A Modest Change To The LCR That Could Substantially Improve Financial Stability

3.21.2019

BANK POLICY INSTITUTE

Bill Nelson
202.589.2454
bill.nelson@bpi.com

THE USABILITY PROBLEM

The liquidity coverage ratio (LCR) requires certain banks and bank holding companies to hold high quality liquid assets (HQLA) sufficient to meet projected 30-day liquidity needs in a situation of severe idiosyncratic and systemic stress. The LCR is calculated as HQLA divided by projected 30-day net cash outflows under stress; the requirement is that banks maintain a ratio of at least 100 percent. A problem recognized by economists and regulators alike is that the LCR can only help a bank meet its liquidity needs under stress if its HQLA can actually be used (that is, sold or used as collateral to borrow funds). However, use of the HQLA would, in many circumstances, cause the bank's LCR to fall below 100 percent and therefore constitute a violation of the regulation, an outcome the bank would work strongly to avoid. This is the problem of "HQLA usability."[1]

There is a broad consensus that usability is desirable. If both banks and their counterparties are comfortable with banks tapping their HQLA under stress, then shocks to the liquidity of money markets will be less likely to lead banks to hoard liquidity and withdraw from term lending, resulting in a broader liquidity crisis. Such a dynamic was an important early driver of the 2007-09 crisis.

Moreover, if banks are fearful of violating the LCR requirement under stress, they will hold a buffer of HQLA above the required level at all times. Bankers report that they maintain LCRs of about 110 percent to reduce the risk that they will fall below the requirement of 100 percent. Indeed, at the end of the third quarter, the eight largest banks were required to hold $1.9 trillion in HQLA and chose to hold $2.3 trillion.[2] There are material social costs associated with banks holding excessive HQLA. Not only would economic activity be higher if more bank assets were loans to businesses and households and fewer were cash and government securities, but also HQLA has valuable alternative uses for others as collateral and as an investment for those needing safety and liquidity.

Reducing banks' need to hold a sizeable buffer above the LCR minimum would also allow the Federal Reserve to operate with a smaller balance sheet. Banks subject to the LCR currently hold about $3 trillion in HQLA, a considerable part as deposits at the Fed (aka "reserves"). In a recent speech, Vice Chairman Quarles noted that the size of the Fed's balance sheet was determined by bank demand for reserves, and that bank demand for reserves was determined importantly

---

1 See The BPI blog post, "Unlocking the Liquidity Coverage Ratio," Bill Nelson, June 7, 2017 for additional information on the usability problem and other suggestions to address it.

2 In addition to concern about violating the LCR, banks hold HQLA above the amount required by the LCR to meet liquidity requirements associated with resolution, other supervisory expectations, or their own internal assessments of necessary liquidity. See the BPI blog post "Missing Answers to a recent Fed survey," Greg Baer and Bill Nelson, January 17, 2019, and "Senior Financial Officer Survey," Federal Reserve Board, September 2018. The adjustment to the LCR suggested here would only impact the HQLA demands of banks that are bound by the LCR; if the same adjustment were made to other supervisory expectations about liquidity, the impact would be larger.
"The Future of the Federal Reserve's Balance Sheet," Vice Chairman for Supervision Randal K. Quarles, February 22, 2019. https://www.federalreserve.gov/newsevents/speech/quarles20190222a.htm.

by the LCR.[3] If banks reduce their approximately $300 billion buffer over the LCR, the Fed will be able to operate with a correspondingly smaller portfolio of securities – something all appear to agree would be a good outcome.[4]

A POSSIBLE SOLUTION – THE DYNAMIC LCR

In this post, we suggest a change to how the LCR is calculated that would make HQLA more usable and thereby promote financial stability. We refer to the proposed change as a "dynamic LCR." The change would help achieve the financial stability objective of the LCR, would not ease the requirement in the normal course of business, and yet would reduce the need for banks to hold liquidity buffers on top of the buffer established by the LCR.

The dynamic LCR is based on a different, but equally valid, way to conceptualize liability outflows. The LCR currently assumes that any pool of liabilities – for example, operational deposits generated by clearing, custody, and cash management activities – is uniform and should be expected to flow out at a constant rate in each period when conditions remain constantly stressed.[5] In contrast, the dynamic LCR assumes that within any pool of liabilities, some are flightier than others (owing to different characteristics of the investors) so that when some flow out, the average "stickiness" of the remaining pool of investors increases. Further below, we discuss some evidence supporting the heterogeneous investor assumption behind the dynamic LCR.

In a nutshell, we suggest that, for each liability type, any outflows that have recently occurred be deducted from projected future outflows when determining the HQLA requirement. The dynamic LCR is simple to implement and, because outflows are small or nonexistent in normal times, the change will have a de minimis effect on a banks' LCR requirements in normal times (which is to say, the great majority of the time). But with the proposed adjustment, it will generally no longer be the case that a bank that uses its HQLA to meet outflows will reduce its LCR. In other words, if a bank does exactly what the LCR expects and is premised on its doing, its LCR will not suffer. As a result, the bank will have a reduced incentive to hoard liquidity.

HOW WOULD THE DYNAMIC LCR WORK?

The LCR identifies multiple different types of deposits and other liabilities and assumes that each runs off at a constant percentage rate over any 30 days of stress based only on the current amount of the liability item. For example, operational deposits are assumed under the LCR to run at a 25 percent rate.

So, if a bank has $1000 in operational deposits, the LCR assumes that 25 percent ($250) will run and requires the bank to hold $250 in HQLA. If $100 then run, the LCR assumes that 25 percent of the remaining $900 will run ($225) and requires the bank to hold $225 in HQLA. Assuming the bank had used $100 in HQLA to meet the $100 run, it now has $150 in HQLA and so must raise an additional $75 in HQLA during the stress period.

The alternative assumption behind the dynamic LCR is that within any $1000 of operating deposits there are 25 percent that will run under LCR-level stress. Once those have left, the remaining deposits are assumed to remain sticky. Again, consider a bank that starts with $1000 in operational deposits. Under the dynamic LCR, 25 percent of those deposits ($250) would be judged to

---

3 "The Future of the Federal Reserve's Balance Sheet," Vice Chairman for Supervision Randal K. Quarles, February 22, 2019. https://www.federalreserve.gov/newsevents/speech/quarles20190222a.htm.

4 The Fed's "Policy Normalization Principles and Plans" state "The Committee intends that the Federal Reserve will, in the longer run, hold no more securities than necessary to implement monetary policy efficiently and effectively…"

5 The last point is important, while runs tend to intensify because people see others running, or stop because of government intervention, the LCR is designed to estimate liquidity needs in precisely the unchanging set of circumstances (neither intensifying nor easing), spelled out in the regulation.

be prone to run in a period of LCR-level stress and the bank would be required to hold $250 in HQLA. If $100 then run, then only $150 of the remaining $900 would be projected to run, so the bank would be required to hold $150 in HQLA. Assuming the bank had used $100 in HQLA to meet the $100 run, it now has $150 in HQLA and so continues to satisfy its LCR.[6]

It is necessary to define a specific window over which the recent change in a liability is measured. We propose 30 days, symmetric to the 30-day forward-looking window of the LCR, but it could be shorter or longer. A longer look-back would do more to support usability because the adjustment to the LCR would be more long-lived. A shorter look-back would focus the adjustment on recent, severe outflows. In either case, the adjustment would provide only a temporary reprieve. The precise formula that could be used to calculate net outflows under the dynamic LCR is provided in the appendix.

Importantly, the LCR and dynamic LCR require two banks with the same balance sheets to hold the same level of HQLA in the normal course of business, so the banking system would enter a stress period with approximately the same aggregate required HQLA. But in a stress period, when there actually are outflows, the current LCR would increase liquidity pressure on banks while the dynamic LCR would allow banks to use their HQLA for the purpose it was set aside – meeting demands on their liquidity.

IS THE DYNAMIC LCR REASONABLE?

Is it more reasonable to think that within any type of liability the investors have a relatively uniform propensity to run under a constant level of stress or that there are some inclined to run and others inclined to stay? Detailed data on depositor behavior under stress are scarce, but recent research on bank runs has emphasized that within a pool of investors the propensity to run depends on depositor characteristics and that the flightier investors run quickly, increasing the density of stickier investors within the pool.

Carlson and Rose (2016) studied the run on Continental Illinois in 1984. They note that the initial run on the bank was massive, fast, and short-lived, with 30 percent of funding leaving in 9 days. They find no evidence that the run was driven by concerns about Continental failing, but rather by investors who had a stronger preference for liquidity, were somewhat less sophisticated, and were located farther away.

Martin et al (2018) examined the experience of a medium-sized bank that came under pressure in the crisis using a unique daily account-level dataset. They too find that account stability depends importantly on account characteristics; for example, a customer receiving direct deposit or customers with a longer relationship is less likely to leave the bank.

Iyer and Puri (2008) examine depositor behavior at an Indian bank when news broke that a larger bank in the same city had committed massive fraud and subsequently failed. The smaller bank experienced a sharp run of deposits during the few days around the failure of the larger bank after which the run stopped. Most of the customers that withdrew funds withdrew all or almost all of their balances, so within a few days the characteristics of the pool of depositors had shifted toward one with a propensity not to run. By examining the personal reference each accountholder gave when opening an account, Iyer and Puri determine that a depositor's likelihood of running depends on whether others in his or her network have run. They also find that borrowers that have received a loan from the bank were less likely to run.

The Clearing House Association (TCH) (2011) (one of the antecedents of BPI) gathered data on bank liability outflows as well as draws on committed lines during the crisis as input for the calibration of the LCR. The report provides monthly data on the banks experiencing some of the largest demands on their liquidity. In the case of the bank in the TCH sample that experienced the

6 As described in the appendix, projected outflows would be restricted to be nonnegative.

worst 30-day wholesale deposit runoff (33 percent), the runoff fell to near zero in the subsequent month, and then reversed the month after that.[7] The largest monthly percentage draw on corporate credit lines (10 percent) fell to zero the following month. Taken together, these studies suggest that investors within any liability class are heterogeneous with respect to their propensity to run and that those that do run do so quickly. That said, there is, of course, contrary evidence – Carlson and Rose (2016), for example, find that the run on Continental Illinois occurred in waves and some that didn't run in the earlier wave ran in a later wave. But given that the assumption in the LCR about outflow propensity – that it is uniform across investors within a deposit category – does not appear to be based on any analysis, the assumption in the dynamic LCR about outflow propensity – that it is heterogenous across investors within a class—appears just, if not more, reasonable and leads to a regulatory metric with much better financial stability properties.

IS THE DYNAMIC LCR SAFE?

Public policy on bank liquidity is about ensuring banks maintain enough liquidity in good times and use their liquidity in bad times. If the dynamic LCR is adopted, when a bank is growing or not experiencing a significant outflow of liabilities, its LCR would essentially be unchanged from the current requirement. Apart from reducing the incentive for banks to hold liquidity buffers on top of their liquidity buffers, this change would still leave the banking system with the intended amount of liquidity going into a period of stress.

By design, with this change, if a bank is experiencing stress, its dynamic LCR requirement will go down, but that is a feature not a bug. All "usability" proposals ease the requirement in stress times; that's the point. Possibilities that have been discussed (although not adopted in the United States), for instance, include declaring an "LCR holiday" when there is systemic stress. Mervyn King, then Governor of the Bank of England, made essentially such an announcement that linked the availability of central bank funding to HQLA usability in a speech delivered in June 2012 during the European financial crisis:[8]

> Next January, the Basel rules for the Liquidity Coverage Ratio will be agreed by central bank governors and heads of supervision. Much work still needs to be done to ensure that those rules are properly integrated with the regime of liquidity provision by central banks. In current exceptional conditions, where central banks stand ready to provide extraordinary amounts of liquidity, against a wide range of collateral, the need for banks to hold large liquid asset buffers is much diminished, and I hope regulators around the world will take note. (p. 4)

To be clear, while the dynamic LCR will be just as effective as the current LCR in ensuring that banks maintain robust liquidity resources in normal times, and better than the current LCR in ensuring that banks can use those resources in stress times, it will be less effective than the current LCR as a tripwire for increased supervisory scrutiny or action when a bank is experiencing liquidity pressures. For example, the dynamic LCR would be less effective than the LCR as a way to avoid another Bear Stearns situation – Bear Stearns alerted the Fed on Thursday that if it were not bailed out it would default on Friday.[9] However, the banking agencies have many ways to monitor the liquidity situation of large banks and their affiliates; large banks, for instance, have to report on their liquidity condition on a daily basis. And of course, a bank experiencing a serious liquidity problem would have an incentive to take remediating actions even if it were in compliance with a dynamic LCR.

---

7 The Clearing House study does not identify the individual banks. The severe deposit outflow may have stopped because of government intervention. See https://www.theclearinghouse.org/-/media/files/association-documents/20111102-tch-study-assessing-the-liquidity-coverage-ratio.pdf p.17

8 Speech given by Mervyn King, Governor of the Bank of England At the Lord Mayor's Banquet for Bankers and Merchants of the City of London at the Mansion House 14 June 2012. https://www.bankofengland.co.uk/-/media/boe/files/speech/2012/mansion-house.

9 See "Report Pursuant to Section 129 of the Emergency Economic Stabilization Act of 2008: Loan to Facilitate the Acquisition of the Bear Stearns Companies,Inc. by JPMorgan Chase & Co.

While the LCR holiday approach has not been adopted in the United States, the implementing regulation for the LCR indicates that

> …supervisory actions should not discourage or deter a banking organization from using its HQLA when necessary to meet unforeseen liquidity needs arising from financial stress that exceeds normal business fluctuations. p. 61517[10]

Unfortunately, the regulation goes on to say that a bank whose LCR falls below one by any amount for any length of time must notify its supervisor immediately and, if the shortfall persists for a few days, provide a remediation plan, a key reason why HQLA is not usable (p. 61537). For this and many other reasons – most notably, reputational and legal risk – it is difficult to find anyone who believes that such an approach would work.

WOULD A DYNAMIC LCR WORK?

The dynamic LCR would, instead, officially and automatically consider a bank in compliance with the requirement if it used its HQLA to meet an outflow as intended by the regulation. As a result, there would be no need for the bank to report to its supervisor as if it had misbehaved. Consequently, a shift to the dynamic LCR should help ensure that supervisory actions do not deter a bank from using its HQLA when needed, as intended by the banking agencies according to the regulation.

The dynamic LCR is modelled on an aspect of the current LCR that is already designed to encourage usability. Changes in short-term loans to, or borrowings from, other financial institutions, regardless of collateral, generally do not change a banks' LCR.[11] For example, an unsecured loan to a financial institution (a fed funds loan) that matures within the 30-day LCR window is assumed not to roll over. As a result, if a bank makes a fed funds loan to another bank, its cash holdings (and thus HQLA) will fall by the amount of the loan, but it will also project an equal cash inflow when the loan matures, resulting in an equal change to HQLA and projected net cash outflows. Conversely, and importantly, if a bank stops making a fed funds loan, it will not boost its LCR. Although a bit more complicated, the same neutrality applies to repo transactions regardless of collateral.[12]

As a result, when a bank uses its HQLA to pay off short-term financial institutions creditors, its LCR does not decline, and when it stops making such loans, its LCR does not increase. The LCR, therefore, does not build in any incentive for banks to pull back from each other, at least for loans with terms less than 30 days; nor does it establish an incentive for banks to sell assets at firesale prices to meet outflows of such creditors. The dynamic LCR extends this feature of the LCR to all liabilities.

CONCLUSION

HQLA usability is a difficult problem but solving it offers the prospect of making the LCR a much more effective tool for promoting financial stability. The usability problem is intrinsically linked to the characteristic of the LCR that a bank that uses the liquidity resources the regulations requires be set aside to meet demands on its liquidity – the objective of the regulation – can put itself out of compliance with the regulation. The dynamic LCR proposed here should correct that flaw in the regulation and make it a damper, rather than an amplifier, of systemic liquidity pressures.

---

10 https://www.federalregister.gov/documents/2014/10/10/2014-22520/liquidity-coverage-ratio-liquidity-risk-measurement-standards

11 See "CGFS Papers No 54, Regulatory change and monetary policy," Committee on the Global Financial System and the Markets Committee, May 2015. https://www.bis.org/publ/cgfs54.pdf

12 There are other internal constraints in the design of LCR, such as a limit on the amount of projected cash inflows that can offset projected cash outflows, and limits on the contribution of certain asset types to HQLA, that can break the symmetry.

APPENDIX

Let L_t be a bank liability at time t with an LCR cash outflow rate assumption of a where $0 \leq a \leq 1$. The projected cash outflow is aL_t. Let $\Delta L_t = L_t - L_{t-1}$, the change in the liability over the previous 30 days. If ΔL_t is negative (the liability declined), then instead of a L_t the projected outflow would be $aL_{t-1} + \Delta L_t$, the projection calculated from the previous period minus the outflow that has already occurred. Of course, if funds flow out in an amount above to the LCR projection, it would be unreasonable to project that some of the funds would flow back in, so $aL_{t-1} + \Delta L_t$ can't be negative. And if funds came in during the previous period (the liability grew), the outflow is just the unadjusted assumption aL_t.

Putting that all together, you get: The outflow assumption is min (aL_t, max (0, $aL_{t-1} + \Delta L_t$)).

REFERENCES

Carlson, Mark A. and Rose, Jonathan, Can a Bank Run Be Stopped? Government Guarantees and the Run on Continental Illinois (2016-02-03). FEDS Working Paper No. 2016-3.

Iyer, Rajkamal and Puri, Manju, Understanding Bank Runs: the Importance of Depositor-Bank Relationships and Networks (August 2008). NBER Working Paper No. w14280.

Martin, Christopher and Puri, Manju and Ufier, Alexander, Deposit Inflows and Outflows in Failing Banks: The Role of Deposit Insurance (May 1, 2018). FDIC Center for Financial Research Paper No. 2018-02.