SPEECH

On the Optimal Supply of Reserves

May 22, 2025

John C. Williams, President and Chief Executive Officer

Remarks at the New York Fed – Columbia SIPA Monetary Policy Implementation Workshop, Federal Reserve Bank of New York, New York City

As prepared for delivery

Let me start by personally welcoming you to the New York Fed. We have enjoyed a long and productive relationship with Columbia's School of International and Public Affairs, or SIPA, and it's great to be here to discuss timely and important issues.

The topic of my talk today is the optimal supply of central bank reserves. Prior to the global financial crisis, this issue was more or less settled. Then, in response to that crisis and the ensuing economic downturn—and again following the COVID-19 pandemic—many central banks expanded their balance sheets through various quantitative easing programs funded for the most part by large-scale increases in central bank reserves. These increases resulted in fundamental changes in ways central banks have approached the provision of reserves while maintaining control of short-term interest rates set by the monetary policymaking body. As a result of these experiences in managing large balance sheets, many central banks have reviewed, and in some cases modified, their strategies for supplying reserves and controlling interest rates. Although their approaches have differed in specifics, they share common elements that reflect the fundamental factors that shape the supply and demand for reserves.[1]

Central banks have multiple goals in supplying reserves to the banking system that frequently involve trade-offs.[2] First and foremost, they target a level of the policy interest rate and aim to minimize the variability of the policy rate around that target. In addition, they have goals related to supporting the functioning of financial markets and financial stability. For example, central banks may see advantages or disadvantages to interbank lending in money markets, as well as costs and benefits related to central bank lending into markets.

In this talk, I will consider this question using a relatively simple analytical framework for the supply and demand of reserves that can be applied to various jurisdictions with differences in institutional arrangements and policy objectives. I see this exercise as being in the spirit of William Poole's seminal analysis of the optimal instrument for monetary policy.[3] My goal is to provide a useful background for the rich discussion ahead of us at this conference and elsewhere.

I should emphasize that the substance of my remarks is based on my joint research with Gara Afonso, Gabriele La Spada, and Thomas Mertens, and I have also benefited greatly from what I have learned from my amazing colleagues at the New York Fed and on the BIS Markets Committee. They deserve the credit, and I get to give the standard Fed disclaimer that the views I express today are mine alone and do not necessarily reflect those of the Federal Open Market Committee or others in the Federal Reserve System.

The Demand for Reserves

The key building block of this analysis is that the demand for central bank reserves is inherently nonlinear and subject to uncertainty. For those who are interested, all the equations, Greek letters, and mathematical proofs are in the paper. For this presentation, I will stick to words and pictures to describe the model and results.[4]

Specifically, a central bank faces a demand curve for reserves that describes the relationship between the quantity of reserves and the spread between the policy rate and the interest rate on reserves. This is illustrated by Figure 1, which shows the relationship between the level of reserves in the United States, normalized by banking assets, and the spread between the federal funds rate and the interest paid on reserves from 2015 to early 2020, before the onset of the pandemic. The spread represents the price of reserves, and the shape of the demand curve is the result of a trade-off: Banks receive benefits from holding reserves, but they also incur costs. Since the net benefits decline with the quantity of reserves held, the curve is nonlinear with a lower asymptote.

As shown in Figure 2, this relationship can be represented in a stylized way by a piecewise linear demand curve characterized by three regions. For a low level of reserves, the slope of the demand curve is steep. A small change in the quantity of reserves results in a meaningful change in the spread. This is referred to as the "scarce reserves" region. At a sufficiently large level of reserves, the slope of the demand curve flattens but is still downward-sloping. This region is called "ample reserves." At an even larger level of reserves, the demand curve is flat. This region is called "abundant reserves." The point at which reserves become abundant is labeled X_A in the figure. For reserves above that level, the spread is constant at $\bar{s}$.

So far, I have described the demand for reserves as if it were fixed and known. However, there is considerable evidence that this demand curve in the U.S. has shifted at times over the past 15 years.[5] I will now introduce uncertainty by including random variation in the demand for reserves, as represented by leftward and rightward shifts in the value of the cutoff for abundant reserves, XA.[6] Such variation is illustrated in Figure 3, where I have simplified the model to two regions—ample and abundant reserves. And the demand shock is indicated by the random variable ε. For a given quantity of reserves X, large shocks have asymmetric effects on the resulting level of the spread, as illustrated in the figure. This asymmetric effect is due to the nonlinearity of the demand curve. A sizable positive shock to demand has a larger effect on the spread than a negative shock of equal magnitude, where in the latter case, the presence of the flat demand curve limits the decline in the spread.

Because positive shocks raise interest spreads more than negative shocks lower them, the net effect of shocks is to raise the spread on average. That is, the combination of random variation in demand and a nonlinear demand curve implies that for a given quantity of reserves, the resulting stochastic mean level of the spread is above the value implied by the demand curve absent uncertainty. This relationship is illustrated by the solid red line in Figure 4, which again assumes there are only two regions. The difference between the mean value of the spread under uncertainty and the corresponding deterministic value, shown by the black line, is greatest for levels of reserves in the vicinity of the cutoff between ample and abundant reserves under the deterministic curve.

The Optimal Supply of Reserves

I'll now turn to the central bank's problem of selecting the optimal provision of reserves given the nonlinear demand curve and uncertainty. A central bank is assumed to target a level of the spread, denoted by $\hat{S}$, that lies in the region of ample reserves. The choice of the target spread reflects considerations related to the efficient functioning of money markets and the provision of interbank lending that facilitates the redistribution of reserves across banks.[7] In the following, I take this as a fixed, exogenous target.

The central bank has potentially two tools that it can use to achieve this objective. First, it sets the aggregate level of reserves available to the banking system, as defined by the variable X. Second, it may make available a lending facility to the banking system that offers loans to banks at an interest rate, which I will represent by L. For ease of exposition, I assume that this lending facility has unlimited capacity, but later, I will discuss the implications of putting constraints or costs on its use.[8]

In a deterministic world without uncertainty, this problem is easy to solve. The central bank sets the optimal level of reserves equal to the value consistent with the intersection of the demand curve with the target spread. This is labeled by X* on Figure 4. In this case, perfect interest rate control is achieved at the targeted level without a lending facility.

In the presence of uncertainty regarding the demand for reserves, the problem is more interesting. If the central bank must set the level of reserves and the rate on the lending facility before the demand shock is realized, it can no longer perfectly control the spread. To make things precise, I assume the central bank has a quadratic objective function—that is, it minimizes the mean squared difference between the realized spread and its target rate. With this specification of the objective, the central bank's problem can be described by the sum of two parts: the squared difference between the mean spread and the target and the variance of the spread.

With all of this preparation, I can now derive the optimal supply of reserves. I will start with the special case where there is no lending facility. This simplifies the exposition while conveying clearly a key result.

Case 1: No Lending Facility

Let me cut to the chase: Absent a lending facility, the optimal supply of reserves under uncertainty is unequivocally higher than it is without uncertainty. This finding is the result of two effects that create a precautionary motive to add to the supply of reserves to guard against a positive demand shock. First, as I discussed earlier, for a given level of reserves, the presence of uncertainty increases the mean spread. Second, increasing the supply of reserves reduces the variance of the spread because there is a greater probability of being on the flat part of the demand curve.

Indeed, if the uncertainty is great enough, the optimal level of reserves can exceed the cutoff for abundant reserves under no uncertainty. That is, uncertainty can imply that it is optimal to supply abundant reserves even if the central bank's target spread is in the ample reserves region, as I have assumed.

Although uncertainty increases the optimal supply of reserves, the resulting mean level of the spread may rise or fall with uncertainty. This ambiguous result owes to uncertainty increasing the mean spread for a given level of reserves, while the higher optimal level of reserves reduces the spread. Therefore, one cannot judge whether the policy is optimal simply because the mean spread equals its target.

Case 2: Lending Facility

I now consider the case where there is a lending facility. There are a number of factors that influence the choice of the facility lending rate and the conditions of use for such a facility.[9] For the present purposes, I do not explicitly analyze these factors but instead assume the spread between the facility lending rate and the interest rate on reserves is set exogenously at a level L above the target spread and is invariant to the quantity of reserves. If the intersection of demand and reserves supply occurs at a spread above L, the incremental demand is assumed to be met by the lending facility. Such a situation is illustrated as the solid green horizonal line in Figure 5.

In addition, I assume that there are frictions or costs associated with the scale of the use of the lending facility. I model this as an additional quadratic penalty in the central bank's objective that's associated with the magnitude of lending facility usage.[10] For example, such a penalty could reflect the view that heavy and regular use of the lending facility fosters an overreliance on the central bank for funding needs that should be tempered.

With a lending facility, the choice of the optimal supply of reserves takes into account the fact that the spread is capped at the lending rate. This mitigates the precautionary motive for supplying reserves to guard against a positive demand shock. This effect is greatest when the lending rate is closer to the target spread. On the other hand, if the lending rate is very high relative to the target spread, then the facility is rarely used, and the optimal supply of reserves is near what it would be absent a facility.

The choice of the optimal level of reserves also reflects the trade-off between reducing excess reserves and incurring the costs of additional usage of the facility: The greater the cost incurred for usage of the facility, the higher the optimal level of reserves. In particular, if the costs associated with usage of the facility are relatively small, then the optimal supply of reserves is lower than it would be absent the lending facility. But if the costs are sufficiently large, the optimal supply may exceed that in absence of a lending facility to guard against extensive use of the facility.

Conclusion

This analytical framework for the supply and demand of reserves under uncertainty provides valuable insights into the factors affecting the optimal supply of reserves and the design of liquidity facilities. First, absent a lending facility, the optimal supply of reserves under uncertainty is greater than it would be absent uncertainty. Second, the ideal choice for the supply of reserves depends on institutional factors, such as the pricing of and frictions related to liquidity facilities as well as preferences over the various goals related to monetary policy implementation. In other words, there is no single best way to supply reserves; rather, the best mix of tools depends on circumstances and policy preferences unique to each jurisdiction.

Figures PDF

---

[1] See, for example, the policy implementation frameworks descriptions for the Bank of Canada, the Bank of England, the European Central Bank, Federal Reserve System, and the Reserve Bank of Australia.

[2] See Afonso, Gara, Gabriele La Spada, Thomas M. Mertens, and John C. Williams. 2023. "The Optimal Supply of Central Bank Reserves under Uncertainty." Federal Reserve Bank of New York *Staff Reports*, no. 1077, revised December, for discussion and references to the relevant literature.

[3] Poole, William, 1968. "Commercial Bank Reserve Management in a Stochastic Model: Implications for Monetary Policy," *The Journal of Finance*, *23* (5), 769–91. https://doi.org/10.2307/2325906.

[4] See Gara Afonso, Domenico Giannone, Gabriele La Spada, and John C. Williams, 2022, "Scarce, Abundant, or Ample? A Time-Varying Model of the Reserve Demand Curve," Federal Reserve Bank of New York *Staff Reports,* no. 1019, revised April 2024, for a detailed description of a model of the demand for reserves and empirical evidence on its shape.

[5] See Afonso et al. (2022, revised 2024) for evidence on such shifts.

[6] The following follows Afonso et al. (2023). That paper also considers other sources of uncertainty about the demand for reserves, including uncertainty about the slope of the demand curve and vertical shifts in that affect the entire demand curve.

[7] See Afonso et al. (2023) for a discussion and references to the relevant literature.

[8] I assume that there is no stigma or similar reluctance to borrow from the facility so that the lending rate provides a firm cap on the spread.

[9] For example, in some jurisdictions, the lending facility is designed to be a backstop that is only used intermittingly and the spread between the lending rate and interest on reserves is set at a relatively high level. In other jurisdictions, the lending facility aims to provide a tighter limit on upward movements in interest rates and the lending spread is set close to or at the target spread.

[10] This assumption of a quadratic cost facilitates the analysis. An alternative approach would be to limit the total amount of lending, and the implications would be similar.