A Mechanism for LIBOR*

Brian Coulter[†] and Joel Shapiro[‡]
University of Oxford

September 2014

Abstract

The investigations into LIBOR have highlighted that it is subject to manipulation. We propose a mechanism that gets banks to reveal their borrowing costs truthfully at no cost to the administrator. The mechanism works even when borrowing does not occur. First, banks report the rates at which they can borrow. Second, a whistleblower bank may contest another bank's report by revealing a transaction or stating a different rate at which the reporting bank could borrow. Third, the whistleblower's claim and the initial reported rate are confirmed or denied by the willingness of other banks to lend at these rates.

Keywords: LIBOR, mechanism design
JEL Codes: D82, G21, G28

[†]Saïd Business School, University of Oxford, Park End Street, Oxford OX1 1HP. Contact: Brian.Coulter@sbs.ox.ac.uk

[‡]Saïd Business School, University of Oxford, Park End Street, Oxford OX1 1HP. Email: Joel.Shapiro@sbs.ox.ac.uk Phone: +44 1865 288973.

1   Introduction

The London Interbank Offered Rate (LIBOR) is supposed to represent the average rate at which banks can borrow in the unsecured market. It is computed by taking the trimmed mean of the reported borrowing rates of the banks on the LIBOR panel. These reported rates are not verified by transactions; moreover, there are sometimes few, if any, transactions to verify. The ongoing LIBOR scandal has already resulted in fines of over \$5.7 billion for inappropriate submissions and repeated attempts to manipulate LIBOR.[1]

There are two reasons why banks may want to manipulate LIBOR. First, manipulating one of the rates by even a fraction of a basis point can bring substantial gains to banks through their LIBOR exposures. The possible gains are large: the market for derivative and loan products that use LIBOR rates has been estimated as larger than \$300 trillion (Wheatley, 2012). Second, banks may be motivated to lower their reported costs of borrowing to hide their credit risk from counterparties and the financial markets.[2]

There is a clear need to reform the process by which LIBOR is determined. In this paper, we propose a mechanism (which we name the 'whistleblower' mechanism) that resolves many of the issues facing LIBOR. Most importantly, the unique pure-strategy equilibrium has all banks revealing their borrowing costs truthfully. Second, despite the presence of fees and punishments in the mechanism, they are never used in equilibrium. Therefore, the whistleblower mechanism has zero cost.[3] Third, the mechanism can elicit borrowing costs even when there are no transactions by using the willingness to lend of other banks. Last, the mechanism does not rely on a specific functional form to compute LIBOR from the banks' reports. Instead, the mechanism ensures that the submissions themselves are accurate, allowing the administrator to choose the desired aggregation function.

The whistleblower mechanism proceeds in three stages. First, in the reporting stage, all banks in the panel report the rate at which they could borrow on the unsecured market. These rates are revealed to the panel and the LIBOR rate is set. Second, in the contestation stage, a panel bank can become a whistleblower by contesting that another bank's report is inaccu-

---

[1]To date, the scandal has hit Barclay's (\$450 million), Deutsche Bank (\$980 million), RBS (\$1.1 billion), UBS (\$1.5 billion), and others. More investigations of other banks are currently under way.

Jill Treanor, "Banks fined record €1.7bn over benchmark interest rate rigging cartel," *The Guardian*, 4 Dec 2013.

[2]We describe the current LIBOR process, how banks manipulated LIBOR, and related measures in greater detail in Section 1.1.

[3]In Section 3.2, we demonstrate that even off the equilibrium path, the mechanism can be set up to be either zero cost or revenue generating.

rate. The whistleblower bank either presents an executed transaction that confirms its assertion, or states the rate at which it believes the initial bank can borrow (a contestation rate). If a transaction is presented that disproves the initial bank's report, the reporting bank is punished (see below) and the mechanism ends. Third, if a report is contested and no transactions are presented, the mechanism enters the verification stage. In this stage, the non-whistleblower panel banks are asked to verify the initial report. These other banks state if they would be willing to lend to the reporting bank at the initially reported rate and/or the contestation rates. If at least one bank is willing to lend at the initially reported rate and none will lend for less, the initial report is confirmed; otherwise, the report is found to have been inaccurate.[4] Banks found to have misreported or falsely accused are fined; banks that accuse accurately or present transactions that disprove the initial report are given a whistleblower payment. The mechanism is implemented using the concept of Perfect Bayesian Equilibrium.

Implementing truthful revelation in a LIBOR reporting mechanism is made particularly difficult by the banks' incentives to distort their private information. Each bank has an incentive not only to distort the rate at which it may borrow, but also to distort the LIBOR administrator's perception of the rate at which every *other* bank may borrow. Furthermore, in a LIBOR reporting mechanism, each bank has three possible sources of private information: the rate at which it can borrow, its exposure to LIBOR, and its payoff from manipulating perceptions of its credit risk. Many commonly used mechanisms can only deal with one source of private information.[5] In our model, both banks' LIBOR exposures and their incentives to indicate that they have low credit risk are completely private information. Banks' borrowing costs are 'semi-private' information: we assume that the lowest rate at which a bank may borrow is observable to (i) the party from which it may borrow and (ii) at least one other panel bank. Therefore the main relaxation of the assumption of private information is that one other panel bank aside from the potential lender may observe the rate at which borrowing may occur. This is key for our results, though we do not believe this assumption to be very strong. If a bank solicits offers for rates at which a loan may occur, for example, other banks may view the outcome of the process. Moreover, the panel banks are presumably well informed about each other's credit risk, as this is their major consideration when evaluating the rate at which to offer

---

[4]Lending would be to the LIBOR administrator, and the promised repayment would be made only if an equivalent loan to the bank would have been repaid. Thus, a synthetic bank loan may be created, which will be viewed by market participants as equivalent to lending directly to the bank in question. We discuss this further in Section 3.

[5]We discuss such mechanisms later in this section.

loans. In particular, the Financial Service Authority's (FSA's) case against Barclays specifically cited evidence that banks knew the borrowing costs of others in the industry.[6]

Another key feature of the LIBOR problem is that at the time at which a bank submits a rate to the administrator, the bank's true borrowing rate does not enter into its payoff function. The *reported* rate enters the bank's payoff function through the bank's exposure to LIBOR and through the market's perception of the bank's credit risk, but the actual rate on which the report should be based is absent. This presents a problem for the mechanism designer, as a typical technique is to use instruments that interact with the variable of interest in the agent's payoff function.[7]

A major concern of both the current LIBOR system and any subsequent mechanism is the issue of collusion. If enforceable contracts cannot be written between banks (as is generally the case for collusive agreements), collusion is impossible in our mechanism. The intuition is simple: sequential rationality causes any collusive agreement to unravel because banks have an incentive to whistleblow on each other, knowing that lending will be truthful in the verification stage of the game. Of course, the model is not an infinitely repeated game, so one might imagine that collusive agreements supported by punishment strategies are possible. We discuss in detail in Section 4 why collusion is difficult to sustain even in this situation.[8]

Our approach is complementary to the proposals made by recent regulatory reviews. Martin Wheatley (2012), in his review at the initiation of the British Chancellor of the Exchequer (the "Wheatley Review"), made

---

[6]One trigger of the recent scandal was Barclays' unhappiness at what it perceived to be underreporting of rates by other panel banks at the onset of the crisis. For example, in the FSA case against Barclays, one email from a LIBOR submitter at Barclays to his or her manager states, "Feeling increasingly uncomfortable about the way in which USD libors are being set by the contributor banks, Barclays included. . . one contributor was paying [x%] in the market at 11am [and setting at y%]. This is not an uncommon phenomenon." (Financial Services Authority, 2012a).

Another example of banks' abilities to observe each other's borrowing rates is DOJ (2013), where Rabobank traders remark several times on LIBOR manipulations by others in the market.

[7]Some recent papers postulate that banks' borrowing rate enters the utility function through a reputation cost (Chen (2012), Snider and Youle (2012), and Youle (2013) use a quadratic penalty), but until recently it was unclear what form a punishment would take for misreporting. In our model, it is not necessary for the actual rate to enter into the bank's payoff function. A reputation cost would only strengthen our results.

[8]Note that while collusion is certainly a concern, much of the recent LIBOR manipulation occured within a given firm, with traders encouraging the *internal* submitter to misreport. In the case against UBS, for example, the FSA found more than 1000 examples of inappropriate submissions by UBS itself.

three principal suggestions. These are to (1) maintain LIBOR as based on a reporting mechanism, (2) tie LIBOR reports more to actual transactions, and (3) reduce the number of tenors (maturities) and currencies for which a LIBOR rate exists based on limited transactions. The Financial Stability Board (2014), in a more recent and wider-reaching review (the "FSB Review"), suggested (1) retaining the LIBOR rate but strengthen it by underpinning "to the greatest extent possible with transactions data" and (2) augmenting LIBOR with additional, nearly risk-free reference rates.

We take the imperative to retain the LIBOR rate as a primary benchmark as given, as both reviews describe the drawbacks of moving to other measures. These drawbacks include the lawsuits that could arise from a wholesale transition to a new benchmark, substantial transaction costs, and even the potential for broader financial instability.[9] Regarding the two reviews' goal to increase the impact of relevance of transactions to the benchmark, our approach clearly links reports to transactions. When there are executed transactions, they may be used in the mechanism to verify or invalidate a bank's report. However, the mechanism also works when there are *no* actual transactions. It works by presuming that although a transaction did not occur, there is a bank willing to lend to the bank reporting its rate and the reporting bank is aware of this (although it decided not to borrow).[10] The bank that is willing to lend has the correct incentives to verify or invalidate a report by lending in the verification stage of the mechanism. Those incentives are simple, as they reflect pure revealed preference. Therefore, the third point of the Wheatley Review, which recommends eliminating LIBOR measures in thinly traded markets, is not necessary given our mechanism. And while we do not dispute the potential benefits of augmenting LIBOR with additional benchmarks, correcting the reporting incentives in LIBOR reduces the need for a costly transition.

Our approach of focusing on panel banks' 'willingness to lend' to a reporting bank has similarities to some reference rates around the world. Australian BBSW, for example, is based on "committed quotes," as they are considered binding offers to transact (Australian Financial Market Association, 2013). Other approaches, like ours, require banks to potentially transact at the rates

---

[9]Moreover, while other benchmarks exist, they measure subtly different things than LIBOR. For example, while a benchmark based on CDS premia would accurately compile information about bank credit risk, it would fail to account for both costs of funding (liquidity) and term premia (maturity). Similarly, a benchmark based on repo rates would reflect bank funding costs, but fail to fully account for bank credit risk. Both reviews point out that LIBOR cannot be replicated by any single pre-existing instrument.

[10]There are potentially situations where there are neither transactions nor offers. This implies that no bank is willing to lend, in which case the offers can be defined as infinite.

submitted. In the Telbor rate published by the Bank of Israel, for example, a panel bank is committed to either borrowing or lending a set amount (that varies by maturity) to/from other banks at a preset spread to their reported rate if one of the other banks wishes to transact (Bank of Israel, 2012). The Swedish Stibor mechanism asks banks to submit transactions with their reports, or indicative quotes if these are not possible. Once the reports are submitted, banks may freely borrow or lend to others at a pre-specified spread to their reported rates, in a set amount of SEK 100million (Swedish Bankers' Association, 2012). Our mechanism is unique, however, in using the willingness of *other* banks to lend to determine the truthfulness of a given bank's report. This makes it more difficult for a single bank, acting alone, to manipulate the benchmark.

In the following subsections, we describe some institutional details of LIBOR and then discuss how our mechanism relates to the literature on LIBOR and mechanism design. In Section 2, we set up the basic model. In Section 3, we present the whistleblower mechanism and our main result. In Section 4, we discuss how collusion would affect the mechanism. In Section 5, we demonstrate that the mechanism continues to work when credit risk signaling is explicit. In Section 6, we conclude.

1.1   The Current LIBOR Process

LIBOR rates were introduced in the early 1980s as a measure of unsecured borrowing costs for banks that would allow them to hedge their funding risk. The growth of interest rate derivatives (such as swaps and forward rate agreements) created demand for uniformity in such a measure. The rates are also now used in foreign exchange options, mortgages, bank loans, and many other products. The rates are for 15 tenors (maturities) and 10 currencies and the process was run from inception until February 2014 by the British Banking Authority (BBA), an organization representing the banking industry. The benchmark is now run by ICE Benchmark Limited, previously part of NYSE Euronext.[11] There are many other -IBORs used throughout the world. Some examples include EURIBOR (European), TIBOR (Tokyo), and NIBOR (Norwegian).[12]

---

[11]HM Treasury, 'ICE Benchmark Administration Ltd take responsibility for administrating LIBOR,' *UK Government - HM Treasury*, 3 February 2014.

[12]There is evidence that many of these benchmark rates have been manipulated as well by their respective panel banks. The following references contain some examples:

Philipp Halstrick, 'Deutsche Bank suspends traders over EURIBOR,' *Reuters*, 6 February 2013.

Ben McLannahan, 'Japanese banks accused of TIBOR fixing,' *Financial Times*, 6 Febru-

For LIBOR, each tenor-currency pair has an associated panel of banks that submit rates on a daily basis. The panel banks are a collection of multinational banks meant to be representative of those banks participating in the interbank market. Each panel bank's submission is a response to the following question:[13]

> "At what rate could you borrow funds, were you to do so by asking for and then accepting inter-bank offers in a reasonable market size just prior to 11 am?"

Under the current mechanism, LIBOR submissions are not required to be based on actual transactions and no justification for the submission is needed. In some instances, no transactions may have occurred for a given bank in a given tenor-currency. It is suggested that banks use the available data that they have on the market and from their other transactions to come up with a rate if necessary. "Reasonable market size" is left intentionally undefined.

Rates are submitted every day between 11:00am and 11:10am. Submissions are unobservable until the LIBOR rate is calculated and subsequently released to the public. LIBOR is calculated by eliminating the top and bottom quartile of the submissions and taking the average of the remaining rates.[14]

Under the current mechanism, LIBOR manipulation generally occurred for one of two reasons. First, banks manipulated their reports to influence the LIBOR rate itself and profit from their exposure to LIBOR through derivative and loan products. This manipulation resulted in either overstating or understating a bank's borrowing cost, depending on the direction of their exposure. Banks' misreporting was timed to coincide with rate fixings on various products, as small changes in the benchmark could yield large profits. In one example, the FSA case against RBS cites the 'inappropriate request' from a derivatives trader to the submitter of the rate, "pls get a very very very low 3m and 6m today pls, we have rather large fixings!" (FSA,

---

ary 2013.

Richard Milne, 'Norway launches interbank rate probe,' *Financial Times*, 16 January 2013.

[13]Interestingly, up until 1998 the question was quite different: "At what rate do you think interbank term deposits will be offered by one prime bank to another prime bank for a reasonable market size today at 11am?" It is unclear whether there has been a significant change in LIBOR manipulation after the implementation of the revised question (note that the original LIBOR question is quite similar to the current EURIBOR question).

[14]Note that the other -IBOR rates vary somewhat in both the question asked and the aggregation procedure for the reference rate.

2013).[15] A rough calculation (shown in Section 3) suggests that a given bank may make as much as $5m by manipulating LIBOR by even a basis point.

Second, banks misreported their borrowing rates to manipulate counter-parties' perceptions of the strength of the bank. This manipulation resulted in understating a bank's borrowing cost. As the rates submitted by individual banks were made public immediately by the BBA, these submissions potentially signaled the banks' credit risk in absolute terms and relative to other banks. This generated incentives to change the rates submitted to appear safer, particularly during the height of the financial crisis. For example, the FSA final report against UBS (FSA, 2012b) cites a manager writing to employees, stating "It is highly advisable to err on the low side with [LI-BOR] fixings to protect our franchise in these sensitive markets." Signaling that UBS could borrow at a low rate was intended to provide counterparties with confidence in the bank's creditworthiness.

1.2   Related Literature

A few very recent papers discuss the evidence on LIBOR manipulation. Kuo, Skeie, and Vickery (2012) find that LIBOR tracks alternative measures of interbank borrowing costs (TAF, Fedwire, NYFR), but is lower at the height of the crisis despite the expectation that it would be higher. LIBOR is also less diffuse than these other measures. These results suggest some degree of manipulation. Similarly, Snider and Youle (2012) demonstrate in a model that rate manipulation should lead to the bunching of reports around the cutoffs for the interquartile range, and then show that the LIBOR data has this property. Finally, using an econometric model identified by banks' rank order of LIBOR submissions, Youle (2013) finds that manipulation may have downwardly biased the benchmark by up to 50 basis points.

Some of these papers examine possible reforms to LIBOR. Eisl, Jankow-itsch, and Subrahmanyam (2014) show that setting LIBOR equal to the median submission (the extreme version of a trimmed mean) lowers the benefits of manipulation. Youle (2013) also argues for using the median submission, and shows that it could reduce manipulation by up to 70%. Duffie, Skeie, and Vickery (2013) show that the absence of current transactions may be mitigated by using a sample window of previous transactions. Unlike these papers, however, we do not focus on the best way to aggregate reports into an accurate LIBOR measure. Instead, we focus on how to incentivize proper reporting in the first place and find a mechanism that eliminates manipulation entirely. This gives the LIBOR administrator flexibility to choose a desired

---

[15]The primary submitters of the LIBOR rates at RBS were money market traders.

aggregation function unconstrained by concerns around manipulation.

Both the Vickrey-Clarke-Groves (VCG) and d'Aspremont and Gerard-Varet (AGV) mechanisms can incentivize agents to truthfully report their private information.[16] These mechanisms do not work in the LIBOR setting, however, as they fail when there is more than one source of private information. In the LIBOR case, there is an important second source of private information: banks' exposures to LIBOR. Chen (2012) demonstrates that the AGV mechanism can solve misreporting incentives when all banks have known exposures in the same direction. We allow for exposures to be private information and to be in either direction for any bank, which makes the AGV and VCG mechanisms inapplicable. Moreover, the AGV and VCG mechanisms require banks' actual borrowing rates to enter their payoff functions; when applied to LIBOR reporting, these rates may not be in the payoff functions, leaving nothing to 'tie' bank's reports to the truth.

The mechanism design literature also includes a number of possible approaches when agents have some information over each other's type. In this type of problem, for example, the principal may wish to maximize social welfare, but does not know the individual agents' utility functions. Demski and Sappington (1984) show that while a traditional mechanism design approach generates a mechanism that can achieve the societal first-best and maximize the principal's payoff, this mechanism is plagued by problems of multiple equilibria. In fact, while "all play fair" is an equilibrium, in general, there is an "all cheat" equilibrium too. This latter equilibrium Pareto-dominates the first from the point of view of the agents.[17] This sets the stage for later papers that design more complex mechanisms, with a joint focus on both achieving the societal first-best *and* uniqueness.

Ma, Moore and Turnbull (1988) consider a similar problem, but generalize the setting such that types are imperfectly correlated. An agent's type thus gives that agent some noisy information over others' types. In this setting, the authors solve the multiple equilibrium problem of Demski and Sappington (1984) by expanding the agents' strategy set. If agents can report not only their own type, but also a probability with which the other

---

[16]The VCG mechanism is implemented in dominant strategies but is not budget balanced; the AGV mechanism is implemented in a Bayesian Nash Equilibrium but is ex-ante budget balanced (Mas-Colell, Whinston and Green, 1995).

[17]Consider the following LIBOR reporting mechanism. First, all banks report their rate. Second, any bank may accuse another bank of cheating. Finally, any bank accused of cheating is fined. In this simple mechanism, all banks reporting accurately and accusing any cheaters is an equilibrium. Demski and Sappington show, however, that it is also an equilibrium for all banks to lie and to ignore others' cheating. This latter equilibrium Pareto-dominates the first from the point of view of the banks.

agents are 'cheating,' then any non-preferred equilibrium can be eliminated by the principal. In similar work, Cremer and McLean (1985) show that an auctioneer may extract the maximum possible surplus from an auction if agents' types are correlated and the auctioneer knows the structure of this interdependence. Ma (1998) considers an analogous problem of hidden action. When agents' actions are unobserved by the principal but observed by one another, the first-best can be achieved *if* a lag exists between action and outcome.[18]

Moore and Repullo (1998) present a mechanism when agents' types are hidden from the principal, but known to each other. The authors design a mechanism in which a randomly chosen agent proposes a rule. All other agents may then either agree to this rule or contest it; a contestation consists of an alternate characterization of the rule. If the initial report is contested, a cost is levied on the initial reporter who must then choose between the two available outcomes. In equilibrium, the contestation may be structured such that the initial agent continues to select his/her initial characterization if and only if it was the correct rule.

All of these mechanisms are similar, in that they use the knowledge of agents to ensure that the correct rule is implemented, even as the principal does not know the exact specification of the rule itself. There are a number of problems in applying these approaches in a LIBOR reporting context, however. The mechanisms of Ma (1998), Cremer and McLean (1985), and Ma, Moore and Turnbull (1988) place too strict a requirement on the information of the principal. In these models, not only must the principal know the ex ante distribution of agents' types, the principal must also know the correlation *between* agent's types to a sufficiently accurate degree so as to condition payments accordingly. We approach the LIBOR reporting problem by assuming that the administrator has *no* information about banks' LIBOR exposures and borrowing rates; for this reason, our mechanism is closest to the work of Moore and Repullo (1998).

The work of Moore and Repullo, however, places too strict a requirement on the information structure between the agents to be applicable as is. In Moore and Repullo (as in the other models), every agent is assumed to have information over the types of all other agents. This is too strong an assumption in the interbank market, as it requires each bank to have information over all other banks' borrowing rates and their incentives to manipulate LIBOR.[19] For this reason, while our mechanism is most similar to Moore and

---

[18]This lag is necessary to allow the mechanism to operate before the outcome realizes; the lag also makes the problem very similar to one of hidden information.

[19]In the general case considered by Moore and Repullo, information that is distributed across the economy cannot be elicited.

Repullo, we must alter it substantially to apply in a LIBOR reporting context.

2 Assumptions and Definitions

In this section, we describe the banks' profit functions and how they depend on lending and LIBOR. Bank lending occurs first. The whistleblower mechanism (analyzed in the next section) takes place after banks have had the opportunity to lend to one another.

2.1 Payoffs from lending

Consider a set of n banks. Each bank j, where $j = 1, ..., n$, has a payoff function from lending:

$$\pi_j^{\Psi}(\Psi_j)$$

Where Ψ_j is the set of loans made by bank j to other banks on the panel. Ψ_j is a vector of length $n-1$ of ordered pairs, (q_k, r_k), $k = 1, 2, ..j - 1, j + 1, j + 2, ..n$. Each pair represents the quantity, q_k, lent at a rate, r_k, to the k^{th} bank. If there is no loan to bank k, we set $(q_k, r_k) = (0, \cdot)$. The set of loans is important to the mechanism in two ways. First, we will use the complete set of loans by non-j banks to bank j to define what the borrowing cost is for bank j. Second, it will describe the willingness to lend of banks, which will be necessary to understand how a contested report by a whistleblower can be verified or invalidated through a loan.

As a point of terminology, it is useful to define $\Psi_j^i(q', r')$ as the set of loans of bank j, Ψ_j, with the loan to bank i set equal to (q', r'). For example, then, consider the incremental payoff to a bank j for changing the terms of its loan to bank i to q'' at a rate r''. This is the payoff to the bank of the new loan book, $\Psi_j^i(q'', r'')$, less the payoff from the old loan book, Ψ_j, $\pi_j^{\Psi}\left(\Psi_j^i(q'', r'')\right) - \pi_j^{\Psi}(\Psi_j)$. This notation is necessary for describing the benefits (or costs) of considering different loans.

To allow the banks' payoff functions to be as general as possible, we impose only a single assumption on the functions themselves.

Assumption 1 *The profit generated by lending to a given bank is increasing in the rate charged to that bank:*

$$\pi_j^{\Psi}\left(\Psi_j^i(q, r')\right) > \pi_j^{\Psi}\left(\Psi_j^i(q, r)\right) \ \ \forall i \neq j; \forall q > 0; r' > r \tag{1}$$

Assumption 1 states that for a given loan amount, a bank makes larger profits by charging higher interest rates.

With this basic assumption, we make two definitions which will be useful in the analysis.

Definition 1 *The minimum rate, ρ_{ji}, at which bank j is willing to lend $\bar{q}$ to bank i is:*

$$\rho_{ji} = \arg\min_{r} \pi_j^{\Psi}\left(\Psi_j^i\left(\bar{q}, r\right)\right) \text{ such that } \pi_j^{\Psi}\left(\Psi_j^i\left(\bar{q}, r\right)\right) \geq \pi_j^{\Psi}\left(\Psi_j^i\left(0, \cdot\right)\right) + \epsilon \quad (2)$$

where ϵ is an arbitrarily small profit level that must be generated to incentivize making a loan.

Definition 2 *The lowest rate, $\underline{b_i}$, at which a given bank i could borrow an amount, $\bar{q}$, is as follows:*

$$\underline{b_i} = \min\left\{\rho_{1i}, ...\rho_{(i-1)i}, \rho_{(i+1)i}, ...\rho_{ni}\right\} \quad (3)$$

Definition 1 describes the lowest rate at which a bank is willing to make a loan to another bank for a fixed loan amount. Definition 2 denotes $\underline{b_i}$ as the lowest rate that *any* bank would be willing to charge to bank i for a fixed loan amount. Formally, $\underline{b_i}$ is a function of the size of the fixed loan amount, $\bar{q}$, though for simplicity we drop the function notation.

By Assumption 1 and Definitions 1 and 2, no bank can profitably lend an amount $\bar{q}$ to a bank i for a rate less than $\underline{b_i}$:

$$\pi_j^{\Psi}\left(\Psi_j^i\left(0, \cdot\right)\right) + \epsilon > \pi_j^{\Psi}\left(\Psi_j^i\left(\bar{q}, r\right)\right) \ \ \forall j \neq i; \ r < \underline{b_i} \quad (4)$$

The above definitions provide a framework to analyze banks' willingness to loan. This is important for the mechanism, as at the time at which the mechanism is run, a given bank may or may not have actually borrowed any money in the inter-bank market. A bank that reports its borrowing cost may not have a transaction to present or may strategically wish to hide a transaction. In cases where there are no transactions presented, the mechanism instead uses the willingness of other banks to lend to the bank in question.

The borrowing cost is defined as the rate paid on the marginal loan. We assume that banks borrow first at the cheapest rate available, and then 'move up' the supply curve to progressively more expensive sources. Therefore, if a given bank has not borrowed any money in the inter-bank market, the

marginal loan is the *lowest* rate at which bank i could borrow, $\underline{b_i}$. If bank i cannot borrow at any price, then the offered rate is defined as infinite.[20]

When a given bank has borrowed, the borrowing cost is once again defined as the marginal loan: the *highest* cost transaction for bank i.[21] Denote the highest rate at which a bank has actually transacted as b_i^t. By definition, $b_i^t \geq \underline{b_i}$, because of the impact of bargaining power. We assume that transactions are costless to present (to the administrator) and are verifiable.

Having specified banks' willingness' to lend and transactions, we now define the borrowing cost b_i^* of a bank i.

Definition 3 *The borrowing cost b_i^* for a bank i is defined as the highest cost transaction b_i^t as long as a transaction exists, and $\underline{b_i}$ otherwise.*

As the incentive to make a loan depends on a bank's current loan portfolio, and the banks' loan portfolios are presumably different, banks will all differ in their willingness to lend to a given bank. For clarity, we make the following a formal assumption.

Assumption 2 *For each bank, i, there is only a single bank willing to lend at $\underline{b_i}$.*

Next, in describing the information structure of the mechanism, we must first define the beliefs of the various panel banks.

Definition 4 *Every bank i has a belief σ_i over the borrowing cost, b_j^*, of all other banks $j \neq i$. Formally, σ_i is an $(N-1)$-dimensional probability distribution in $\mathbb{R}_{\geq 0}$.*

We make one final assumption on the information structure:

Assumption 3 *The lowest rate at which a bank can borrow is known by at least two other banks.*

This is the key informational assumption and allows the mechanism to work when there are no transactions. While we allow for exposures to LIBOR to be private information for each bank, we assume that the lowest rate at which a bank can borrow is observed by at least two other banks (and the

---

[20]Any other interpretation questions the actual existence of the unsecured loan market, which we cannot address.

[21]This may provide an incentive for bank i to hide its past transactions, which will be taken into account in the mechanism.

bank itself). Formally, this implies that for a given bank i, there exist (at a minimum) two other banks j and k whose belief about i's borrowing cost places all probability on the correct borrowing cost, $\underline{b}_i$. Uninformed banks, for simplicity, are assumed to have a prior that places non-zero probability on the full support of possible borrowing costs.[22]

It is natural (although not necessary) to assume that one of the informed banks is the bank that actually offers the lowest rate, which means that there only needs to be one other bank that observes the lowest rate. In practice, banks have shown themselves to be able to observe, to at least some degree, the borrowing rates of their competitors. This is likely the case in a competitive environment where several banks are offering loans. Banks will generally have very accurate information about the borrowing of other banks, though we require only Assumption 3 for the proper functioning of our mechanism.

2.2   Payoffs from LIBOR

The whistleblower mechanism and LIBOR determination take place subsequent to lending. If the administrator were to have full information, LIBOR would be calculated using the borrowing rate for each bank that borrowed, and/or the rate at which the banks could borrow (having chosen not to). Thus, with $b^* = \{b_1^*, b_2^*, ..., b_n^*\}$, the LIBOR rate, L, is calculated as follows:

$$L = \mathbb{L}\left(b^*\right)$$

This is the LIBOR rate that would be reported absent any market frictions or manipulation.

As the vector of true borrowing rates, b^*, is hidden from the regulator, however, each bank i instead sends a report s_i where $i \in \{1, ..n\}$. We define the vector $s = \{s_1, s_2, ..., s_n\}$. The goal of the mechanism is $s = b^*$, so that $L = \mathbb{L}\left(b^*\right) = \mathbb{L}\left(s\right)$.

We do not specify the report aggregation function $\mathbb{L}\left(\cdot\right)$, as this will not be necessary for our mechanism to work. As mentioned above, the current LIBOR rate is calculated using a trimmed mean. We could incorporate this, or any other suggested function.[23]

---

[22]Our results are unchanged with respect to the priors of the uninformed banks. If uninformed banks have more precise priors, such that with some probability they can detect an incorrect submission by a reporting bank, then they simply serve as a de facto larger population of informed banks.

[23]The choice of aggregation function will quantitatively impact Equation P3 in Proposition 1, but the mechanism itself can accommodate any desired function.

Suppose that there are n banks in the LIBOR panel. We define the payoff from a given LIBOR rate L for bank j, where $j = 1, ..., n$, as:

$$\pi_j^L(L)$$

The bank's payoff function depends on LIBOR through its derivative and loan exposures to LIBOR. As this exposure may vary in magnitude and direction across banks, we place no assumptions on how a given bank's payoff depends on LIBOR. Nevertheless, this dependence drives the bank's incentive to manipulate LIBOR.

Notice that there is no argument in the bank's payoff function for the accuracy of the bank's LIBOR submission. Similarly, there is no argument in the payoff function for the actual rate at which a bank borrowed (if it borrowed). Allowing for either of these would make our proposed mechanism simpler, as it would be easier to provide incentives for truthful reporting.

For now, we exclude the incentive for a bank to manipulate its rate to hide credit risk and focus on the incentive to profit from derivatives exposures. In Section 5, we include the incentive to hide credit risk (thus adding a term for the reported rate to the utility function), and show that the mechanism that we propose easily withstands this manipulation incentive as well.[24]

A bank's payoff from LIBOR occurs subsequent to its lending in the interbank market; we therefore assume that the payoffs from LIBOR and its loans are additively separable.

$$\pi_j^\Psi(\Psi_j) + \pi_j^L(L) \tag{5}$$

The mechanism requires one further assumption. We assume that the time between the submission of reports and the possible verification of these reports through the mechanism is sufficiently short such that banks do not reoptimize their loan books.[25] This implies that in the few minutes over which the mechanism takes place, the willingness of a bank to lend to another will not shift materially. In the majority of situations, this assumption is appropriate. In situations where there is a material change in the strength of banks during this time, the initial LIBOR reports are definitionally out-of-date. Our mechanism is imperfectly able to deal with this regime change,

---

[24]Alternatively, the incentive to manipulate reported rates for hiding credit risk can be suppressed by hiding bank reports from the public for some period of time. This was suggested in the Wheatley Review and has been in place since the beginning of July 2013. This would work equally well in our mechanism.

[25]Theoretically, this could be solved by having banks input their entire strategy at once, though this could present practical difficulties.

but then again, *any* non-instantaneous mechanism is vulnerable to structural breaks of this sort.

3   The Mechanism

We now describe our proposed mechanism.

The mechanism uses as its starting point a very similar question to the one asked to the current LIBOR panel:

> "At what rate could you borrow funds of size $\bar{q}$ just prior to 11AM?"

The original LIBOR question states that the loan should be of reasonable market size. This reasonable market size is left to judgment, which may adjust for fluctuating markets and may vary for tenor-currency pairs.[26] The question in our proposed mechanism pins down a level, $\bar{q}$, which may be adjusted.

The mechanism consists of communications, payments and loans. The term w represents the whistleblower payment, the amount that a bank receives for correctly contesting a report. The terms ϕ_r and ϕ_c represent punishment fees for inappropriate reporting and contesting, respectively. The variables $\bar{q}$, w, ϕ_r, and ϕ_c are all positive constants.

Panel banks can take on specific roles in the mechanism, which are defined as follows:

- **Reporting bank**: a reporting bank i reports (1) the rate at which it can borrow, s_i, and (2) either from whom it can borrow at s_i, t_i, or presents a transaction.

- **Offering bank:** the bank j willing to lend at the lowest rate to reporting bank i (at $\underline{b_i}$).

- **Informed bank**: a bank that knows the lowest rate at which the reporting bank could borrow.

- **Named bank:** the bank t_i that the reporting bank i claims it can borrow from, if the reporting bank makes such a claim.

---

[26]For example, in the Commodities Futures Trading Commission case against Barclays, it says, "Prior to the financial crisis period, Barclays' U.S. Dollar LIBOR submitters considered "reasonable size" for a transaction by Barclays in the London U.S. Dollar money markets to be \$250-\$500 million. During the financial crisis period a reasonable sized transaction was significantly lower" (Commodities Futures Trading Commission, 2012).

- **Contesting bank**: a bank that contests the report of a reporting bank.

- **Bystander bank**: a bank in the panel that has no role in the mechanism (and is not informed).

We present the mechanism as succinctly as possible, and then describe the most important details in the following pages.

The **whistleblower mechanism** is as follows:

1. (Reporting Stage)
   Simultaneously, each bank $i \in \{1, 2, ..., n\}$ reports its own borrowing rate, s_i. Each bank also either presents a transaction in which it borrowed an amount of at least $\overline{q}$ at the rate s_i, or names a bank that would be willing to lend to it at this rate, $t_i \in \{\{1, 2, ..., n\} \setminus i\}$.

2. LIBOR is set, $L = \mathbb{L}(s)$, where $s = \{s_1, s_2, ..., s_n\}$.

   The following steps are run for all panel banks i simultaneously, $i \in \{1, 2, ..., n\}$.

3. (Contestation with executed transactions)
   If bank i presented a transaction in stage 1, any non-i bank that lent at least $\overline{q}$ to bank i at a rate greater than s_i may present the transaction. If bank i did not present a transaction in stage 1, any non-i bank that lent at least $\overline{q}$ to bank i at any rate may present the relevant transaction. Non-i banks act simultaneously.

4. If any transactions are presented, either by bank i in stage 1 and/or non-i banks in stage 3:

   (a) The 'true' rate, $\hat{b}_i$, is the highest rate across all transactions.
   (b) If the 'true' rate $\hat{b}_i \neq s_i$, then:

      i. A randomly selected bank that presented a transaction at $\hat{b}_i$ is paid w.
      ii. Bank i is fined $w + \phi_r$.

   (c) The mechanism ends for bank i's report.

5. (Contestation without executed transactions)
   If no transactions are presented:

(a) Any non-i bank k may contest the report of bank i by reporting $\tilde{s}_i < s_i$. Any non-i, non-t_i bank k may contest the report of bank i by reporting $\tilde{s}_i > s_i$. Contestations occur simultaneously and are observable to the panel.[27]

(b) (Verification Stage)
   If there is at least one contestation:

   i. Consider the set of $m \geq 1$ contestations, $\tilde{s}_i$ and the one initial report, s_i. For the m lowest values in this set, all non-i, non-contesting banks may state willingness to lend an amount $\overline{q}$ to (synthetic)[28] bank i at any of these rates.

   ii. The lowest of the m rates at which banks are willing to lend (or, if no bank will lend at any of them, the remaining, highest rate) is the 'true' rate, $\hat{b}_i$.

   iii. One bank (randomly chosen) from those that contested $\tilde{s}_i$ equal to the 'true' rate $\hat{b}_i$ receives w.

   iv. Any bank that contested at any rate not equal to the 'true' rate $\hat{b}_i$ is fined ϕ_c.

   v. If bank i did *not* report the 'true' rate, it is fined an amount $w + \phi_r$.

   vi. With equal probability, one of the m rates is chosen. If no bank expressed willingness to lend at this rate, the mechanism ends. Otherwise, one bank, selected randomly from the set of banks that expressed willingness to lend at this rate must lend $\overline{q}$ to (synthetic) i at the rate.

We depict the mechanism in Figure 1.

Next, we give a brief overview of how the mechanism works. We then present Proposition 1, which proves that given a positive whistleblower payment and a condition on the size of the punishment fees, ϕ_r and ϕ_c, truthful reporting is the unique pure-strategy equilibrium of the mechanism.

In the mechanism, each bank i reports the rate s_i at which it can borrow an amount, $\overline{q}$. At the same time, it either presents a transaction where it borrowed at this rate, or states that this is the rate at which it *could* borrow from a bank t_i. These reports are shown to all banks in the panel and the LIBOR rate is fixed. This is the reporting stage. While any initial

---

[27]The reason why the t_i bank can contest only by naming a higher rate (and not a lower rate) is discussed after Proposition 1 in the main text and in detail in the Appendix.

[28]A bank that states its willingness to lend and gets called up on to do so by the mechanism will lend an amount $\overline{q}$ to the administrator, who repays with interest if and only if a similar loan to bank i would have been repaid. We discuss this in detail below.
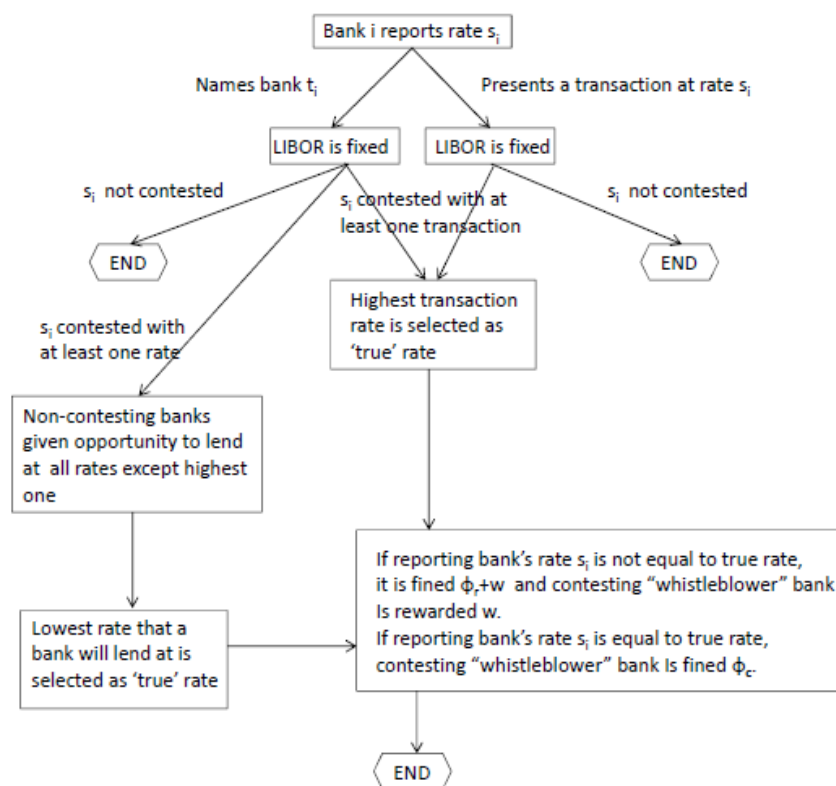
Figure 1: Overview of the Whistleblower Mechanism

report may subsequently be discovered to be false, it is *not* replaced by the 'true' rate uncovered by the mechanism. The reason for this upfront fixing of LIBOR rate is to eliminate incentives for the panel banks to contest or verify/invalidate rates solely for the motive of manipulating the LIBOR rate.[29]

Next, in the contestation stage, any other bank may present a transaction that it had with the reporting bank, if disclosing this would disprove the report of that bank. The 'true' borrowing rate of the bank is the highest rate at which it borrowed: this is interpreted as the marginal borrowing rate of the bank. If a bank that reported it has not borrowed is shown to have borrowed, or a bank is shown to have borrowed at a higher rate than it disclosed, it is fined. The non-reporting bank that presented the

---

[29]The early fixing does present the possibility that a bank misreports, knowing that it will be punished. This provides the incentive-compatibility constraint: the punishment for misreporting must be sufficiently high to discourage this. This is discussed more fully in Proposition 1.

transaction with highest loan rate (chosen randomly if there is a tie) receives the whistleblower payment.

Any report that was not based on a verifiable transaction (and has not been disproved by presentation of one) may be contested. A contestation consists of stating a rate $\tilde{s}_i \neq s_i$. Any bank may contest the reported rate, except that the bank t_i may not contest that the reported rate is too low. The reason for the named bank, t_i, and the restrictions placed on it, are left until the end of the mechanism description.

If no contestations are received, the mechanism ends. If contestations are made, lenders are sought at the m lowest of the $m + 1$ reported and contested rates. If lenders can be found at rates lower than that reported, or if no lender can be found at the reported rate, then the bank's report is shown to have been false, and the bank is fined.

Finally, one of the m rates is chosen randomly. If no banks expressed willingness to lend at this rate, the mechanism ends. Otherwise, one of these banks willing to lend is chosen randomly. This bank will 'synthetically' lend an amount $\bar{q}$ to the reporting bank. It would be difficult for an administrator to force a reporting bank to take on a loan, so we envision the administrator taking on the loan, which is why we call the loan a synthetic one.[30] The administrator of LIBOR demands a loan of $\bar{q}$ which is repaid with interest if and only if a similar loan to bank i would have been repaid. This would therefore be payoff-equivalent to lending to the reporting bank itself, and could be achieved through some sort of escrow account.[31] This concludes the final phase of the mechanism, the verification stage. We will demonstrate that as loans do not occur in the unique pure-strategy Perfect Bayesian Equilibrium of the mechanism, it is unlikely that the administrator would need a large facility to handle loans.

At the beginning of the mechanism, a bank is named by the reporting bank as a 'willing lender.' The role of this part of the mechanism is not to learn which bank made the best offer to lend; rather, it prevents the offering bank from falsely contesting that a correctly-reporting bank's borrowing cost is too low. The offering bank could then receive the whistleblower payment as no other bank will step in to invalidate the contestation. Therefore, in equilibrium, the reporting bank will indeed name the bank from which it can borrow to prevent this false contestation from occurring.

---

[30]Note that the administrator does not have to force the lending bank to make the loan, as the lending bank will find it strictly beneficial to do so. This is true due to a revealed preference argument. We discuss this in more detail below.

[31]Another potential way to structure this exposure would be to have the lender sell a credit default swap on the reporting bank to the administrator, with the payment adjusted accordingly. This would mean less cash out initially.

In summary, the mechanism augments the transaction-based approach of the Wheatley and FSB Reviews with a revealed preference argument to allow LIBOR rates to be calculated for currencies and tenors at which trading is thin. Rather than *requiring* transactions for the calculation of LIBOR (as the Wheatley Review does and the FSB Review recommends), our mechanism requires the use of transactions insofar as they exist, but is robust to situations where no trading occurs. Notice that even for markets where trading is heavy, some reporting mechanism is needed as reporting banks may have incentives to hide one or more transactions (if they hide all transactions, they actually may claim that trading is thin).

We provide some simple examples of how some of the details of the mechanism work in Section 3.1.

We solve for the pure-strategy Perfect Bayesian Equilibrium of the model, as defined in Fudenberg and Tirole (1991). We show that the whistleblower mechanism incentivizes truthful revelation of borrowing costs by all banks.

Proposition 1 *Given assumptions 1, 2, 3, and*

$$w > 0 \tag{P1}$$

$$\phi_c > 0 \tag{P2}$$

$$\pi_i^L \left(\mathbb{L}(s_i, b_{-i}^*) \right) - \phi_r - w < \pi_i^L \left(\mathbb{L}(b^*) \right) \quad \forall s_i \tag{P3}$$

the unique pure-strategy Perfect Bayesian Equilibrium outcome of the whistleblower mechanism is truthful reporting for all banks, $s = b^*$.

Proof. *See Appendix.* ∎

Proposition 1 states that as long as the whistleblower payment, w, is strictly positive, the punishment for incorrectly contesting, ϕ_c, is strictly positive, and Equation [P3] holds, then there does not exist any profitable deviation from the truthful reporting of a bank's borrowing costs. The intuition for Equation [P3] is that the punishment imposed on a misreporting bank must be sufficiently large that profit cannot be made by misreporting and being discovered.

We relegate the proof of Proposition 1 to the Appendix. However, we provide a brief sketch of the proof here. The proof of uniqueness proceeds in five parts.

First, we describe the equilibrium strategies. These are that (i) banks report truthfully, (ii) non-offering banks contest any reports known to be inaccurate, and (iii) banks verify correctly in the verification continuation game - that is, they offer to make loans at rates only at which they are

actually willing to lend. Second, we illustrate that these strategies, coupled with the appropriate beliefs, constitute a Perfect Bayesian Equilibrium. With existence proved, we next prove the uniqueness of this equilibrium.

Third, we prove that any sequentially rational strategies of the continuation game following on from a correct report cannot include any contestation. Effectively, any contestation is necessarily fallacious, and this is confirmed during the verification stage. Fourth, we prove that any sequentially rational pure strategy of the continuation game following on from an *incorrect* report must result in the discovery of the incorrect report and the punishment of the reporting bank.

Fifth, and finally, we combine these results. Given that the constraint Equation [P3] is satisfied, it is never optimal for a bank to manipulate its report if it will be discovered and punished. Since banks that manipulate their reports are always discovered in equilibrium, and banks that report correctly are never punished, there cannot exist a pure-strategy equilibrium involving report manipulation. Therefore, the unique pure-strategy equilibrium is proper reporting.

Note that when banks are found to have misreported, the imposed fine is simple: a flat fee of $w + \phi_r$. So long as ϕ_r is large enough, there is no possible deviation. A fine ϕ_r that satisfies Equation [P3] may be quite large, however, as it needs to work for all inappropriate reports.[32] This can be modified easily, by adding another punishment that is proportional to the degree of the LIBOR manipulation, such as $\kappa \left| \mathbb{L}(s_i, b^*_{-i}) - \mathbb{L}(b^*) \right|$, where κ is a price for distortion of the LIBOR calculation. This would allow banks that slightly misrepresent to be fined less than those that do so to a larger extent, and may be a more just approach. We use this modification in the calibration in Section 3.1.

As shown by Proposition 1, the mechanism incentivizes truthful disclosure from the various banks. There is a subtle issue with the mechanism, though, in that the reports from banks with actual loans and banks without loans may be slightly different, i.e. $b^t_i \geq \underline{b_i}$. Banks that have actually borrowed money report the rate at which they borrowed money, which is impacted by the bargaining power of the involved parties. Banks that have not borrowed anything, however, submit reports that are unimpacted by bargaining power. When the borrowing banks' bargaining power approaches 100%, the two are the same. In any situation, given the competition in the unsecured lending market, we expect any variation caused in LIBOR by changes in bank

---

[32]For example, if a given bank can manipulate its report by up to 50bps, under the current LIBOR trimmed mean, this could necessitate a ϕ_r on the order of $25m.

bargaining power to be minor.[33,34]

3.1   Examples

In this subsection, we consider three examples that display how the mechanism works. We also calibrate the mechanism to quantify how large a punishment fee would be needed for deterrence.

Possible overstating: Suppose that there are no transactions presented by the reporting bank or by any contesting bank. The reporting bank reports that its borrowing rate is 2%, and one bank contests, claiming that the rate is 1%. The other panel banks are given the chance to lend at 1%. If at least one does, 1% is the true rate, the contesting bank gets w, and the reporting bank is fined $w + \phi_r$. One of the banks that expressed willingness to lend is randomly selected to make the loan. If no banks are willing to lend at 1%, then 2% is the true rate, and the contesting bank is fined ϕ_c. Note that as the LIBOR rate is fixed just after the initial reports, the contesting bank and the lending bank cannot influence it through their actions.

Possible understating and the role of the named bank: Suppose that there are no transactions presented by the reporting bank or by any contesting bank, and the bank named by the reporting bank is bank A. The reporting bank reports that its borrowing rate is 2%, and one bank contests, claiming the correct rate is 3%. The other panel banks are given the chance to lend at 2%. If at least one does, 2% is the true rate, and the contesting bank is fined ϕ_c. If none do, 3% is the true rate, the contesting bank gets w, and the reporting bank gets fined $w + \phi_r$.

By the rules of the mechanism, bank A cannot be the bank that contests. Suppose that this rule was not in place and that bank A could contest that the report was understated (and contest with, for example, 3%). There would be no bank willing to lend at 2%, even though 2% was the actual rate at which the reporting bank could borrow. Bank A would then receive the whistleblower payment w. For this reason, the mechanism allows the reporting bank to stop Bank A from contesting that the report is too low. Notice that this blocking procedure is not needed to stop Bank A from contesting that the report is too high, as Bank A knows that it will be caught in this case (no bank will lend at the low rate it contests is the true rate).

---

[33]Note that this is an issue both for the version of LIBOR in current use, and for nearly every proposed alternative. Any use of loan transactions necessarily exposes LIBOR to random variations caused by differing levels of bargaining power.

[34]This problem also diminishes as the size of the panel increases, in which case we would expect $b_i^t \to \underline{b_i}$. We discuss below other benefits of a larger panel for our mechanism.

Truthful revelation of willingness to lend: Here we demonstrate that there is no reason for the offering bank to strategically state its willingness to lend. Suppose that there are no transactions presented by the reporting bank or by any contesting bank. The reporting bank reports that its borrowing rate is 2%, and two banks contest, claiming that the proper rate is 1% and 3%, respectively. The mechanism asks whether any bank will lend at either 1% or 2%. Suppose that the offering bank is the only bank willing to lend and it is willing to lend at both 1% and 2%.[35] If the offering bank states it is willing to lend only at 2%, then the mechanism calls on it to do so with probability $\frac{1}{2}$ (and with probability $\frac{1}{2}$ there is no lending at 1%). If the offering bank states it is willing to lend at both 1% and 2%, then the mechanism calls on it to lend at 1% with probability $\frac{1}{2}$ and to lend at 2% with probability $\frac{1}{2}$. As should now be clear, the offering bank's statement does not affect the probability that a rate will be selected for a loan, and therefore it does not act strategically. As it prefers to lend at both rates, it will state that truthfully.[36]

Calibration: As discussed earlier, consider an additional component for the misreporting punishment, $\kappa \left| \mathbb{L}(s_i, b^*_{-i}) - \mathbb{L}(b^*) \right|$, where κ is a price for distortion of the LIBOR calculation. To calibrate the punishment for misreporting used in the mechanism, consider that a bank will have a maximum transaction or set of transactions that has the floating rate fixed at 3 month LIBOR worth \$200bn.[37] The exposure is then $200[(1 + r)^{\frac{1}{4}} - 1]$, where r represents 3m LIBOR. Taking the derivative, the change in exposure with a change in the interest rate is $50(1 + r)^{-\frac{3}{4}}$. When $r = .03$, the change in exposure for a 1 basis point change is \$4.89m. When $r = .003$, the change in exposure for a 1 basis point change is \$4.99m. The approximate gain from a 1 basis point change in LIBOR is thus about \$5m. Thus, $\kappa > \$5$m per basis point is a sufficient restriction for 3m LIBOR (one of the most liquid tenors). Note that this is larger than an estimate from a transaction in an email cited by the U.S. Department of Justice (2012) in their case against Barclays. In the email, a trader states that, "[F]or every 0.25 bps tomorrows [sic] fix is below 4.0525 we lose 154,687.50 usd," which is about \$.6m for one basis point.

---

[35]If it is not willing to lend at 1%, there is no reason for it to state that it would, because with positive probability it would be forced to do so.

[36]If the mechanism were modified such that a loan was always demanded at the 'true' rate, this would provide banks with an incentive to misrepresent their willingness to lend. Banks would feign a lack of willingness to lend at low rates, in an attempt to increase the rate at which they lent money.

[37]Simone Foxman, "How Barclays Made Money on LIBOR Manipulation," *Business Insider*, 10 July 2012.

Considering many of the recent LIBOR manipulations to be of roughly 0.5 basis points magnitude, our estimate would result in an approximately $2.5m fine per incident. This suggests a roughly similar order-of-magnitude to (though still five times larger than) the $940m in fines assigned to UBS for roughly 2000 "requests for inappropriate submissions."[38]

3.2 Discussion

One strength of our mechanism is that it can be applied even to tenors where transactions are generally infrequent or are temporarily so because of extreme market events. By using a revealed preference approach to verify the accuracy of contested reports, the measure can remain up-to-date, even when transactions are not occurring. This removes the need to rely on old estimates of borrowing costs, as might otherwise be required.

The mechanism keeps separate the verification of reports and the calculation of LIBOR. Once reports are sent, the LIBOR rate is calculated, and any further contestation and verification of the individual reports submitted does not impact the final rate. This implies, paradoxically, that even if a LIBOR rate is known to be incorrect (because certain of the submitted reports have been found to be fallacious), it remains as is. As demonstrated, this ex-post inaccuracy (which is off the equilibrium path) leads to greater ex-ante incentives to submit accurately, because the LIBOR calculation does not influence banks' choices in the contestation and verification stages.

Standard mechanisms from the literature cannot implement truthful reporting in this environment. A VCG mechanism or an AGV mechanism will not work in this model, as banks' exposures to LIBOR are private information both in terms of direction and magnitude. Furthermore, the actual borrowing rate that is being solicited is semi-private information and does not appear in the reporting bank's payoff function – a necessity for those mechanisms.[39] And while our mechanism has much in common with the Subgame Perfect Implementation of Moore and Repullo (1988), they do not allow for asymmetric information between agents as we do.

From an implementation standpoint, the mechanism described here has a large advantage over several other mechanisms: truthful reporting is the unique pure-strategy equilibrium. First, this implies that along the equi-

---

[38]Jill Treanor, "Two former UBS employees charged in US over LIBOR," *The Guardian*, 19 December 2012.

[39]Chen (2012) proposes an AGV mechanism. This is possible because in his model, exposures are known and are all in the same direction and the borrowing cost of the bank enters the payoff function through a quadratic reputation cost. As we note in the introduction, adding reputation costs would only make our mechanism easier to implement.

librium path, the mechanism is budget-balanced. A simple restriction on the allowable range of possible contestations ensures that the mechanism is budget-balanced even off the equilibrium path. Consider the following - when a contestation occurs, the administrator generates a cash inflow of at least the minimum of ϕ_c and ϕ_r, less any amount lost on the acceptance of a synthetic loan. So long as the excess interest that the administrator pays on the synthetic loan is less than the earnings on the punishment fine, the mechanism is budget-balanced even off the equilibrium path. This is a significant improvement over the ex-ante budget balance of the AGV mechanism, which can often result in ex-post budget concerns.

Second, while the mechanism is at times complex, the complexity is entirely off the equilibrium path. In equilibrium, banks simply report the rates at which they could borrow accurately, and no contestation or verification occurs.

The mechanism implicitly assumes that the best loan offer to a given bank is from another bank within the panel. This implies that the panel must not be too small.[40] However, the mechanism can easily be modified to ensure that the panel is sufficient. One possibility is to include all prime banks and rephrase the LIBOR question to report only on how much it costs them to borrow from other prime banks. The current EURIBOR panel and the LIBOR panel up until 1998 only asked banks to consider borrowing costs from prime banks.[41] A second is to allow *any* bank (including banks not in the panel) to lend in the verification stage. This could be an issue if by revealing the panel banks' credit risk, it altered their incentives, but Section 5 demonstrates that the mechanism can easily take this into account.

4   Collusion

In this section, we consider possible collusion between the various banks. First, we show that collusion is not possible in the model. Second, we discuss the possibility of collusion in a repeated game framework.

---

[40]The Wheatley Review proposes a larger panel size, but for a different reason: to decrease the incentive to manipulate.

[41]However, these questions are/were more vague than the current LIBOR question, asking how much it costs one prime bank to borrow unsecured from another prime bank, rather than asking how much it costs *your* prime bank to borrow unsecured from another prime bank.

4.1 Collusion in the model

Collusion is defined as an agreement by at least two banks that is not enforceable in a court of law. In the appendix, we show in detail that no such collusive agreement could exist in the model for a simple reason: there would always be a bank that wished to deviate when its turn to act was reached, i.e. collusion is not sequentially rational.

To briefly sketch the important points in the appendix, consider first that any reporting bank cannot sustain collusion with a bank that may contest or verify/invalidate a contestation. Once reporting banks have reported, the LIBOR calculation is fixed and they cannot credibly offer anything in the game contingent on a collusive agreement. Reporting banks cannot collude with each other, as they submit their reports simultaneously and only consider their individual incentives. Second, banks could try to collude to collect the whistleblower payment w; however this would involve collusion between a bank at the stage where reports are contested and a bank at the stage where contestations were verified/invalidated. The bank at the later stage would always act in its own self interest.

There are gains to be had by colluding, but the precise structure of the mechanism prevents these gains from being realized.

4.2 Collusion in a repeated game

Collusion is potentially more concerning when the mechanism is extended to a repeated game model. Still, there are a number of mitigating factors that suggest that collusion will be very difficult to maintain in a repeated game.

In any finitely repeated game, collusion remains impossible. Banks would have no incentive to collude in the final period, unraveling any possible attempt at collusion. Collusion is thus only possible if the mechanism is repeated an infinite number of times.

When banks attempt to collude to manipulate LIBOR, they face both asymmetric information about each others' exposures to LIBOR and the fact that the exposures (and hence incentives) will vary over time. While banks' borrowing rates will be visible by some other banks in the industry, the same is not necessarily true of LIBOR exposures. Therefore, any bank that wishes to manipulate LIBOR in a given direction must first find another bank with an exposure in the same direction that is willing to collude. Even if another bank is found, the banks are unlikely to be able to collude in every period for the simple reason that the magnitude and direction of each bank's exposure changes over time. This reduces the benefits to collusion.[42]

---

[42]This assertion is supported by the action of Barclays, in whistleblowing on its fellow

Importantly, changes in banks' exposures to LIBOR also reduce the banks' ability to impose punishments to sustain collusion. A punishment phase of trying to manipulate LIBOR in a fixed direction to hurt another bank (i) may not be credible as the banks doing the punishing may have exposures that move in the opposite direction and (ii) may not hurt the target bank as its exposure may change as well, making this 'punishment' perversely benefit the target.

An alternative punishment would involve a first bank contesting the deviating bank's report and a second verifying the first bank's contestation with a loan. The punishing banks would have to ensure that a non-colluding bank does not interrupt this (or that there are no non-colluding banks), even though there may be explicit monetary incentives (the whistleblower payment) to do so. As the verifying bank loses money (and the contesting bank gains money) in this punishment phase, its credibility depends on whether the roles could be rotated among the banks that are punishing and whether the gains outweigh the losses.[43]

The most credible punishment would be to revert to the equilibrium from Proposition 1. In this case it will still be hard to reach a collusive agreement, as the colluding banks would have to derive a scheme to reveal their derivative exposures and communicate on how to collude given the set of exposures. The group of colluding banks would also need to include almost if not all of the panel banks. An administrator can therefore mitigate the likelihood of this by increasing the panel size, thus making coordination more difficult and the impact on the LIBOR calculation smaller. Setting the whistleblower payment sufficiently high also makes this agreement more prone to profitable deviations.

Of course, over a large number of periods there are also empirical techniques to detect whether submitted rates are accurate or not (at least in the aggregate). The wealth of other measures of the risk free rate, credit risk, and liquidity can also raise red flags as they did in the recent scandal.

In conclusion, then, while collusion in an infinitely repeated game is possible, there are a number of obstacles that banks would need to overcome.

---

banks. Barclays made the active decision *not* to collude with others (or at the very least, to stop colluding with them). Under our mechanism, whistleblowing is remunerated, increasing the incentives of banks to reveal extant collusive agreements.

[43]Also, from a practical standpoint, if such a punishment phase were ever entered, it would be likely to be noticed by the administrator and non-colluding institutions.

5   Incorporating credit risk signaling incentives

In the whistleblower mechanism, it is assumed that there is no incentive for banks to misrepresent their borrowing rate to hide their credit risk. This incentive was an important factor during the recent financial crisis and the recent LIBOR inquiries have found evidence that banks biased down their reported rates to appear safer than they were. In this section, we demonstrate that it is simple for the mechanism to incorporate the incentive to hide credit risk (and remain robust to this manipulation incentive). There are two ways to address this through the payoff functions of the banks.[44]

First, suppose that a bank's report enters its payoff function directly[45] if it is judged by the mechanism to be the true borrowing rate for the bank. If the report is found to be inaccurate, market participants disregard it. (LIBOR is still manipulated by a report found to be inaccurate, then, but it has no *signaling* value.) In effect, the report would affect the bank's payoff if and only if the bank's report was not contested and then found to be misrepresented. This approach would leave nothing changed from the main model. As there is only a benefit to manipulating the report if it is undiscovered, and Proposition 1 shows that an inaccurate report is always discovered, there remains no profitable deviation.

Second, it is possible that the report itself directly enters into the payoff function of the bank. It is unclear why exactly this would be the case - if a report is later found to have been untrue, does the report still serve any signaling value? Nonetheless, we will consider this situation and show that the proposed mechanism remains robust to manipulation for this reason.

Consider a more general payoff function for bank i :

$$\pi_i^L\left(L\right) + \pi_i^\Psi\left(\Psi_i\right) + \pi_i^s\left(s_i\right) \tag{6}$$

In this more general payoff function, L and Ψ_i are defined as earlier, while s_i is the rate report sent by bank i. To focus on misreporting to hide credit risk, we could assume that the derivative of π_i^s with respect to s_i is negative (the lower the perceived rate at which a bank can borrow, the lower

---

[44]As we mention earlier in the paper, an alternative method for minimizing the incentive of banks to misreport to signal low credit risk is to reveal their signals only to the other panel banks. This would reduce the proportion of the economy that witnesses the signal, and therefore mitigate the incentive to signal inappropriately. In this section, we consider alternative methods, within the mechanism itself, to ensure that banks' incentives to signal low credit risk do not result in an inaccurate LIBOR rate.

[45]The report is already in the payoff function indirectly through the LIBOR function, whether it is true or not.

its perceived credit risk). However, as will be seen, this assumption is not necessary.

Suppose that we implement the whistleblower mechanism. When considering possible deviations from the equilibrium path, it is clear that nothing changes at either the contestation or verification stage. Conditional on the reporting stage, the payoff functions are equivalent to those of previously: as there was no possible off-equilibrium path deviation in the primary model, there is none possible here.

The major consideration, then, is whether or not a bank may report improperly. In the main model, the only possible incentive for misreporting is in the impact on the calculated LIBOR rate. In this extension, this is no longer true. For example, with a trimmed mean calculation algorithm for LIBOR, a bank that wants to hide its credit risk may misreport its rate knowing that the report will have no impact on LIBOR (i.e. it is in the trimmed area).

Therefore, to ensure that banks *never* have an incentive to misreport the rate at which they can borrow, the mechanism requires the following to hold:

$$\pi_i^L\left(\mathbb{L}(s_i, b_{-i}^*)\right) + \pi_i^s\left(s_i\right) - \phi_r - w < \pi_i^L\left(\mathbb{L}(b^*)\right) + \pi_i^s\left(b_i^*\right) \quad \forall s_i \qquad (7)$$

Equation [7] follows analogously from Equation [P3]. When banks may misreport both for profit on derivative exposures and for credit risk signaling reasons, the punishment imposed for misreporting must be sufficiently high as to outweigh both of these benefits.

As before, however, it is possible that the necessary ϕ_r is so large as to be impractical. (This is again because a given punishment fine must be sufficient to discourage all possible misreports.) Then, a small modification to the proposed mechanism may be helpful to calibrate the size of the punishment: an extra term can be added to the punishment imposed on banks found to have misrepresented their rate. The punishment levied on misreporting banks should be as follows:

$$\tau\left|s_i - b_i^*\right| + w + \phi_r \qquad (8)$$

Where $\tau\left|s_i - b_i^*\right|$ is a punishment imposed for the difference between the signaled rate, and that which is determined to be the true rate in the verification stage. Thus, to ensure that misreporting is never optimal, the following constraint must hold for all possible misreports, $s_i \neq b_i$.

$$\pi_i\left(\mathbb{L}(s_i, b_{-i}^*), \Psi_i, s_i\right) - \phi_r - w - \tau\left|s_i - b_i^*\right| < \pi_i\left(\mathbb{L}(b^*), \Psi_i, b_i^*\right) \quad \forall s_i \qquad (9)$$

Equation [9] is simply Equation [7] restated under the expanded punishment. This allows banks that slightly misreport to be punished differently than those that significantly do so.

Again, our punishment mechanism does not account for any additional reputational punishment imposed by the media (which could be significant, given that the signaling incentive is driven by a desire to mislead the public). This would only strengthen our results (and decrease the required size of the punishment fine).

6 Conclusion

In this paper, we propose a robust LIBOR reporting mechanism that incentivizes banks to reveal their borrowing costs truthfully. For a bank's report, the mechanism solicits actual transactions when they exist. When there are no transactions, the mechanism solicits the borrowing cost from the other LIBOR panel banks' willingness to lend. It does this by allowing the panel banks to contest the initial report, and uses a revealed preference argument to incentivize other panel banks to lend to verify the true borrowing cost of the reporting bank.

Our approach has several benefits. First, truthful reporting is the unique pure-strategy equilibrium. Second, there is no cost to implementing the mechanism. Third, the mechanism works even when there are no actual transactions. Fourth, it takes into account private information that banks have about the magnitude and direction of their derivative exposures and the incentive of banks to under-report their borrowing cost to hide credit risk. Fifth, even in an infinitely-repeated game, rewards for whistleblowing and the coordination required to collude effectively suggest that collusion would be very difficult to maintain.

There are a few aspects of the mechanism that merit further study. Given the punishments involved, it would be worth learning how quickly agents could converge to equilibrium play in this mechanism. This could be done by running experiments. Also, while it is clear that the loan size $\bar{q}$ should vary depending on market conditions, empirical investigation is needed to set it appropriately.

References

[1] Australian Financial Market Association (2013), "AFMA BBSW: A Guide to the Bank Bill Swap Reference Rate," Sydney, Australia,

http://www.afma.com.au/afmawr/_assets/main/lib90012/afma%20bbsw%20guide.pdf

[2] Bank of Israel (2012), "Standardization of Telbor Quotations," Jerusalem, Israel: http://www.boi.org.il/en/Markets/TelborMarket/Documents/telbordef_eng.pdf.

[3] Chen, Jiakai (2012), "LIBOR's Poker: Interbank Borrowing Costs and Strategic Reporting", mimeo, UC Berkeley.

[4] Cremer, Jacques and Richard McLean (1985): "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," Econometrica, 53, 345-361.

[5] Commodities Futures Trading Commission (2012), "CFTC Docket Number 12-25", Washington D.C., http://www.cftc.gov/ucm/groups/public/@lrenforcementactions/documents/legalpleading/enfbarclaysorder062712.pdf

[6] Demski, Joel and David Sappington (1984): "Optimal incentive contracts with multiple agents," Journal of Economic Theory, 33, 152-171.

[7] Department of Justice (2012), "Appendix A, Statement of Facts", Washington D.C., http://www.justice.gov/iso/opa/resources/9312012710173426365941.pdf

[8] Department of Justice (2013), "Rabobank Admits Wrongdoing in Libor Investigation, Agrees to Pay $325 Million Criminal Penalty", Washington D.C., http://www.justice.gov/opa/pr/2013/October/13-crm-1147.html.

[9] Duffie, Darrell, Skeie, David, and James Vickery (2013), "A Sampling Window Approach to Transactions-Based Libor Fixing," Working Paper, Federal Reserve Bank of New York.

[10] Eisl, Alexander, Jankowitsch, Rainer, and Marti Subrahmanyam (2014), "The Manipulation Potential of Libor and Euribor," mimeo NYU.

[11] Financial Services Authority (2012a), "Final Notice; Reference Number 122702", London, United Kingdom.

[12] Financial Services Authority (2012b) "Final Notice; Reference Number 186958", London, United Kingdom.

[13] Financial Services Authority (2013) "Final Notice; Reference Number 121882", London, United Kingdom.

[14] Financial Stability Board (2014) "Reforming Major Interest Rate Benchmarks", Basel, Switzerland.

[15] Fudenberg, Drew, and Jean Tirole (1991), *Game Theory*, MIT Press: London, England.

[16] Kuo, Dennis, Skeie, David, and James Vickery (2012), "A Comparison of Libor to Other Measures of Bank Borrowing Costs," Working Paper, Federal Reserve Bank of New York.

[17] Ma, Ching-To (1988), "Unique Implementation of Incentive Contracts with Many Agents," The Review of Economic Studies, 55, 555-572.

[18] Ma, Ching-To, Moore, John and Stephen Turnbull (1988): "Stopping Agents from 'Cheating'," Journal of Economic Theory, 46, 355-372.

[19] Mas-Colell, Andreu, Whinston, Michael Dennis and Jerry R Green (1995), *Microeconomic Theory*, Oxford University Press, Oxford.

[20] Moore, John, and Rafael Repullo (1988): "Subgame Perfect Implementation," Econometrica, 56, 1191-1220.

[21] Snider, Connan, and Thomas Youle (2012), "The Fix is In: Detecting Portfolio Driven Manipulation of the LIBOR," mimeo, UCLA.

[22] Swedish Bankers' Association (2012). Framework for Stibor. http://www.swedishbankers.se/web/bf.nsf/$all/C2778E705D5E100AC1257ADB00375311.

[23] Wheatley, M. (2012). The Wheatley Review of Libor. http://cdn.hm-treasury.gov.uk/wheatley_review_libor_finalreport_280912.pdf.

[24] Youle, T. (2013). "How Much Did Manipulation Distort the Libor?, " mimeo, University of Minnesota.

7 Appendix

7.1 Proof of Proposition 1

The proof of uniqueness proceeds in five parts. First, we specify the equilibrium strategies and corresponding updating about beliefs. Second, we illustrate that the strategies and beliefs as specified constitute a Perfect Bayesian Equilibrium. Third, we prove that any sequentially rational strategies of the continuation game following on from a correct report by a reporting bank cannot involve contestation. Fourth, we prove that any sequentially rational

pure strategies of the continuation game following on from an incorrect report by a reporting bank must result in the incorrect report being discovered and punished. Fifth, and finally, we combine these earlier insights to conclude that any pure-strategy PBE of the entire game must result in truthful reporting. In summary, the specified strategies combine to form the unique Perfect Bayesian Equilibrium outcome.

7.1.1 Specification of equilibrium strategies and beliefs

The proposed equilibrium strategy for bank i, where $i = 1, ..., n$ is as follows:

- If funds have been borrowed, report $s_i = b_i^t$, and show proof of this (present the transaction).

- If no funds have been borrowed, report $s_i = \underline{b_i}$, and name the offering bank correctly.

- If any bank loans (transactions) to bank $j \neq i$ can be presented where the rate is larger than s_j, present them.

- For non-offering banks: if any bank $j \neq i$ is known to have reported inaccurately, contest the report (if possible), with $\tilde{s}_j = \underline{b_j}$.

  For the offering bank: if any bank $j \neq i$ is known to have under-reported, contest the report by $\tilde{s}_j = \underline{b_j}$. If it is possible to contest when the reported rate is correct, strategically contest the report by $\tilde{s}_j > s_i$.

- In any verification stage, state willingness to lend to any bank $k \neq i$ at any rate greater than or equal to ρ_{ik}.

It is also necessary to define the equilibrium beliefs.

- The beliefs of a bank i that is informed about the borrowing cost of bank j do not update conditional on the report of bank j.

- The beliefs of any bank i that is *not* informed about the borrowing cost of bank j update, conditional on j's submission s_j, such that all probability mass is on s_j. (Uninformed banks, as stated earlier, have priors that place non-zero probability on the full support of possible borrowing costs.)

7.1.2 Proof of Existence

We now prove that the equilibrium strategies and beliefs specified constitute a Perfect Bayesian Equilibrium. We first consider the beliefs. Then, we illustrate that given the belief structure, the strategies constitute a PBE.

Belief updating process

First, it is straightforward to show that the belief updating process satisfies the requirements for a PBE. Informed banks' beliefs do not update conditional on a report. Informed banks know the true borrowing rate of the reporting bank, and therefore the report cannot contain any additional information. Uninformed banks, on the other hand, initially place non-zero probability on all possible reports, and update to place all probability on the report of the reporting bank being accurate. As the reporting banks always report correctly in equilibrium, this belief updating is 'correct.'

Next, we show that the specified strategies constitute a PBE. We consider the three different stages of the mechanism in turn.

Sequential Rationality in the verification stage

At the verification stage, banks' strategies are defined such that the proof of sequential rationality is trivial. Banks offer to verify if and only if the profit generated by lending at a given rate would be strictly positive. As banks that have the opportunity to verify have definitionally not contested (otherwise they would not be allowed to verify), any strictly positive profit is preferred to the zero profit that would be generated otherwise. Additionally, LIBOR exposures are immaterial in the verification stage, as the LIBOR rate is already set and cannot be changed. Therefore, the specified strategies are sequentially rational and there exist no profitable deviations either on or off the equilibrium path.

Sequential Rationality in the contestation stage (after presentation of executed transactions)

There are two types of actions at the contestation stage. Banks may contest a reporting bank's report either through the disclosure of a transaction at a rate higher than the reported rate, or they may contest an unproven report with an assertion of what the true borrowing rate is. We consider these two types of actions in sequence.

Contesting by disclosing a transaction struck at a rate higher than the reported rate is impossible along the equilibrium path. This is because the reporting bank reports accurately, and so there therefore exist no transactions that could be disclosed to disprove the report. Even off the equilibrium path,

there exists no profitable deviation from the specified strategy.

The banks' strategy stipulates that a report is contested with a proven transaction whenever possible. The sequential rationality of this strategy is straightforward. If the executed transaction that is presented is the highest rate at which the reporting bank borrowed, then the revealing/contesting bank receives the whistleblower payoff, w. If there is another transaction revealed at a higher rate, then the actions of this particular bank are irrelevant - by revealing the transaction it neither receives the whistleblower payoff nor is punished. Therefore, the transaction-revealing strategy of the banks is sequentially rational.[46]

Next, we consider the strategy of contesting reports that have not been confirmed through the presentation of executed transactions. We show that the specified strategy is sequentially rational by illustrating that there exist no profitable deviations either on or off the equililbrium path.

Sequential Rationality in the contestation stage (after no executed transactions are presented)

In the contestation stage, there are four possible deviations. Banks may (falsely) contest an accurate report either by reporting that the original report was too high ($\tilde{s}_i < s_i$), or by reporting that the original report was too low ($\tilde{s}_i > s_i$). Or, banks may fail to contest a bank that has reported either too high ($s_i > \underline{b_i}$), or too low ($s_i < \underline{b_i}$). We consider each of these possible actions in turn.

First, consider the situation where the reporting bank reports its borrowing rate accurately, such that $s_i = \underline{b_i}$. Any other bank *not named* as the potential offering bank may contest this report by contesting that the report is too low. (To reiterate, the mechanism stipulates that the named bank t_i cannot accuse the reporting bank of reporting too low.) Thus, a contestation may be of the type, $\tilde{s}_i > s_i = \underline{b_i}$. Regardless of the total number of contestations, *part* of the verification stage will always consist of the other banks being asked if they are willing to lend $\overline{q}$ to the reporting bank at a rate s_i. If the offering bank does not contest this report (and is thus available to lend), then it will profit by lending at the reported rate. Therefore, regardless of the total number of contestations, it cannot be optimal for any non-offering

---

[46]It is possible that a bank could choose not to present an executed transaction, preferring the payoff from lending in the verification stage. This requires that the bank believes that (i) it is the only bank with an executed transaction, (ii) there exists a bank that will contest the bank's report in the second part of the contestation phase, and (iii) the present bank will be chosen to lend in the verification stage. As this still results in the reporting bank being discovered having misreported, it generates the same incentives not to misreport at the reporting stage, and is therefore effectively equivalent for our purposes.

bank to accuse an accurate report of being too low. A spurious contestation will be discovered, and the bank that raised it punished.

However, there *may* be an incentive for the offering bank to contest the report. If the offering bank is alone in its willingness to lend at $s_i = \underline{b_i}$, then it profits by (falsely) accusing. This is because the only bank willing to lend at s_i, itself, will not do so. Effectively, the offering bank is accusing the reporting bank of lying, and as it represents the only possible witness, its accusation is sustained. For this reason, the offering bank's strategy includes the contestation of an accurate report if it is not the named bank.

In summary, no non-offering bank can profit by accusing an accurate report of being too low.[47] The sole offering bank can profit, and this will provide the incentive for the reporting bank to correctly name its 'potential lender.'

Second, consider again if the reporting bank reports accurately. Any other bank (even that named by the reporting bank) may contest that the initial report was too high. Thus, a report may be of the type, $\tilde{s}_i < s_i$. For a bank to profitably lend at $\tilde{s}_i$, the following would need to hold for some bank x:

$$\pi_x^{\Psi}\left(\Psi_x^i\left(\overline{q}, \tilde{s}_i\right)\right) > \pi_x^{\Psi}\left(\Psi_x^i\left(0, \cdot\right)\right) + \epsilon \tag{10}$$

By Equation [4], however, because $\tilde{s}_i < \underline{b_i}$, there is no bank for which this holds. No bank can ever profitably lend to bank i at a rate less than $\underline{b_i}$. If any bank falsely claims that an accurate report is too high, then no bank exists that will lend at the rate reported in the contestation. Regardless of the number of banks that contest the given rate, then, any bank that accused the rate of being too high is proved to have contested falsely. These banks must pay the fee, ϕ_c. Thus no bank, whether it be an informed bank, the offering bank, or a bystander bank, can profit by contesting that an accurate report is too high because no bank will verify the contestation.

Third, consider if the reporting bank reports too high. Thus, $s_i > \underline{b_i}$. Whether it was named or not, any informed bank may accuse the bank of reporting too high, $\tilde{s}_i = \underline{b_i} < s_i$. As was shown earlier, the offering bank can credibly support the accuracy of this contestation. This is independent of the total number of contestations leveled, such that regardless of the total number of informed banks, they all have an incentive to contest and potentially receive the positive whistleblower payment. Therefore, if the reporting

---

[47]If the offering bank were to contest the report, it could be profitable for a non-offering bank to also contest the report. As we will demonstrate that it is either unprofitable or impossible for the offering bank to contest an accurate report, this is not a profitable deviation.

bank reports too high, the informed bank contests the rate. The offering bank does not contest (as there is not necessarily anyone to verify its contestation), and instead waits until the verification stage to verify the informed bank's contestation.

Note that if the offering bank knows that there exist other banks willing to lend for less than the reported rate, then there are other sequentially rational strategies of the continuation game. For example, it is possible that the offering bank could contest at the lowest rate at which the *next* bank would be willing to lend, and have that bank verify in the verification stage. There are other possibilities. Note importantly, however, that any sequentially rational pure strategies of this continuation game must generate verified contestation, so the outcome for the purposes of this proof is equivalent.

Fourth, consider if the reporting bank reports too low. Thus, $s_i < \underline{b_i}$. In this situation, the offering bank and any informed banks, whichever were not named by the reporting bank, could accuse the reporting bank of having misreported. This would consist of a report, $\tilde{s}_i = \underline{b_i} > s_i$. Lenders would be sought to lend $\overline{q}$ at a rate s_i to the reporting bank. For the reasons previously mentioned, none will be found as no bank is willing to lend to the reporting bank at its reported rate. A (randomly chosen) contesting bank will then receive the whistleblower payoff, providing an incentive for all informed banks and the informed bank to contest (if possible.) Therefore, similarly to before, if the reporting bank reports too low, it is caught. This is again independent of the total number of contestations.

Finally, it can be noted that in the latter two of these situations (where the reporting bank reports incorrectly), it is not necessary that the contesting banks contest with $\tilde{s}_i = \underline{b_i}$. For example, if the initial report of the reporting bank was too high, $s_i > \overline{b_i}$, then any report $\tilde{s}_i \in [\underline{b_i}, s_i)$ can lead to a positive payoff. However, the contestation $\tilde{s}_i = \underline{b_i}$ weakly dominates all other possible contestations. If the contestation is alone, then the precise report $\tilde{s}_i$ is irrelevant. If there are other contestations, though, only the most accurate of the contestations receives the whistleblower payment. Therefore, it is weakly best for a contesting bank to contest with $\tilde{s}_i = \underline{b_i}$.

In the contestation stage, the bystander banks' actions are conditioned on their beliefs about the accuracy of other banks' reports. The bystander banks update their beliefs such that the reporting banks' report is given probability one of being its true borrowing cost. This arises because the strategy of the banks is to truthfully report borrowing costs; therefore, absent some special information, bystander banks place all posterior probability on the report having been accurate.

Note that while bystander banks can be 'fooled' by a false report, the informed banks are not. Informed banks' belief distributions are initially

degenerate, and do not update with the report of the reporting bank. Therefore even out-of-equilibrium contestations which are not accurate will be punished, and those that are accurate will be rewarded (with positive probability).

This concludes the proof of the optimality of the specified contestation strategy. In equilibrium, banks optimally contest every report known to be fallacious, and never contest a correct report *unless* they are an un-named offering bank. Next, we consider the reporting strategy.

Sequential Rationality in the reporting stage

By this point, we have shown the optimality of the specified strategies at the verification and contestation stages of the mechanism. These strategies dictate that any correct report is un-contested, while any incorrect report is contested, the contestation verified and the reporting bank punished. Therefore, if a bank reports its rate incorrectly, it benefits through its exposure to LIBOR, but it is punished with the punishment fine, ϕ_r. This tradeoff is represented by Equation P3:

$$\pi_i^L \left(\mathbb{L}(s_i, b_{-i}^*) \right) - \phi_r - w < \pi_i^L \left(\mathbb{L}(b^*) \right) \quad \forall s_i \tag{P3}$$

So long as Equation P3 is satisfied, the punishment from mis-reporting outweighs the benefit from manipulating LIBOR. Therefore, the given the optimal strategies at the contestation and verification stages, the only optimal strategy is for the reporting banks to report accurately and to name the offering bank correctly.

This concludes the proof of existence. The specified strategies are optimal both on and off the equilibrium path and form a Perfect Bayesian Equilibrium.

7.1.3 Proof of Uniqueness

Proof that correct reports are not contested in equilibrium In this section, we prove that the only sequentially rational strategies of the continuation game following on from a correct report do not include any contestation. The prior section showed that the leaving correct reports un-contested is an equilibrium; in this section we show that this is the only possible sequentially rational pure strategy of the continuation game.

First, as was shown in the previous section, the only sequentially rational strategy in the verification stage is for a given bank j to offer to lend to bank i at any rate ρ_{ji} or higher. Effectively, 'correct' verification is the only equilibrium strategy of the final phase of the mechanism.

We next show that any contestation will not be supported in the verification stage. Accordingly, any contesting bank will be shown to have contested falsely and be punished. As a bank that does not contest generates a payoff of zero, this proves that no sequentially rational strategies of the continuation game can sustain contestation. In effect, when a bank reports its borrowing rate accurately, and names the offering bank appropriately, the report is uncontested.

To start, consider if the contestations come only from the non-offering banks. We consider the alternative situation next. First, consider that the reporting bank has reported correctly, and that some $m \geq 1$ bank(s) contest that the reported rate is too high. Thus, $\tilde{s}_i < s_i$ for all m contestations. All other banks are asked if they would lend to the reporting bank at each rate $\tilde{s}_i$. A bank x that states willingness to lend at a given rate $\tilde{s}_i$ will, with some probability, be required to lend $\overline{q}$ at that rate. (The probability is at most $1/m$. If there are μ banks that all express willingness to lend at the same rate, then the probability is $1/m\mu$.) With complementary probability, the statement has no impact. Therefore, a bank x conditions its response on the states in which the response is relevant, and states willingness to lend if and only if the following holds:

$$\pi_x^{\Psi}\left(\Psi_x^i\left(\overline{q}, \tilde{s}_i\right)\right) > \pi_x^{\Psi}\left(\Psi_x^i\left(0, \cdot\right)\right) \tag{11}$$

However, because the reporting bank has reported correctly, the lowest rate at which any bank will lend to it is $\underline{b_i} = s_i > \tilde{s}_i$. There is no bank x that is willing to lend at any rate below $s_i = \underline{b_i}$. Thus the false contestation(s) will not be supported - it will be shown that the reporting bank reported correctly.

Second, consider that the reporting bank has reported correctly, and that some $m \geq 1$ banks contest that the report is too low. Thus $\tilde{s}_i > s_i$ for all m contestations. All other banks are asked at which of a set of rates they would be willing to lend an amount $\overline{q}$ to the reporting bank. The set of rates includes the initial report, s_i, and the lowest $(m-1)$ contestation reports, $\tilde{s}_i$. The offering bank j is always willing to lend at *all* of these rates because, by definition,

$$\pi_j^{\Psi}\left(\Psi_j^i\left(\overline{q}, r\right)\right) > \pi_j^{\Psi}\left(\Psi_j^i\left(0, \cdot\right)\right) \ \forall r \geq s_i = \underline{b_i} \tag{12}$$

Therefore, as the offering bank was not one of the contesting banks, the offering bank will express willingness to lend at *all* m rates. (Other banks may express willingness too, but an examination of the offering bank is sufficient.) This will show both that the original report was correct, and that all contestations were spurious.

Third, consider that the reporting bank has reported correctly, and that some $m > 1$ banks contest, simultaneously, that the reported rate is *both* too high and too low. Thus, for some subset of the contestations, $\tilde{s}_i < s_i$; for the others, $\tilde{s}_i > s_i$. As per the mechanism, of the m contested rates and the one initial report, the lowest m will be selected, and banks will be asked at which they would be willing to lend money. From above, it is clear that so long as the offering bank was not one of the contesting banks, it will be willing to lend at both s_i, and all $\tilde{s}_i > s_i$. (Other banks may be willing as well, though this does not change the result.)

There is no bank x that is willing to lend at any rate below $s_i = \underline{b_i}$, by Assumption 1. Therefore, even if a truthful report is contested in two directions simultaneously, the remaining banks have incentives to lend only at the correct rates, revealing the truth of the initial report.

Therefore, it cannot be an equilibrium for any subset of banks other than the offering bank to contest the correct report.

Now, consider if the offering bank contests. The offering bank is unable to contest that the initial report is too low (because it has been named by the reporting bank, and the rules of the mechanism prohibit named banks from contesting that a rate is too low). Therefore, the offering bank can only contest that i's initial report was too high, $\tilde{s}_i < s_i$. This can only be confirmed by another bank lending at the contestation rate. However, there exists no bank that will lend for less than $\underline{b_i}$ (as was shown above). Therefore, if the correctly named offering bank contests the report, its contestation is determined to be false and the offering bank is punished. The offering bank cannot contest in equilibrium.

Therefore, as there cannot be an equilibrium in which the offering bank contests (either alone or with others), and there cannot be an equilibrium in which only other banks contest, the proof is complete. In the continuation game following on from a correct report by the reporting bank, sequentially rational contestation cannot occur.

Proof that incorrect reports are discovered and punished in equilibrium   In this section, we prove that the only sequentially rational pure strategies of the continuation game following on from an incorrect report must include (verified) contestation. The proof earlier showed that incorrect reports are contested in equilibrium; this section shows that any equilibrium must include verified contestation.

This proof proceeds in two parts. First, we show that if the reporting bank names the offering bank incorrectly, then the report is necessarily contested (and the contestation verified). This holds whether or not the reported rate is

correct. Second, we show that if the reporting bank names the offering bank correctly but misreports its rate, then the report is necessarily contested (and the contestation verified). This completes the proof that any misreporting is necessarily contested, and the misreporting bank punished.

For all of these situations, we use the unique sequentially rational strategies of the verification stage. We show that there cannot exist a pure-strategy equilibrium without contestation. As un-verified contestation also cannot be an equilibrium (banks would prefer not to contest than to contest and have their contestation un-verified because of the punishment fine), this proves that the only equilibria contain verified contestation. This results in a punishment fine levied on the reporting bank.

First, consider that the reporting bank names a bank other than the offering bank. We will show that regardless of whether the reporting bank also reports the rate correctly or incorrectly, there cannot exist an equilibrium without contestation. There is always a bank with an incentive to deviate. We consider accurate reporting, under-reporting, and over-reporting in turn.

If the reporting bank names a bank other than the offering bank, and reports its rate correctly, the offering bank is able to contest strategically. If the offering bank, j, contests that the reporting bank i, has under-reported, $\tilde{s}_j > s_i$, all of the *other* banks will be polled to determine if they are willing to lend at the initially reported rate. Since with probability one the offering bank is the only one willing to lend at its precise rate, no other bank will offer to lend to disprove the offering bank's (strategic) contestation. The offering bank is paid the whistleblower benefit for contesting, as its contestation is upheld at the verification stage. Therefore, if a reporting bank reports its rate correctly but names incorrectly, it cannot be an equilibrium for there to be no contestation as the offering bank has an incentive to deviate.

Next, consider if the reporting bank, i, names a bank other than the offering bank, and simultaneously under-reports the rate at which it can borrow. In this situation, a no-contestation equilibrium cannot exist because every un-named informed bank, j, (including the offering bank) has an incentive to contest the initial report, $\tilde{s}_j = \underline{b_i} > s_i$.[48] As the reporting bank has under-reported, when its report is contested, there exists no bank that will be willing to lend at the reported rate. This failure to procure a loan at the reported rate confirms the verification, and results in a punishment fee being levied on the reporting bank. Therefore, if the reporting bank names incorrectly and under-reports its rate, a no-contestation equilibrium is not possible

---

[48]The named bank, if it is informed, also has an incentive to contest, but by the rules of the mechanism, the named bank cannot contest that the reporting bank has under-reported.

because the informed banks will always deviate to contest the reported rate.

Third, consider if the reporting bank, i, names a bank other than the offering bank and simultaneously over-reports the rate at which it can borrow. In this situation, the naming is irrelevant, because the named bank is still able to contest that the reported rate is too high. Therefore, a no-contestation equilibrium is again impossible because the informed bank(s) have an incentive (and the ability) to contest at the correct rate: $\tilde{s}_j = \underline{b_i} < s_i$. The offering bank does not contest (as it would be the bank required to verify a true contestation), and instead waits to offer to lend in the verification stage. In the verification stage, the offering bank confirms the informed bank's contestation by offering to lend at the correct rate.[49]

This concludes the proof that if the reporting bank names incorrectly, then its report is necessarily contested, the contestation verified, and the bank punished. We have proved that there cannot exist a pure strategy equilibrium without contestation, as there always exists a bank with a profitable deviation. Given that unverified contestation cannot exist as part of an equilibrium, this concludes that incorrect naming results in punishment. Next, we show that if the reporting bank names the offering bank correctly, but reports incorrectly, this mis-reporting is necessarily discovered and punished.

Consider if the reporting bank, i, correctly names the offering bank but under-reports the rate at which it can borrow. In this situation, a no-contestation equilibrium is not possible because the informed bank(s) will contest that the initial report was too low, $\tilde{s}_j = \underline{b_i} > s_i$. As no bank will offer to lend at the reported rate (because it is too low), the contestations are confirmed. This results in the whistleblower benefit being paid to one of the contesting banks, and the reporting bank being punished. Therefore, a no-contestation equilibrium is not possible because the informed bank(s) have an incentive to contest the initial report, and their contestation is upheld in the verification stage.

Finally, consider if the reporting bank, i, correctly names the offering bank but over-reports the rate at which it can borrow. In this situation, the informed bank(s) have an incentive to contest the reported rate, $\tilde{s}_j = \underline{b_i} < s_i$. In the verification stage, the offering bank will lend at the contested

---

[49]Note that there do exist multiple possible equilibria of this continuation game. For example, it is possible that the offering bank could contest as well, at the lowest rate at which the next bank would be willing to lend, and have that bank verify in the verification stage. However, any sequentially rational pure strategies of this continuation game must involve verified contestation, so the initial report is still found to be inaccurate and punished. Therefore, all sequentially rational pure strategies still generate the same incentive *not* to manipulate reports at the reporting stage.

rate, confirming the contestation, and resulting in the reporting bank being punished. Therefore, there cannot be a no-contestation equilibrium when the reporting bank over-reports, because the informed bank(s) have an incentive to deviate and to contest the report. This contestation is upheld in the verification stage.

This concludes the proof. Regardless of the way in which the reporting bank mis-reports, there exist no equilibria in which the report goes uncontested. Since an equilibrium cannot include a contestation that goes unverified, this implies that the only equilibria of the continuation game following on from an incorrect report result in the reporting bank being punished.

Proof of uniqueness of proposed equilibrium outcome Finally, these various points can be combined. We have shown that the strategies as defined form a Perfect Bayesian Equilibrium. We have shown that the only sequentially rational strategies of the continuation game following on from a correct report do not include any contestations. Finally, we have shown that the only sequentially rational pure strategies of the continuation game following on from a correct report must include verified contestations.

Combined with Constraint [P3], these combine to prove that all pure-strategy equilibria of the mechanism result in truthful reporting. Any bank that manipulates its report is punished, and the punishment more than outweighs the benefit of manipulation. Therefore, every bank optimally chooses to report truthfully, and this is the only pure strategy Perfect Bayesian Equilibrium outcome.

This concludes the proof of Proposition 1.

7.2   Proof that collusion cannot occur in the model

In this section, we illustrate that collusion cannot occur in the model. Collusion is impossible because of sequential rationality: at any point at which collusion may occur, the parties to the collusion have an incentive to break their agreement at a later stage in the mechanism.

Collusion can conceivably take place at four points: at the reporting stage, at the loan transaction presentation stage, at the contestation stage, and at the verification stage. If banks collude at the stage where contestations are verified, they must *also* collude at another stage. Otherwise, the verification stage is not reached, and the collusion is immaterial. Therefore, we begin with collusion at the reporting stage, where the goal of the colluding banks is to manipulate LIBOR. Then, we move on to considering collusion beginning at the contestation stage, with the aim of falsely accusing banks that reported accurately. We find that neither type of collusion is sustainable in a Perfect

Bayesian Equilibrium, even if coupled with collusion in the verification and presentation of transactions stages.

7.2.1 Reporting Stage

We begin by illustrating the impossibility of collusion at the reporting stage. Consider the most extreme situation: all banks know that all banks have an incentive to manipulate LIBOR in the same direction. This is the situation in which collusion is most profitable and presumably easiest to sustain: we will show that collusion is still impossible in this case.

First, consider the outcome if all banks are known to have an incentive to manipulate LIBOR lower. Any number of banks may report $s_i < b_i^*$ in an attempt to influence the LIBOR calculation. Banks that have borrowed have two options: either to reveal the wrong transaction at the reporting stage, or to claim that they did not borrow at all. Once banks report, however, they are exposed to contestations by the other banks. Critically, once the reports are made, the LIBOR calculation is complete: the banks therefore no longer have any reason to collude. After a collusive submission, all banks have an incentive to contest their compatriots' submissions, either through revealing the correct loans, or by contesting the reports of banks that did not borrow. After these contestations are made, in the verification stage, banks have strict incentives to act honestly, and to refuse to lend at the reported rates (thereby supporting the contestations). Therefore, with the contestations supported, each bank that misreported in the first stage of the game will be punished.

Banks may still profit by misrepresenting, even if their misrepresentations are later discovered and punished. This is because the benefit of all banks misreporting LIBOR may be greater than the punishment imposed on banks for misrepresenting their own report (with perfectly correlated LIBOR exposures, in effect, banks' misreporting generates a positive externality for others, while the punishment imposed is calibrated only on the personal benefit of misreporting). Therefore, if all banks could collude to misreport, even though they would later accuse each other of lying and be punished, they would be net beneficiaries.

However, colluding to misreport in the first stage is not individually rational because of the simultaneous submission of reports. While each bank receives a positive payoff if every bank misreports, each bank considers only the impact of its own action. Since the punishment for misreporting is greater than the benefit (for the particular bank), every bank has an incentive to renege on any collusive agreement it makes. The banks want to encourage others to lie, but do not have an incentive to do so themselves. Similarly,

colluding banks may threaten to punish others that do not collude with them. However, once a bank is seen not to have misreported, the only sequentially rational action is *not* to contest the bank's report. Punishment strategies, insofar as they are even possible, are not sequentially rational.

The same argument holds when the banks wish to manipulate LIBOR higher. To manipulate LIBOR higher, banks would claim not to have borrowed, and report $s_i > b_i^*$. (Banks cannot report the wrong transaction, because this necessarily results in a *lower* LIBOR.) Again, any possible collusion unravels, because once the reports are submitted, banks have an incentive either to reveal the extant transactions, or to contest $\tilde{s}_i = b_i^* < s_i$. Importantly, while we have described this collusion as between a number of reporting banks, this approach also proves the impossibility of any form of collusion between reporting banks and any other banks. Therefore, while we have shown that collusion between reporting banks is not sequentially rational, the same approach shows that collusion between reporting banks and the offering bank, between reporting banks and contesting banks, and between reporting banks and banks that have already lent to them, is also not sequentially rational.

Any collusive attempt unravels both because of the early setting of the LIBOR rate (encouraging the colluding banks to whistle-blow on each other) and because of the simultaneous reporting (encouraging banks to renege on the entire agreement). This occurs even if banks are able to make cash transfers between themselves: in a finitely-repeated game, no bank can commit to pay a promised cash payment. Next, we consider possible collusion at the contestation stage.

7.2.2 Contestation Stage

In this section, we consider that all banks have reported accurately. Then, we determine if there exists any possible sequentially rational form of collusion that could allow banks to falsely accuse others of improperly reporting. Collusive contestations may then take one of two broad forms: either accusing a rate of being too high, or accusing a rate of being too low. Importantly, *neither* of these has any impact on the LIBOR rate. Therefore, banks' exposures are immaterial.

Any two banks may collude and accuse a third bank's report of being too high. To do so, the first bank accuses the report of bank i of being too high, and reports a new rate, $\tilde{s}_i < s_i = \underline{b}_i$. The second bank would then lend at this lower rate, generating a net payoff to the colluding banks of w, less the cost of the below-market loan. However, the bank that issues the loan incurs only the cost of making the loan, while the benefit of the (falsely

obtained) whistleblower payment accrues solely to the first bank. Therefore, this form of collusion unravels: neither bank is willing to make the necessary below-market loan.

Similarly, banks may collude and accuse another bank's report of being too low. This would involve the offering bank and another bank colluding: the second bank would accuse the reporting bank of having reported too low, while the offering bank simply refuses to lend in the verification stage. This, too, is easily shown not to be sequentially rational, however. After the second bank accuses the reporting bank of having reported too low, the offering bank has an incentive to breach its collusive agreement and lend at the reported rate. A second option is for the offering bank to accuse the reporting bank of reporting too high, at the same time as the other bank accuses it of reporting too low. This would effectively tie the hands of the offering bank, such that it cannot lend in the verification stage (because it is one of the contesting banks). The offering bank would not agree to this, however, given that its incentives are always to breach any such agreement (and no punishment can be levied for so doing).

Overall, then, the separation of the determination of LIBOR and the assignment of punishments for misreporting unravel any possible form of collusion in the model due to sequential rationality. Even though the whistleblower payment, w, may provide a high incentive to collude, any possible attempt to collude unravels because of the timing of the reporting mechanism.