the expectation of higher rates, the core of the predictive association would be with the expected part of the counterparty sampling series, because that is the part that has a pattern relative to conditions in the network. In contrast, we next show a strong predictive relationship in the negative direction for the surprise series—more surprises in the extent of linkages, and lower interest rates in the future. A negative relationship is not consistent with the idea of banks' broadening their lending cliques to take advantage of expected higher rates, but is consistent with the learning interpretation.

To set up this investigation, we begin by computing predictions of a standard long short-term memory (LSTM) deep learning neural network model for bank's lending cliques and compare them to the actual lending cliques that are observed during each week of our data. Simple neural network models can be thought of as approximations of complex non-linear data-generating functions (Hornik et al. 1989). Feedforward neural networks can approximate non-linear ARMA models with standard time-series decay properties, for example. Hochreiter and Schmidhuber (1997) make the case that LSTM networks are especially good for fitting complex sequences of data, such as time series. In our situation, we have the possibility that links depend on long-past links that sometimes become important. Dixon et al. (2022) provide an overview of LSTM models that emphasizes how their combination of forget gates, output gates, and input gates enables them to propagate forward a smoothed hidden state that essentially contains a long-term memory. This is very appealing for our situation where we think of conditions of certainty and uncertainty that may repeat over time with long and variable lags. Additionally, LSTM models do not rely on explicit assumptions of stationarity, which we cannot be sure of for the case of interbank networks, and they can deal with sparse data matrices like our adjacency matrices.

Past adjacency matrices are the source of input features for our application of an LSTM model. Even though our LSTM model is a standard implementation, we implement important configuration and hyperparameter choices that help us take account of the changeability of adjacency matrices. At each date, the LSTM model updates a hidden state computed from past adjacency matrices. We append a decoder to convert the hidden state to link probability predictions. The overall model setup involves specific choices in terms of the loss function that we use to guide the estimation of parameters, especially to take account of the sparsity of adjacency matrices. We explain our setup and choices in detail in the Appendix. We train the model in a rolling window fashion, based on 10 weeks' of prior adjacency matrices. After each training of the model, we use it to predict bank-specific linkages for the next week of our data set.[14] Then, to compute a network-wide summary statistic of the extent of surprises at a given time, we measure the difference between the average of the predicted probabilities for each possible pairwise link and the average of the 0–1 indicator variables for the actual existence of links.

With the surprise measure in hand, we then follow the recipe of the second part of our empirical design, using the surprise series as the counterparty sampling measure in a set of Granger causality models. For comparison with our earlier results, we predict week-ahead one-day changes and also weekly changes from Thursday to Thursday. The results are presented in Table 9.

Panel A presents the results of using LSTM surprises to predict week-ahead one-day changes in 3-month LIBOR. Column (1) is directly comparable to Table 4, with the VAR system having the same setup except that now the counterparty sampling measure is the LSTM surprise. Like the results in that table, the level effect of the counterparty sampling measure is small and statistically insignificant. In contrast to the results in that table, the coefficient on the interaction term is negative and fairly strongly statistically significant (just shy of the 5% p-value level). Column (1) VAR results establish that the effect of LSTM surprises during PTTP uncertainty periods is significantly less that the insignificant level overall. Column (2) provides a second perspective by removing the level effect from the VAR and instead including interaction terms for both PTTP and non-PTTP periods. Consistent with the impression given by

Table 9. Non-Italian bank network vector-autoregressions showing the influence of the aggregate intensity of surprises in the overnight lending linkages to Italian banks for the week-ahead one-day changes in LIBOR and Thursday-to-Thursday weekly changes.

Regressors	(1)	(2)
Panel A: Predicting one-week-ahead daily changes in 3-month LIBOR		
Surprise in extent of counterparty sampling relative to LSTM model predictions	−0.0103 (−0.78)	
Surprise in extent of counterparty sampling relative to LSTM model predictions × non-PTTP Dummy		−0.0103 (−0.78)
Surprise in extent of counterparty sampling relative to LSTM model predictions × PTTP Dummy	−0.0495*** (−3.15)	−0.0392*** (−4.46)
Other endogenous variables included?	Y	Y
Day of week dummies included?	Y	Y
Observations	471	471
R-square	0.053	0.055
Panel B: Predicting weekly Thursday-to-Thursday changes in 3-month LIBOR		
Surprise in extent of counterparty sampling relative to LSTM model predictions	0.0185 (0.23)	
Surprise in extent of counterparty sampling relative to LSTM model predictions × non-PTTP Dummy		0.0185 (0.23)
Surprise in extent of counterparty sampling relative to LSTM model predictions × PTTP Dummy	−0.1740* (−1.92)	−0.1554*** (−3.66)
Other endogenous variables included?	Y	Y
Observations	93	93
R-square	0.240	0.240

The table reports selected statistics from vector autoregressions (VARs) based on lags of 5 variables: network characteristics, the change in overnight LIBOR, the number of active lending banks, the number of active borrowing banks, and the total volume of lending. The network characteristic is the surprise in the network-wide proportion of links relative to the forecasts of a long short-term memory (LSTM) deep learning model. The VARs additionally include an interaction term involving the network statistic. The VAR systems reported here thus have 6 equations. Day-of-the-week dummy variables are also included as exogenous variables in the Panel A VARs. The table reports only on the equation having the 3-month LIBOR change as the dependent variable. Figures in parentheses are *t*-statistics and the significance levels of 10% (*) and 1% (***) are used in the analysis.

Column (1), Column (2) shows that the non-PTTP effect of the LSTM surprises is insignificant and the PTTP effect is strongly significant and negative. Thus, overall, the PTTP effect is both relatively more negative than the TTP effect and is also negative in an absolute sense.

Panel B presents results of a similar analysis as Panel A, except where the dependent variable being predicted is the Thursday-to-Thursday weekly change in 3-month LIBOR. The results are fully consistent with those in Panel A. The level effect of LSTM surprises is insignificant in Column (1) and the interaction term for the effect during the PTTP uncertainty period is negative and strongly statistically significant. The non-PTTP effect in Column (2) is small and insignificant, and the PTTP uncertainly effect is negative, large in absolute value, and statistically significant.

The finding that surprises in linkages are negatively associated with future changes in rates is not consistent with the idea that banks broaden their lending cliques to take advantage of expected future rate increases. This idea is ruled out. The clique changes we focus on here are deviations from the pattern uncovered by the deep learning model, not expected. And when there are a lot of them, rates tend to decrease, not increase. In contrast, the learning idea allows for such a pattern, and so is not ruled out.

8. Conclusion

Interbank networks are important for several reasons relating to systemic risk and the services the banking system provides to the economy. In this paper, we have suggested an additional reason for their importance: information flow and learning. We have provided statistical evidence that changes in lending banks' use of the interbank network during periods of stress are predictive of future interest rates. In particular, changes that reflect an increase in the breadth of lenders' counterparty cliques, a phenomenon we call "counterparty sampling," are the key. Our findings suggest that ideas about networks from sociology and information theory are applicable to bank networks in that ties that bridge otherwise-distant parts of the interbank network are important for information flow. Banks learn from lending in the interbank network.

Notes

1. One example is the economically-pivotal U.S. federal funds (or "fed funds") market. Interbank lending connections can increase systemic risk by propagating one bank's failure to other network banks that have lent to it. But they can also help the banking system to be more reliable for customers who want to withdraw funds because banks can borrow to cover any temporary cash shortage, and also to be more reliable for outside borrowers because banks can easily obtain funds to lend if they don't have enough.
2. The September 2019 fed funds market situations arose as a combination of exogenous and endogenous influences. Corporate taxes and a large Treasury bond auction drained reserves from the banking system at a time when reserves were already at historic lows (Anbil et al. 2020). The mix makes causal inference more challenging than in our setting.
3. This market traded substantial volumes of deposits until the Global Financial Crisis, but eventually withered as liquidity for banks' short-term needs became readily available *via* ECB facilities during the more-recent quantitative easing and ample reserves regimes (Pontus et al. 2021). Though e-MID is not the only means for interbank loans during this period, transacted e-MID rates closely track other interbank rates due to banks' opportunities to arbitrage (Angelini et al. 2011).
4. The requirement was instituted in 2002 by the Italian government *via* Decree 63/2002, Article 1. Specifically it required transmittal of 80% of all taxes collected by the third business day after collection. Typically this would be the 23rd of each month (or the following day if the 23rd is a weekend or holiday). Thus there are large changes in government deposits at Italian banks during the few days preceding the 23rd.
5. Denbee et al. (2021) emphasize that whether a bank network amplifies or dampens shocks depends on what

banks infer from each other's liquidity holdings more than on network topology, similar to our idea that information flow is more key than topology.

6. For example, see Anderson et al. (2019) for the U.S. banks in the nineteenth century, Craig and von Peter (2014) for German banks, Blasques et al. (2018) for the Netherlands, and Iori et al. (2008) for Italy.
7. Beaupain and Durré (2008) discuss the differences. In a nutshell, the ECB's main refinancing operations (MROs)—the most important source of overnight liquidity in the eurozone—effectively discouraged bids for funds during periods when participants expected target rate cuts, which led to occasional liquidity shortages. Subsequently, maturities of MRO funds where shortened and the timing of reserve maintenance periods (the spans over which required reserve levels much be satisfied) were adjusted.
8. The actual liquidity drain comes during the tax transfer period, which could be a surprise in itself. In fact, in the U.S. case noted in the introduction, a corporate tax payment event seems to have been one of the sources of the volatility. Below we provide statistical support for our notion that, in the context of the eurosystem, it is actually a resolution of uncertainty.
9. The t-statistics are simple ts that do not adjust for the overlapping data. They are intended only to be indicative of the pervasiveness of the sharp differences. For this exercise, as well as for analogous rolling weekly calculations later in the paper, we measure days as business days, accounting for bank holidays.
10. Formally, our basic daily VAR setup can be written as

$$\mathbf{y}_t = \boldsymbol{\Gamma}\mathbf{y}_{t-5} + \boldsymbol{\Theta}\mathbf{d}_t + \boldsymbol{\varepsilon}_t, \tag{1}$$

where $\mathbf{y}_t = (y_{1,t}, y_{2,t}, y_{3,t}, y_{4,t}, y_{5,t})$ is the vector of (1) current observations on daily changes in 3-month LIBOR, (2) a specific counterparty sampling measure (i.e. either the proportion of restorations, the standard deviation of clique size, or entropy), (3) the number of distinct borrowers, (4) the number of distinct lenders, and (5) the number of euros lent in aggregate. $\mathbf{d}_t$ is a 4-element vector of dummy variables for the days of the week Tuesday through Friday, and $\boldsymbol{\varepsilon}_t$ is a random error vector. $\boldsymbol{\Gamma}$ is a 5×5 matrix of VAR time-series coefficients to be estimated, and $\boldsymbol{\Theta}$ is a 5×4 matrix of coefficients. As noted above, we sometimes add an important sixth equation to the VAR system, appending $y_{6,t}$ to the $\mathbf{y}_t$ vector, where $y_{6,t} = y_{2,t} \times P_t$, with P_t being a dummy variable registering whether date t is in a PTTP uncertainty-case period. As an alternative for some analyses, we instead set up a basic weekly VAR as

$$\mathbf{y}_t = \boldsymbol{\Gamma}\mathbf{y}_{t-1} + \boldsymbol{\varepsilon}_t, \tag{2}$$

where time is indexed in weeks (instead of in days as in our basic daily VAR), so $y_{1,t}$ becomes current observation on *weekly* changes in 3-month LIBOR, and other elements of y_t are also weekly observations.
11. Mechanically, an adjacency matrix for an N-bank network is an $N \times N$ matrix of zeros and ones, where a one indicates that a lending relationship from a row bank to a column bank existed on at least one day during the time period.
12. The first principal component is computed in the standard way, as the vector that maximizes the variance of the projection of all three counterparty sampling measures.
13. We have checked to be sure that we obtain similar results with the non-Italian bank sub-network. This is in fact the case.
14. We lose some weeks at the beginning due to the need to train the model the first time.

Acknowledgments

The authors sincerely thank the Editor, the Associate Editor, and three anonymous referees for their valuable comments that helped improve our manuscript significantly.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

All authors acknowledge partial support from the JPMC Fellowship; Wei Qian is also partially supported by NSF (DMS-1916376).

References

Aciero L, Heijmans R, Heuver R, Massarenti M, Picillo C, Vacirca F. 2016. How to measure the unsecured money market: the Eurosystem's implementation and validation using TARGET2 data. Int J Cent Bank. 10:187–211.

Afonso G, Kovner A, Schoar A. 2014. Trading partners in the interbank lending market. Federal Reserve Bank of New York Staff Reports, no. 620.

Allen F, Gale D. 2000. Financial contagion. J Polit Econ. 108(1):1–33.

Anbil S, Anderson A, Senyuz Z. 2020. What happened in money markets in September 2019? FEDS Notes. 2020(2527). https://doi.org/10.17016/2380-7172.2527

Anderson H, Paddrik M, Wang JJ. 2019. Banking networks and systemic risk: evidence from the National Banking Acts. Am Econ Rev. 109(9):3125–3161.

Angelini P, Nobili A, Picillo C. 2011. The interbank market after August 2007: what has changed, and why? J Money Credit Bank. 43(5):923–958.

Ashcraft A, Duffie D. 2007. Systemic dynamics in the federal funds market. Am Econ Rev. 97(2):221–225.

Beaupain R, Durré A. 2008. The interday and intraday patterns of the overnight market: evidence from an electronic platform. ECB Working Paper Series No. 998.

Blasques F, Bräuning F, van Lelyveld I. 2018. A dynamic network model of the unsecured interbank lending market. J Econ Dyn Control. 90:310–342.

Brunetti C, Harris J, Mankad S, Michailidis G. 2019. Interconnectedness in the interbank market. J Financ Econ. 133(2):520–538.

Burt R. 1992. Structural holes: the social structure of competition. Cambridge, MA: Harvard University Press.

Caballero R, Simsek A. 2013. Fire sales in a model of complexity. J Finance. 68(6):2549–2587.

Craig B, von Peter G. 2014. Interbank tiering and money center banks. J Financ Intermed. 23(3):322–347.

Denbee E, Julliard C, Li Y, Yuan K. 2021. Network risk and key players: a structural analysis of interbank liquidity. J Financ Econ. 141(3):831–859.

Diamond D, Dybvig P. 1983. Bank runs, deposit insurance, and liquidity. J Polit Econ. 91(3):401–419.

Dixon M. 2020. Machine learning for finance.

Dixon M, Halperin I, Bilokon P. 2022. Machine learning in finance: from theory to practice. Cham: Springer Nature Switzerland.

Eisenberg L, Noe T. 2001. Systemic risk in financial systems. Manage Sci. 47(2):236–249.

Erol S, Ordoñez G. 2017. Networks reactions to banking regulations. J Monet Econ. 89:51–67.

Finger K, Fricke D, Lux T. 2013. Network analysis of the e-mid overnight money market: the informational value of different aggreagation levels for instrinsic dynamic processes. Comput Manag Sci. 10(2–3):187–211.

Glasserman P, Young HP. 2016. Contagion in financial networks. J Econ Lit. 54(3):779–831.

Goodfellow I, Bengio Y, Coureville A. 2016. Deep learning. Cambridge (MA): MIT Press.
Granovetter M. 1973. The strength of weak ties. J Sociol. 78(6): 1360–1380.
Granovetter M. 1983. The strength of weak ties: a network theory revisited. Sociol Theory. 1:201–233.
Granovetter M. 2005. The impact of social structure on economic outcomes. J Econ Perspect. 19(1):33–50.
Greenwood R, Landier A, Thesmar D. 2015. Vulnerable banks. J Financ Econ. 115(3):471–485.
Hatzopoulos V, Iori G, Mantegna RN, Miccichè S, Tumminello M. 2015. Quantifying preferential trading in the e-mid interbank market. Quant Finance. 15(4):693–710.
Hirshleifer D, Sheng J. 2022. Macro news and micro news: complements or substitutes? J Financ Econ. 145(3):1006–1024.
Hochreiter S, Schmidhuber J. 1997. Long short-term memory. Neural Comput. 9(8):1735–1780.
Hornik K, Stinchcombe M, White H. 1989. Multilayer feedforward networks are universal approximators. Neural Netw. 2(5):359–366.
Iori G, De Masi G, Precup OV, Gabbi G, Caldarelli G. 2008. A network analysis of the Italian overnight money market. J Econ Dyn Control. 32(1):259–278.
Jacobs P, Lewis P. 1978. Discrete time series generated by mixtures. J R Stat Soc B Methodol. 40(1):94–105.
Kaminska I. 2019. A story about a liquidity regime shift, *The Financial Times, Alphaville Column,* September 20.
Monticini A, Ravazzolo F. 2011. Forecasting the intraday market price of money. Working Paper No. 2011/06, Norges Bank.
Pfaltz J. 2012. Entropy in social networks. SocInfo2012, International Conference on Social Informatics.
Pontus Å, Corsi M, Grossmann-Wirth V, Hudepohl T, Mudde Y, Rosolin T, Schobert F. 2021. Demand for central bank reserves and monetary policy implementation frameworks: the case of the Eurosystem. Occasional Papers No 282/September, European Central Bank.
Shleifer A, Vishny R. 2011. Fire sales in finance and macroeconomics. J Econ Perspect. 25(1):29–48.

Appendix: Setup of the LSTM Model

In this study, we measure surprises in counterparty sampling by comparing link outcomes to the link predictions of a standard long short-term memory (LSTM) recurrent neural network (RNN) model and then investigate whether the extent of surprising new linkages is also predictive of future interest rates. Though, as with most deep learning approaches, LSTM-based RNNs have the disadvantage that they are challenging to interpret, they also have advantages that are important for our application. In contrast to standard time series models, LSTM does not rely as much on assumptions about the data, such as linear link connections and stationarity. Further, with appropriate loss functions, we have found that LSTM models deal well with the relatively sparse linkage matrices that characterize interbank networks.

The generic logic and setup of LSTM models are provided in, for example, Goodfellow et al. (2016) and, with particular application to finance, in Dixon (2020). A standard LSTM model as applied to the sequence of adjacency matrices updates the model's "hidden state," which is a complex non-linear combination of data on historical adjacency matrices that carries the information needed to predict adjacency matrices. We also add a fully-connected layer decoder, incorporating a loss function as described below, to transform the hidden state into adjacency matrix predictions.

For our implementation, several specific considerations are significant. In particular, these are the inclusion and nature of a decoder, the loss function that underlies parameter estimates, the training and testing sample setups, and specific hyperparameter choices. In the following paragraphs, we describe these considerations and the design choices that they lead to. We also describe how we benchmark our LSTM model to substantiate that it is a high-quality link prediction process that we can depend on to determine the link surprises that we use in the paper.

Inputting a history of adjacency matrices $\{A_{t-k}, \ldots, A_{t-1}\}$, we want to output a prediction for the next adjacency matrix, $\widehat{A}_t$. To the LSTM, we append a fully connected layer to transform the output hidden state h_t to obtain the one-step-ahead prediction. The decoder is the sigmoid function σ in the expression.

$$\hat{p}_t = \sigma(W_h h_t + b),$$

where W_h and b are the weight and bias terms applied to the fully connected layer h_t. The output $\hat{p}_t$ is interpreted as a probability value which we then use to predict links based on a maximal area under the ROC curve (AUC) criterion. The use of a decoder in this way is added to the generic LSTM model.

We need a specific loss function and optimizer to train the model to best improve link prediction accuracy. There are two special considerations in our setting. First, based on an examination of the e-MID data, the network snapshots are sparse in that there are many more zero elements than nonzero. Therefore, we use a weighted squared error loss function as follows:

$$Loss = \sum_t^k \sum_i^N \sum_j^N (a_{i,j,t} - \hat{a}_{i,j,t})^2 * \lambda_{i,j}$$

where $a_{i,j,t}$ is the element in the adjacency matrix A_t and $\hat{a}_{i,j,t}$ is the element in the output probability matrix $\widehat{A}_t$, and where $\lambda_{i,j}$ is the parameter that controls the loss weight. For simplicity, the $\lambda_{i,j}$ value does not change with different time *t*. The point for adopting different $\lambda_{i,j}$ for null or existing links is to bias the system to choosing low probabilities, which will translate into null links, to match the bias toward sparsity in the system. We have also experimented with using the unweighted squared error loss function to optimize parameters in our LSTM, finding lower accuracy.

As an important consideration in the context of an interbank network is that positive links relate to systemic risk, as discussed in the text—false negatives should therefore be considered more costly than false positives. In our training process, we assign a higher $\lambda_{i,j}$ value to cases with positive links and a lower $\lambda_{i,j}$ value to cases without positive links. To avoid overfitting, we also employ the regularization term L_{reg} which is the sum of squares of the weight parameters. Therefore, the total loss function is defined as:

$$Loss_{total} = Loss + \beta Loss_{reg}$$

where β is the trade-off tuning parameter. To minimize the total loss $Loss_{total}$, we make use of the Adam optimizer.

To train the LSTM model, we feed *k* historical interbank network snapshots $(A_{t-k}, \ldots, A_{t-1})$ to predict A_t and use the estimated parameters we obtain from the training process to feed into the network snapshots $(A_t, \ldots, A_{t+k-1})$ to obtain the prediction for A_{t+k}. With aggregated weekly data from 2005 to 2006, we train and test the performance on a rolling window basis with window size $k = 10$.

There are several particular choices that must be made for the training process, which we have chosen based on gird searches. The number of the hidden layers of the LSTM model is $d = 12$. The weight decay parameter of the Adam optimizer is 1×10^{-5}, and the learning rate is 0.01. The penalty value $\lambda_{i,j}$ is 4 for the positive links in the loss function. The trade-off tuning parameter β is 0.01. With these model settings, we apply the LSTM model to $N = 164$ banks in the network to get predictions that become the basis for measuring whether a particular link is a surprise, as used in the final section of the paper.

The point of our LSTM application is to get strong performance in link prediction within the setting of the sparse and dynamic interbank network. We have compared the performance of our model as described here with that of a standard discrete autoregressive (DAR) model as originated by Jacobs and Lewis (1978). The average out-of-sample AUC from the DAR model, with windowing setup to match our design, is 0.71. The average out-of-sample AUC from the LSTM model is 0.86. Thus, we believe it is a superior choice to use the model described in this Appendix as the basis for evaluating surprises for the analysis in our paper than to use the standard model.