

July 13, 2020

Design and Implementation of Yield Caps or Targets to Reinforce Forward Guidance¹

This memo lays out design and implementation considerations for the use of yield caps or targets (YCT) to reinforce forward guidance (FG) about the policy rate. YCT represents a commitment to purchase Treasury securities up to a specified tenor, in order to maintain the yield level that the FOMC views as consistent with its FG policy. YCT could clarify the FOMC’s view on the expected path of rates and support its commitment to FG. However, in practice, the yield curve implications of an FG policy can be ambiguous, especially if uncertainty in the economic outlook translates into uncertainty about the time horizon over which the FOMC’s FG objectives will be met. Therefore, setting the tenor and yield for a YCT program is not necessarily straightforward. Maintaining alignment between YCT and FG over time may also be challenging when the expected path of the federal funds rate shifts with changes in the outlook.

In some circumstances, the commitments to FG and YCT may be sufficient to influence yields without large open market operations. However, in other situations, large purchases could be required. Such purchases could support the FOMC’s policy objectives if they forestalled a premature increase in yields, but would be less beneficial if they resulted from a misalignment between YCT and FG or from anticipation of exit from YCT.

The first section of the memo describes possible strategies for setting and communicating the YCT tenor, including how the tenor could evolve with new information along the path to the eventual exits from FG and YCT. The second section describes choices associated with the setting of the yield level for a specified tenor. The final section discusses potential implications of YCT for the Federal Reserve’s balance sheet. The companion memo on “Financial and Macroeconomic Impacts and Effectiveness of Yield Caps or Targets” discusses how YCT would affect financial conditions and the macroeconomy.

1. The YCT tenor

The choice of tenor is an essential element of aligning YCT with FG. This section examines how the YCT tenor could be determined under different forms of FG and how changing economic information could affect the evolution of the tenor including, but not

¹ Alyssa Anderson and Zeynep Senyuz, Federal Reserve Board; Sam Schulhofer-Wohl, Federal Reserve Bank of Chicago; Patricia Zobel, Federal Reserve Bank of New York. We thank David Bowman, James Clouse, Rochelle Edge, Thomas Laubach, Lorie Logan, David Lopez-Salido, and Trevor Reeve for helpful comments.

limited to, the exits from YCT and FG. A key distinction is between date-based and outcome-based FG, as these two approaches react differently to incoming information.

Specifying the tenor is not straightforward in the presence of uncertainty about the time horizon over which the FOMC’s objectives will be met. At the start of a YCT program, the FOMC would need to specify a particular time horizon for keeping yields low, even though economic conditions might ultimately call for a different path of rates. Over time, depending on how the FG reacts to incoming economic information, the YCT tenor might also need to adjust to changes in the economic outlook.

Importantly, changes in the economic outlook have asymmetric implications. If new data suggest that it will take longer than previously anticipated to achieve the FOMC’s objectives, markets will expect a shallower path of policy rates and—at least to the extent consistent with the effective lower bound on nominal rates—will bring down yields whether or not YCT is in place. But if incoming data suggest the FOMC’s objectives will be achieved sooner than expected, markets will expect higher rates, and the FOMC could need to make large asset purchases to prevent an increase in yields, at least until the YCT horizon or cap level can be adjusted.

To illustrate the key concepts, we assume here that the yield cap or target would be set at a relatively low level, corresponding to a time period when the FG indicates that the policy interest rate will remain close to zero.² Section 2 discusses the choice of this yield level.

A. Date-based FG and YCT

YCT aligns naturally with unconditional date-based FG. If the FOMC is committed to maintaining a policy rate near zero at least until a certain date regardless of economic outcomes, it could reinforce the implications of that commitment for the yield curve by capping yields on securities maturing before the specified date. Because unconditional date-based FG fixes the path of rates to the specified date independent of incoming information, there would be relatively little risk that the YCT parameter settings could become misaligned with the FG based on evolving economic information. (In practice, date-based FG would not be entirely unconditional, so there could be circumstances in the FOMC would adjust the FG and, by extension, the YCT.)

In a YCT program aligned with unconditional date-based FG, the maturity *date* of securities subject to the YCT would remain fixed over time, and correspondingly the *time to maturity* of securities subject to YCT would smoothly roll down or “telescope” with time. As liftoff of the policy rate neared, the YCT program would cap or target yields

² The FOMC could also set a tenor corresponding to a period during which the Committee expects to increase rates, and set the yield level to reflect that path of rates. Such a strategy could reinforce guidance on the pace of policy rate increases after liftoff. The underlying challenges for setting YCT parameters and maintaining alignment with FG would be the same as for the strategies illustrated here.

only on securities with very little time remaining to maturity, which would allow for a smooth exit from the YCT program.

A YCT program linked to date-based FG could be simpler to explain and implement than one linked to outcome-based FG, as the YCT horizon would be constant unless the FOMC modified its FG. However, this simplicity could lead to an undesired stance of policy if the economy did not evolve as expected, and potentially to challenges for YCT. If the economy recovered much faster than expected, it might become increasingly difficult to communicate an unwavering commitment not to raise the federal funds rate target or the yield cap before a certain date. Even though the YCT would still be aligned with the existing FG, large purchases could be required as market participants anticipate a change to the FG and YCT. On the other hand, if the economy recovered more slowly than expected, market participants might expect the FOMC to extend the FG date and YCT tenor, but large asset purchases would not be likely to occur ahead of that extension.

B. Outcome-based FG and YCT

With outcome-based FG, the path of the policy rate adjusts to incoming economic information: The sooner the FOMC’s objectives are met, the sooner interest rates are expected to rise. Implementing YCT when there is uncertainty about how the economy will evolve can provide clarity about the FOMC’s expected path of rates. However, the responsiveness of the policy path to the evolution of the economy also makes it more difficult to set the initial YCT tenor. In addition, the responsiveness of outcome-based FG to the evolving outlook makes it more challenging to maintain alignment with YCT over time. The horizon of yields being controlled could need to be adjusted—perhaps frequently—as new information shifts the distribution of possible liftoff dates, as illustrated in the memo on financial and macroeconomic impacts. Such adjustments can have benefits and costs. If modifying the time horizon for YCT communicates the FOMC’s new expectation for the FG horizon, shifts may be beneficial in reinforcing the FG. However, if market participants come to expect frequent shifts in the YCT parameters, there could be greater uncertainty about the level of yields or about the FOMC’s degree of commitment, which could increase risk premiums and reduce the benefits of YCT.

A range of strategies are possible for setting the YCT tenor to support outcome-based FG, each with different implications for the FOMC’s communications and policy commitments:

Forecast-based YCT caps yields on securities maturing at or before the date at which the FOMC expects to meet its conditions for liftoff of the policy rate.

As an example, the FOMC could communicate that it intends to maintain the current target range for the federal funds rate at least until the unemployment rate reaches

4 percent, and that it currently projects the unemployment rate to reach this level around December 2023; accordingly, the FOMC would cap yields on Treasury securities maturing before December 2023.³ The FOMC would subsequently adjust the YCT horizon as it updated its projection for the liftoff date.

By communicating the FOMC’s outlook for the likely date of liftoff of the policy rate, forecast-based YCT might help align market expectations with the FOMC’s own outlook. However, forecast-based YCT would require the FOMC to have a Committee forecast, and this forecast would need to extend over a horizon long enough to include the likely liftoff date. Agreeing on a singular Committee forecast for the path of rates could be difficult, especially if the economic outlook is highly uncertain or the FOMC’s objectives are not likely to be met for several years.⁴ By contrast, outcome-based FG alone does not require such a forecast, as the FG specifies only how the FOMC will react to certain outcomes whenever they occur.

If the economy evolved as expected, the target maturity for forecast-based YCT could roll down smoothly as the date of liftoff neared, providing a smooth exit. In practice, news about the economic outlook would likely require revisions in the forecast and in the YCT tenor, with asymmetric implications for the YCT. Negative data would likely push down market yields without any open market operations. But if positive data led market participants to anticipate a reduction in the tenor of the yield cap, they might sell large quantities of securities to the Federal Reserve until the FOMC reduced the tenor.

Visibility-based YCT would cap yields out to a horizon where the FOMC is reasonably certain that liftoff is not likely. For example, the FOMC could communicate that it intends to maintain the current target range for the federal funds rate at least until the unemployment rate reaches 4 percent, and that it would cap the 2-year yield because it is very likely to take more than 2 years for unemployment to fall to 4 percent.

Communicating the time horizon over which the FOMC is highly unlikely to raise rates would require less precision in forecasting than forecast-based YCT. Accordingly, qualitative outcome-based FG, which does not give quantitative thresholds for increasing the policy rate, could more easily be paired with visibility-based YCT. As yields would be capped only over a horizon where the policy rate path is relatively certain, visibility-based YCT would also be less vulnerable to the asymmetric risks posed by incoming

³ The choice and level of the unemployment threshold as well as the December 2023 forecast are illustrative; other threshold variables such as inflation, other threshold levels, and other forecasts could be used.

⁴ As forecast-based YCT is tied to a point forecast, it could be difficult to communicate that the YCT program would help to manage downside risks to the economy—even though yields would automatically adjust in response to negative shocks to the outlook. In addition, as noted in the memo on financial and macroeconomic impacts, YCT would force financial market prices to agree with the FOMC’s economic forecast, not just with the FOMC’s policy rate commitments.

information. Correspondingly, however, visibility-based YCT would provide less insurance against undesired increases in yields. In addition, if the YCT horizon were significantly shorter than the horizon at which the FOMC’s objectives are likely to be achieved, visibility-based YCT could be seen as timid or could undermine the FOMC’s commitment to maintaining policy rates at the effective lower bound as long as necessary.

Under visibility-based YCT, the FOMC could for some time roll forward the YCT horizon at each meeting, continuing to cap yields over a constant number of years as long as liftoff was not likely within the specified window.⁵ However, as liftoff became more imminent, the YCT program would need to be removed or the parameters would need to adjust to incoming information about the economic outlook. At that stage, the same basic asymmetry in effects of a changing outlook would emerge.

Lastly, *outcome-based YCT* would cap yields and specify that the yield cap itself would be removed based on realized outcomes. For example, the FOMC might state that it would cap the 2-year Treasury yield at least until unemployment fell below a certain level or inflation rose above a certain level.

Because term yields depend on expected future overnight rates, the cap on term yields would need to be removed or increased well before liftoff on the policy rate—two years before liftoff, in the case of a cap on the 2-year yield. Otherwise, as liftoff on the policy rate approached, investors would also expect a discrete jump in term yields, and large open market operations would be needed to maintain the yield cap. This sequencing poses difficulties for specifying the threshold for ending outcome-based YCT. If the same threshold is used for both YCT and FG, liftoff on the overnight rate would have to be delayed until well after the threshold is reached. Alternatively, if a higher unemployment threshold or lower inflation threshold is specified for the yield cap than for the policy rate, the difference between the thresholds would need to be calibrated such that sufficient time is expected to elapse between reaching them; that calibration would require forward visibility on the pace of economic recovery, and multiple thresholds could be more challenging to communicate. To avoid some of these challenges, the FOMC could communicate a qualitative threshold for removing the YCT program, though the ambiguity of a qualitative threshold might also add to uncertainty.

The choice of tenor for outcome-based YCT is fundamentally different from the choice of tenor for forecast-based or visibility-based YCT. With forecast-based or visibility-based YCT, the tenor is based on the economic and policy outlook—specifically, the predicted time until the FG threshold is achieved, or a time horizon in

⁵ The Bank of Japan and the Reserve Bank of Australia communicate their YCT programs through fixed tenors that roll forward.

which the threshold is unlikely to be achieved. With outcome-based YCT, the FOMC would instead choose a tenor based primarily on how capping or targeting yields at that tenor was likely to be helpful in achieving the Committee's goals, though low rates at the chosen tenor would also need to be consistent with the FG.⁶ For example, the tenor could be one where the FOMC has particular concerns about the risk of a premature rise in yields.

2. The YCT yield level

A YCT program uses open market operations to implement the FOMC's judgment of the yield level associated with its FG policy out to a particular tenor. In contrast to quantity-based asset purchases, YCT maintains a set yield level by allowing the balance sheet to adjust. This section focuses on two questions that determine how yields are controlled. First, what yield level is consistent with the FOMC's FG on the policy rate? Second, how firmly should yields be controlled relative to the level viewed as consistent with FG, and should control be one-sided (a cap) or two-sided (a target)?

A. The yield level consistent with FG on the policy rate

With fully credible and unconditional date-based FG, it is straightforward to calculate the term yields consistent with the FG out to the specified date of liftoff. Under fully credible date-based FG, there should be no uncertainty about nominal interest rates out to that date, and term interest rates should theoretically equal expected average overnight rates, which are constant until the liftoff date. In practice, the term yield might deviate slightly from this calculation as a result of liquidity premiums in the Treasury market.⁷

Outcome-based FG complicates the calculation, even when the FG is fully credible. First, over any horizon where there is a non-zero probability that the FOMC's FG threshold will be achieved, the policy rate target range could increase but is unlikely to decline further. Therefore, the expected average overnight rate over any horizon where liftoff might occur is greater than the current overnight rate. Indeed, the expected average overnight rate depends on the probability of liftoff as well as the distribution of potential rate outcomes should liftoff occur. Second, because the path of overnight interest rates is uncertain, term interest rates should also include term premiums. The appropriate term

⁶ Thus, the forecast-based and visibility-based approaches are similar to Delphic FG, which predicts the future, while outcome-based FG is more comparable to Odyssean FG, which commits to a course of action. See Jeffrey R. Campbell, Charles L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano, "Macroeconomic Effects of Forward Guidance," *Brookings Papers on Economic Activity*, Spring 2012.

⁷ These liquidity premiums have typically been small at the short end of the Treasury yield curve.

premium given the FOMC's FG is difficult to estimate, as it depends on a number of factors, including the economic consequences of the FG itself.⁸

To the extent that the FOMC judges that current term yields appropriately reflect its FG, those yields might also be used as a guide for establishing a yield cap consistent with FG. As an example, the current 2-year yield is currently around 15 basis points, or 6 basis points higher than the federal funds rate. This spread could be interpreted as reflecting some modest probability of a policy rate increase within two years, along with some degree of negative term premiums.⁹

B. Firmness of control: caps, targets, ranges and operations

YCT can be designed only to limit upward movement in yields or to also prevent undue declines in yields. In either case, YCT can allow for modest deviations of yields from the level calculated to be consistent with the FG, or can impose a tight cap (and floor, if desired) on yields.

Yield caps provide the primary benefit typically associated with YCT, by maintaining at least the FOMC's desired degree of accommodation even if market participants misunderstand or doubt the FOMC's commitment or are more optimistic about the economic outlook.

Yield targets or ranges would set a floor on yields as well as a ceiling. The floor would further limit volatility and prevent yields from declining below levels consistent with the FOMC's expected policy stance. In particular, a floor of zero could reinforce the FOMC's commitment to maintain positive policy rates, although to the extent that zero is already a credible lower bound on the policy rate, there may be a floor on term yields not far below current yield levels even if the YCT program does not include a formal floor.¹⁰ While selling securities to put a floor on yields is feasible, sales could be challenging to communicate when the FOMC is trying to maintain an accommodative policy stance. The Bank of Japan has chosen for that reason not to sell securities when yields fell below the floor in its YCT policy.

With either caps or targets, the precision of interest rate control depends partly on the YCT parameter settings and partly on the design of open market operations. The yield

⁸ Term premiums depend on the correlation of interest rates with other outcomes and on the market price of risk. Outcome-based FG changes the correlation of interest rates with macroeconomic variables and could affect the price of risk; the sign and magnitude of these effects depend on numerous factors including how the FG policy itself influences macroeconomic outcomes. The memo on financial and macroeconomic impacts further discusses these issues.

⁹ See the memo on financial and macroeconomic impacts.

¹⁰ See Joseph E. Gagnon and Olivier Jeanne, "Central Bank Policy Sets the Lower Bound on Bond Yields," Peterson Institute for International Economics Working Paper 20-2, January 2020.

cap could be set very close to the level viewed as consistent with FG, or somewhat higher to allow for modest variation in yields. A narrower range could further limit the volatility of yields, which may result in lower term premiums, but could also limit private trading.

Regarding operations, the FOMC could use *hard caps* that commit to purchase securities in unlimited amounts at the specified yield,¹¹ or *soft caps* that set fixed daily purchase amounts to limit movements in yields. In either case, the FOMC's commitment to maintaining YCT, and market confidence in that commitment, would likely be as important an influence on yields as the actual open market operations. If market participants think the YCT and the FG are well aligned and perceive the FOMC's commitment to the YCT as strong, then yields could adjust in advance of operations, resulting in small purchase sizes.

With hard caps, full-allotment operations would likely put a strong ceiling on rates but would offer no control over the quantity purchased on any given day. These operations could result in very large purchases if there were sharp changes in expectations for policy rates, substantial increases in Treasury debt issuance at the capped tenor, or events that precipitated strong selling pressure in Treasuries.¹²

Soft caps can provide more certainty about the daily purchase size and have been effectively used by other central banks; however, they would not shield the YCT program from large *total* purchases over the medium run. Setting each day's purchase amount would require an assessment of market conditions. Yields could, on occasion, vary more than desired. If operations did not achieve the objective, purchase sizes would need to be increased. And if operations were perceived as unsuccessful, investor confidence could diminish, requiring even larger purchases. In practice, fixed-size operations could be used most of the time, with full-allotment operations added when needed to enhance control.¹³

3. Balance sheet implications of YCT

As discussed above, a credible YCT program that is aligned with clear and credible FG should, under most circumstances, lead market prices to adjust without

¹¹ The Desk would translate yields to prices on a range of securities out to the YCT tenor, and commit to purchasing all securities offered from primary dealers at those prices. The range of securities to be purchased is an important choice. For example, it may not be necessary to purchase Treasury bills, which have a strong and unique investor base, or securities that are very close to maturity.

¹² In one operation during the period of stress in early April, the Desk received almost \$60 billion in offers to sell coupon securities with 0 to 2.25 years remaining to maturity. All of these offers would have been accepted if the operation had been conducted on a full-allotment basis. Counterparty or proposition limits could offer some control over quantities but might reduce control over yields.

¹³ The Bank of Japan and Reserve Bank of Australia have used soft caps with generally successful results. The Bank of Japan has conducted full-allotment purchases on occasion to enhance control.

requiring large changes in investors' securities holdings. As a result, such a program might require only modest purchases of securities. However, a YCT program could require significant purchases if investors came to doubt the FOMC's commitment to FG and YCT. The balance sheet expansion in that instance could be viewed as meeting the objective of aligning yields with the FOMC's FG. Other factors could also result in large purchases. For example, increased government debt issuance within the capped tenor, such as happened during the historical U.S. experience with YCT, could put upward pressure on yields that would result in large purchases. YCT could also become misaligned with outcome-based FG if the economic outlook shifted.

There are \$4.1 trillion of Treasury coupon securities outstanding with remaining maturity of 3 years or less, exclusive of Federal Reserve holdings.¹⁴ Given this large investible universe, holdings could expand significantly even with a cap at a short tenor.¹⁵

Purchasing securities in short-dated tenors for YCT may not have the same persistent balance sheet impact as large-scale asset purchases. YCT purchases would not be intended to directly change the amount of duration risk held by the public, but rather to change expectations about yields by demonstrating a commitment to purchases. Accordingly, to achieve the desired effect of YCT, it might not be necessary to reinvest proceeds of maturing securities; the FOMC could conceivably allow securities to roll off the balance sheet even as YCT was still ongoing, reducing the overall expansion of the balance sheet. However, the Treasury would then need to issue more debt to the public, so the appropriate pace of redemptions would depend on how this increased issuance might affect yields. Once YCT ended, the purchased securities would mature and the balance sheet could shrink quickly, assuming the YCT had targeted a short tenor.

¹⁴ This figure excludes floating rate notes (FRNs) and Treasury inflation-protected securities (TIPS), as purchases of these securities would likely not directly support nominal yield control. In practice, the Desk may need to purchase only securities with more than 1 year of remaining maturity. There are \$2.8 trillion in outstanding coupon securities with remaining maturities of 1 to 3 years, exclusive of FRNs and TIPS.

¹⁵ If the Treasury Department responded to YCT by issuing more securities at the target maturity, the outstanding universe would expand. However, Treasury is currently emphasizing longer-maturity issuance. Provided the yield target is not unreasonably low or inconsistent with market yields at other points on the curve, there should not be strong incentives to change debt management practices.