

Methods of Policy Accommodation at the Interest-Rate Lower Bound

Michael Woodford

Recent events have confronted many of the world's leading central banks with a situation that was regarded a few decades ago as merely a theoretical curiosity—a situation in which they have reached a lower bound on the level to which they are able to push overnight interest rates, despite an undesirably low level of capacity utilization, and low inflation or even fears of deflation. The theoretical possibility of reaching such a situation first became an all-too-real challenge for the Bank of Japan in the late 1990s, when even an eventual reduction of the BOJ's target for the call rate (the overnight rate that had been its operating target until then) to zero was insufficient to halt deflation in Japan. But in the wake of the global financial crisis, other central banks, notably including the Federal Reserve, have found that even reductions of their policy rates to the lowest levels that they are willing to contemplate have been insufficient to spur satisfactory recoveries. Most worrisome of all for the Fed is the fact that, as with Japan, the situation has proven not to be merely a momentary anomaly; instead, slow growth and lower-than-desired inflation have continued, despite a zero to 25-basis-point target band for the federal funds rate since December 2008, and there is little optimism about exit from the situation within the coming year.

It is true that, in these more recent cases, one cannot quite say that overnight rates have reached their lowest feasible levels, as was arguably true of Japan. What we have seen in countries like the United States is a situation in which overnight rates are reduced to (or even slightly below) the rate of interest paid on overnight balances at the central bank, so that further expansions of the supply of bank reserves cannot bring about any additional material reduction in the level of overnight rates, given the rate of interest paid on reserves.¹ The rate of interest paid on reserves is not necessarily at its lowest feasible level, but may be set at a level that the central bank is unwilling to go below, because of fears about the consequences for the functioning of the money markets of further shrinkage in the small spreads that remain. This is a prudential concern, rather than an issue of technical feasibility;² but to the extent that a central bank determines that such concerns are important, it establishes an effective lower bound on the policy rate that may be slightly above the technical lower bound, and the considerations discussed below become relevant. And in any event, even if a further reduction in the rate of interest paid on reserves should be listed among the available options for further policy easing in such a case, there clearly is a lower bound on how far the policy rate can be pushed through further reductions in the rate of interest paid on reserves, as long as it remains possible to hold currency that, for institutional reasons, must earn a zero nominal interest rate. Hence, the question whether other options for policy accommodation exist, apart from additional cuts in the current level of overnight interest rates, has become a pressing one for central banks like the Federal Reserve.

This paper discusses two of the main alternatives that have been the focus not only of considerable recent discussion, but a fair amount of policy experimentation, in a number of countries. The first is forward guidance—explicit statements by a central bank about the outlook for future policy, in addition to its announcements about the immediate policy actions that it is undertaking. While this is not necessarily a dimension of policy that becomes relevant only at the interest-rate lower bound, the experience of reaching the lower bound has undoubtedly increased the willingness of central banks like the Fed to experiment with more explicit forms of forward guidance, making

statements about future policy that are both more precise and quantitative and that refer to policy decisions much farther in the future than was understood to be intended in the case of past (relatively cryptic) statements about future policy.

A second broad category of additional dimensions of policy is balance-sheet policies, in which the central bank varies either the size or the composition of its balance sheet, even in the absence of any change in its target for overnight interest rates, rather than operating in financial markets purely for the purpose of implementing its interest-rate target. Some of these additional dimensions of policy are also available in principle even when the policy rate is not at its lower bound, even if some traditional doctrines about prudent central banking, such as the “bills only” doctrine (Luckett 1960) would preclude their use.³ But these too have become a focus of much greater interest as central banks have sought to provide additional policy accommodation after reaching the interest-rate lower bound.

To preview the paper’s main arguments, I shall suggest that in the case of each of these broad classes of unconventional measures, caution on the part of central bankers has frequently led to a preference for versions of the policies that are less likely to be effective. In the case of forward guidance, it has been tempting for central bankers to believe that they can affect financial conditions simply by offering forecasts of likely future policy, while not really tying their hands with regard to future policy decisions. But instead, I shall argue that the most effective form of forward guidance involves advance commitment to definite criteria for future policy decisions.

Similarly, with regard to balance-sheet policies, it has been tempting to believe that it is possible to use the central bank’s balance sheet in a way that is practically equivalent to conventional interest-rate policy—and that can accordingly influence general financial conditions, without involving the central bank in the allocation of credit to particular classes of borrowers, or requiring it to purchase assets outside some narrow class that it conventionally deals in. But instead, I shall argue that the most effective type of balance-sheet policies seek to more directly channel credit to particular segments of the market. I shall further argue that the idea that balance-sheet policies can

be used as a substitute for forward guidance, and, therefore, excuse a central bank from any need to make commitments about future policy, is mistaken. Much of the effect of balance-sheet policies seems to have resulted from their being taken as a signal about likely future policy; and to the extent that this is true, explicit forward guidance will surely allow the desired effect to be obtained more reliably.

I consider first the uses of forward guidance (Section I), then balance-sheet policies focused on the liabilities of the central bank (“quantitative easing,” Section II), and finally balance-sheet policies focused on the composition of the central bank’s assets (Section III). In each case, I begin by reviewing theoretical arguments for the usefulness of the additional dimension of policy in question, and then turn to the evidence regarding their effectiveness that can be gleaned from recent experience. Section IV offers concluding reflections on the challenges currently faced by central banks like the Federal Reserve.

I. Forward Guidance

Even when a central bank is unable, or at any rate unwilling, to further reduce the current policy rate, it remains possible for it to change what it communicates about how the policy rate is likely to be set *in the future*. This provides, at least potentially, an additional dimension of policy. But how should it be used? Does not prudence counsel that a central bank should speak as little as possible about what it might do under circumstances that it has not yet reached? And if forward guidance is to be provided, what form is most likely to have desirable short-term effects without unnecessarily distorting policy decisions later? I shall first consider theoretical reasons to provide forward guidance, and then consider the available evidence regarding its effectiveness in practice.

I.i Relevance of Forward Guidance in Theory

Should it matter at all what a central bank may say about future policy decisions, as opposed to what it actually does, or what it may announce about actions that it has already determined to take, as soon as they can be implemented?⁴ It is important to recognize first that according to standard macroeconomic theory, people’s expectations about future policy are a critical aspect of the way in which

monetary policy decisions affect the economy. The overnight interest rates (such as the federal funds rate in the United States) that central banks seek directly to influence through their routine market interventions—and decisions about which were the main focus of monetary policy deliberations, before the interest-rate lower bound was reached—are not in themselves of such import for the economic decisions (about spending, hiring and price-setting) that the central bank ultimately wishes to influence.

By this I mean that the level of the overnight rate for the next month or so (which is all that is ordinarily decided upon at a given meeting of the policy committee) would not greatly affect these decisions, in the absence of any change in expectations about short-term interest rates farther in the future. It is instead the anticipated path of short-term rates, years into the future—as well as longer-term interest rates, the exchange rate, and other asset prices, all of which should be linked by arbitrage relations to the expected *path* of short-term interest rates, rather than being determined simply by the current *level* of short rates—that is a more important determinant of these decisions. Hence, even under historical approaches to monetary policy that did not involve much central-bank communication, the fact that policy-rate decisions were able to move markets and the economy as much as they did should be attributed mainly to the fact that a change in the current policy rate would typically have been taken to have implications for the forward path of interest rates as well, extending far beyond the next scheduled meeting, even if the central bank did not explicitly comment on this.

It follows from this view that, even when the current policy rate is constrained by the lower bound, a variety of different short-run outcomes for the economy should remain possible, depending on what is expected about future policy. Indeed, theory implies that expectations about future policy should matter even more than usually in that circumstance—or more precisely, when not only is the lower bound a currently binding constraint, but there is reason to expect that it may continue to constrain policy for several more quarters. The reason is that an expectation of an unchanged nominal interest rate for several quarters, that will be largely insensitive to the precise

evolution of aggregate conditions over that time, creates a situation in which expectations of aggregate conditions *after* the interval over which the nominal rate is expected to be fixed have a particularly large effect on the current economy.

Standard New Keynesian models imply that a higher level of expected real income or inflation in the future creates incentives for greater real expenditure and larger price increases now;⁵ but in the case of a conventional interest-rate reaction function for the central bank, short-term interest rates should increase, and the disincentive that this provides to current expenditure will attenuate (without completely eliminating) the sensitivity of current conditions to expectations. If nominal interest rates instead remain unchanged, the degree to which higher expected real income and inflation later produce higher real income and inflation now is amplified. If the situation is expected to persist for a period of time, the degree of amplification should increase exponentially. Hence, it is precisely when the interest-rate lower bound is expected to be a binding constraint for some time to come that expectations about the conduct of policy after the constraint ceases to bind should have a particularly large effect on current economic conditions—to the extent, that is, that it is possible to shift expectations about conditions that far in the future.⁶

But even granting that expectations about future conditions should matter, can a central bank's forward guidance do anything to change them? There are two reasons why it should matter what the central bank says about its future policy. The first is that, even in the case of a clear intention on the part of the central bank, it may not be easy for its intentions to be discerned by the public, and for their implications for likely future outcomes to be understood, without explicit guidance from the central bank. This is especially likely to be an issue if what one wants people to expect is that, following a period in which the interest-rate lower bound has required policy to be tighter than would otherwise have been desired, policy will be looser than it would otherwise have been (so that the expectation of looser policy later mitigates the effects of the undesirably high short-term real rates while the constraint binds).

In such a case, one wants people to understand that the central bank's policy will be *history-dependent* in a particular way—it will behave differently than it usually would, under the conditions prevailing later, simply because of the binding constraint *in the past*. But this is a complex type of behavior for people to have come to anticipate simply from observing the bank's typical conduct, and the situation in question is one that has seldom, if ever, arisen before. Moreover, if the intention to behave in this way going forward is formulated only after the lower bound has been reached, one would be wishing for people to understand an intention that could not actually be put into practice until later. This is unlikely to occur without explicit discussion by the central bank of its intention to conduct policy later in the history-dependent way.

A second reason why forward guidance may be needed—that again has particular force when the interest-rate lower bound is reached—is in order to facilitate *commitment* on the part of the central bank. As Krugman (1998) emphasizes using a simple two-period model, and Eggertsson and Woodford (2003) show in the context of a more fully articulated dynamic model, the future policy that one wishes for people to anticipate is one that the central bank will not have a motive to implement later, if it makes its decisions then in a purely forward-looking way, on the basis of its usual stabilization objectives. Hence, a desirable outcome requires commitment, just as in the analysis of Kydland and Prescott (1977)—even though in this case, the problem is a lack of motive *ex post* to be as expansionary as one wanted people earlier to expect, rather than a lack of motive *ex post* to control inflation as tightly as one wanted them to expect. In practice, the most logical way to make such commitment achievable and credible is by *publicly stating the commitment*, in a way that is sufficiently unambiguous to make it embarrassing for policymakers to simply ignore the existence of the commitment when making decisions at a later time.

These considerations establish a straightforward case for the benefits that should be attainable, at least in principle, from the right kind of advance discussion of future policy intentions. On the other hand, some caution is appropriate as to the conditions under which such

an approach should be expected to work. It does not make sense to suppose that *merely expressing* the view of the economy's future path that the central bank would currently wish for people to believe will automatically make them believe it. If speech were enough, without any demonstrable intention to act differently as well, this would be magic indeed—for it would allow the central bank to stimulate greater spending while constrained by the interest-rate lower bound, by telling people that they should expect expansionary policy later, and then also fully achieve its subsequent stabilization objectives, by behaving in a way that is appropriate to conditions at the time and paying no attention to past forecasts. But there would be no reason for people believe central-bank speech offered in that spirit.

Hence, it is important, under such an approach to policy, that the central bank not merely give thought to the future course of conduct that it would like for people to anticipate, and offer this as a forecast that it would like them to believe. It must also think about how it intends to approach policy decisions in the future, so that the policy that it wants people to anticipate will actually be put into effect, and about how the fact that this history-dependent approach to policy has been institutionalized can be made visible to people outside its own building. These matters are not simple ones, and require considerable attention to the way the central bank communicates about its objectives, procedures and decisions. The problem is all the more difficult when one must communicate about how an unprecedented situation will be dealt with.

1.ii Effectiveness of Forward Guidance in Practice

It seems clear enough in theory that, if a central bank *can* influence expectations about future policy, this should be an important addition to its tool kit. But to what extent are central-bank announcements actually able to influence expectations in the way that a central bank desires? The question is not a simple one to answer, but recent events provide many more examples of attempts at forward guidance, so that at least some grains of empirical evidence are now available.

I.iii Does Central Bank Speech Matter?

A first empirical question is simply, how confident can we be that attempts at forward guidance matter at all? Do statements by a central bank actually change the expectations of market participants, and hence economic outcomes, or do only the bank's actual trades matter, and not what it may say about them? The most influential approach to this question has been the one pioneered by Gürkaynak, Sack and Swanson (2005). Their work looks at whether market expectations of the forward path of the U.S. federal funds rate seem to change over a narrow time window around the release of a post-meeting statement by the Federal Open Market Committee (FOMC); the idea is that if the window is narrow enough, one can be fairly confident that the only important "news" that should have changed expectations over this time interval was the news in the FOMC statement.

The method cannot, by its nature, reveal anything about *why* market participants forecast a different forward path for interest rates after release of the statement, or which aspect of the statement constitutes the news that changes their beliefs; but it can test the null hypotheses that FOMC announcements do not change the expectations of market participants at all (that speech is irrelevant), or that the only news in a post-meeting statement is the revelation of the new (current) operating target for the federal funds rate. Any effects on market prices during a sufficiently narrow window must indicate an effect of the speech, since the Fed will not yet have conducted any trades to implement the new policy; and even over a longer window (say, a two-day window), any market movements that cannot be predicted by the news about the new operating target alone must indicate an effect of speech, since the change in the Trading Desk's behavior in the market will depend only on the new operating target. Movements of the latter kind further provide evidence that the announcement of the new target is not the only kind of speech that influences expectations, and so justify consideration of what else a central bank might speak about.

Gürkaynak et al. use changes in fed funds futures prices to infer the change after each announcement in market expectations for the funds rate at various future horizons. They use principal components analysis to extract the two most important “factors” explaining movements in the forecasted funds rate at the various horizons, and orthogonalize these two factors so that the loading on one factor (the “target” factor) is equal to the change in the forecast of the current fed funds target (the one that will apply immediately after the meeting), while the other factor (the “path” factor) involves no change in the forecast of the current target, only changes in forecasts of the funds rate at horizons farther in the future. Under the null hypothesis of no effect of the statements on expectations, there should be no appreciable variation in either factor. Under the null hypothesis that the only news is the revelation of the current target, all variations in the forecasted path of the funds rate should be accounted for by the “target” factor alone.

Instead, Gürkaynak et al. find that the “path” factor accounts for an important degree of variation in funds rate forecasts.⁷ More recently, Campbell et al. (2012) extend the work of Gürkaynak et al. to a longer data sample, and find similar results. For their sample of statements between February 1994 and June 2007 (i.e., from the time that the FOMC began issuing a statement about the policy decision after each meeting, until the onset of the subprime crisis), they find that the “path” factor accounts for 67 percent of the variation in the expected funds rate two quarters in the future, and 90 percent of variation in the expected funds rate four quarters in the future. For their sample of statements between August 2007 and December 2011 (treated separately because of the numerous novel aspects of communication policy during and since the crisis), the “path factor” is associated with changes in the expected funds rate farther in the future, but continues to be important: it accounts for 53 percent of variation in forecasts four quarters in the future, and 79 percent six quarters out.

This indicates that FOMC announcements were able to shift expectations about the future path of the funds rate, and not simply through the announcement of a new current target. Some other

aspect of the announcement must have been conveying information about future policy, over and above whatever inference about future policy could be made on the basis of the new funds rate target itself. These changes in expectations about future policy furthermore affected behavior, at least in asset markets, for Gürkaynak et al. also find that their “path” factor is correlated with changes in Treasury yields over the same time window. Campbell et al. confirm this, and also find highly significant effects on corporate bond yields.

Nonetheless, an important limitation of this approach is that it provides no information about *what aspect* of FOMC statements influences expectations. Do market participants accept at face value what the FOMC declares about future policy, or do they form their own inferences about likely FOMC policy from other clues in the statements? More importantly, do forecasts of the future funds rate change because *beliefs about the FOMC’s reaction function* change as a result of the statement, or because *forecasts of future economic conditions* that are expected to determine FOMC policy change, as a result of inferences that are made about information that must be available to the FOMC? The latter question is important in order to determine whether statements can change expectations about the way that a central bank will conduct policy in the future, the goal of “forward guidance.”

In at least some cases, the timing of the forecast changes does coincide with attempts by the FOMC to provide explicit forward guidance about policy. For example, Campbell et al. note that the largest value of the Gürkaynak et al. (2005) path factor occurred on Jan. 28, 2004, which was a meeting at which the funds rate target (which had been held constant at a floor of 1 percent since the previous June) was not changed, but the reference to maintaining policy accommodation “for a considerable period,” included in each post-meeting statement since the previous August, was replaced by a declaration that “the Committee believes it can be patient in removing policy accommodation.” It seems likely that the substantial change in funds rate expectations (despite no change in the current target and no surprise in that regard) was mainly due to this change in language, which was evidently taken to indicate that the FOMC would begin

raising the funds rate target soon than had previously been expected. But even in such a case, one cannot easily say whether this reflected successful signaling of a change in the FOMC's reaction function, or simply an inference that the change in language indicated that the FOMC's information predicted a stronger economy.

Reasons for doubt are provided by the results of Campbell et al. on the extent to which the news in FOMC statements predicts revisions (in the next month's survey) of forecasts of unemployment and CPI inflation in the Blue Chip Economic Indicators forecast survey. They find that positive values of both the target factor and path factor are associated with *downward* revisions of unemployment forecasts, and *upward* revisions of inflation forecasts, in the next month's survey after the FOMC statement in question. Both signs are opposite to what one would expect if the news that lead to a higher expected path of the federal funds rate was a shift in the FOMC reaction toward tighter policy under given economic conditions, but exactly what one would expect if there were no change in beliefs about the reaction function, but news that the economy was likely to be stronger than previously expected. Of course, there could be some news of both kinds; but one cannot say that these results provide clear evidence of an ability to change beliefs about the reaction function.

This is a pervasive problem with attempts to infer from the empirical evidence what the effects of forward guidance have been; but it is particularly severe when there is no way to judge what sort of signal about future policy a given central-bank announcement should have been. For this reason, in what follows, I shall focus on occasions on which central banks not only made public statements, but deliberately attempted to send a particular message about future policy.

1.iib Consequences of Explicit Forward Guidance

The occasions during the recent crisis on which central banks have indicated that they expected to maintain a fixed policy rate for a specific period of time are of particular interest for purposes of our inquiry. These are especially dramatic examples of attempts at forward guidance, making a clear break from "business as usual;" moreover, the import of what is said for the future path of the policy rate

is quite explicit and easily summarized. It is, therefore, of interest to consider what has happened on these occasions, even if one cannot do formal hypothesis tests with such a small sample of events, each rather unique.

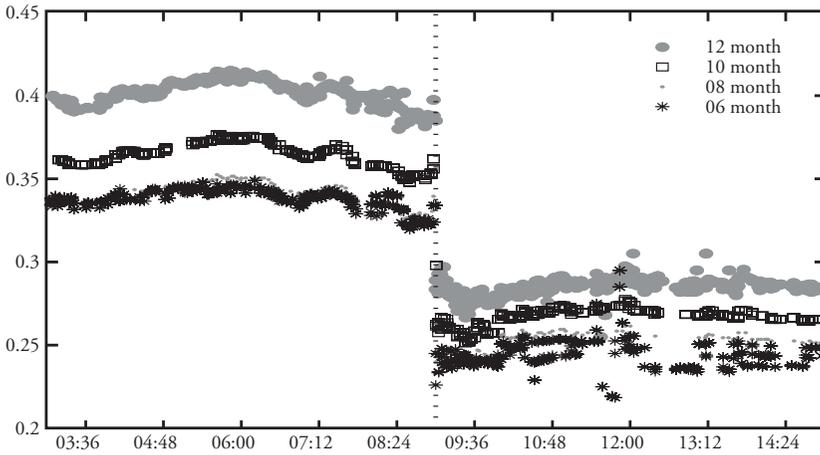
A particularly explicit example of forward guidance was the Bank of Canada's statement on April 21, 2009, which announced the following:

The Bank of Canada today announced that it is lowering its target for the overnight rate by one-quarter of a percentage point to 1/4 per cent, which the Bank judges to be the effective lower bound for that rate.... With monetary policy now operating at the effective lower bound for the overnight policy rate, it is appropriate to provide more explicit guidance than is usual regarding its future path so as to influence rates at longer maturities. Conditional on the outlook for inflation, the target overnight rate can be expected to remain at its current level until the end of the second quarter of 2010 in order to achieve the inflation target.

While the statement included the announcement of a reduction in the current target rate, it *also* offered explicit guidance about where the target should be expected to be, extending more than a year into the future. The release of the statement had an almost instantaneous effect on market expectations about the future path of the policy rate, as indicated by trading in overnight interest-rate swap (OIS) contracts (Chart 1).

The tick-by-tick transactions data plotted in Chart 1 show that market OIS rates fell almost instantaneously at the time that the announcement was made (9 a.m. EST, shown by the vertical line). This was evidently an effect of the statement; yet since the statement included the announcement of an immediate target rate reduction, one might wonder if the moves in the OIS rates reflected simply the typical implications of a cut in the current target for rates months in the future, rather than any additional effects of the "conditional commitment." It is useful to note not only that OIS rates for maturities as long as six to 12 months fall, but that the longer maturities *fall*

Chart 1
Intraday OIS Rates in Canada on April 21, 2009



Note: The dotted vertical line indicates the time of release of the Bank of Canada's announcement of its "conditional commitment" to maintain its policy rate target at 25 basis points through the end of the second quarter of 2010.
 Source: Bloomberg.

more; that is, not only does the OIS yield curve fall in response to the announcement, but it *flattens*. This implies either that expectations of policy rates for months in early 2010 fall even more than do near-term expectations, or that uncertainty about the path of the policy rate over the coming year has been substantially reduced (reducing the term premium). Either of these interpretations is a plausible consequence of the Bank's unprecedented (albeit conditional) commitment to a particular value for the policy rate over the coming year, on the assumption that it is (at least partially) *believed*; neither would be expected to follow from a simple announcement of a cut in the current policy rate, which would typically steepen the yield curve.

The apparent effect on expected future interest rates persisted for at least several weeks following this announcement. Chart 2 plots the path over the course of 2009 of a forward rate $f_t^{(+6,12)}$ defined implicitly by the equation

$$(1 + i_t^{(12)})^{12} = (1 + i_t^{(6)})^6 (1 + f_t^{(+6,12)})^6,$$

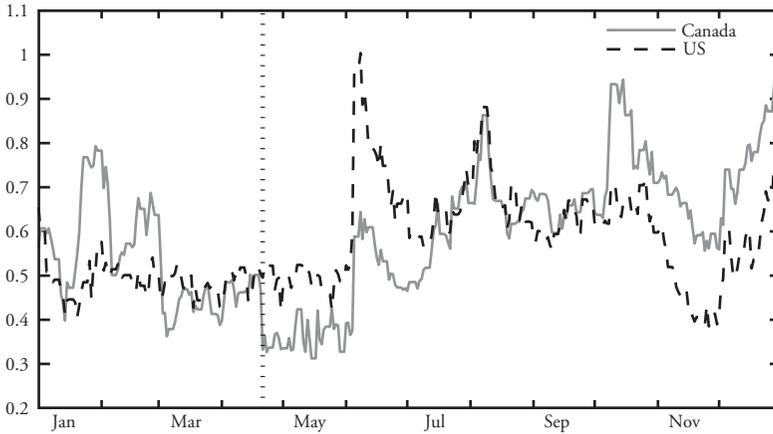
where $i_t^{(n)}$ is the n -month OIS rate. If the n -month OIS rate is interpreted as a market forecast of the average overnight policy rate

over the next n months,⁸ then $f_t^{(+6,12)}$ would correspond to the market forecast of the average policy rate over a time window between six and 12 months in the future. The chart shows that this forward rate falls by 10 to 15 basis points on the date of the announcement (shown by the vertical line), and also that it remains at roughly its new level for the next several weeks. Moreover, there is no similar decline in the corresponding U.S. forward rate during those weeks (as noted by Chehal and Trehan 2009); this suggests that changed expectations about future Bank of Canada policy, rather than news about the economic outlook (which is typically highly correlated with the outlook for the U.S.) are responsible.

This seems a fairly clear example of interest-rate expectations being changed by explicit forward guidance from a central bank. It should not surprise one that the clearest such evidence occurs in the case where a central bank most clearly indicated its intention to provide such guidance—both referring to its statement as having made a “conditional *commitment*”⁹ rather than simply offering a forecast, and stating its intention to “provide more explicit guidance” in order to “influence (longer-term) rates.” Yet even in this case, market beliefs do not simply come to accept that the announced path for the policy rate will be followed with certainty. One observes in Chart 1 that while the OIS rates for maturities between six and 12 months all fall, the rates for 10-month and 12-month maturities do not fall all the way to 25 basis points, even though the announced path involves a policy rate of 25 basis points extending more than 12 months into the future.

One might say that this means that the Bank’s commitment is not completely credible. Actually, the Bank did not purport to make an ironclad commitment; it consistently refers to having made a “conditional commitment,” and the conditionality on “the inflation outlook” is clear in the part of the statement quoted above. It appears that, at the time of the announcement, the escape clause was not expected to be invoked with any very great probability within the coming six months, but that a somewhat higher chance of a rise in inflation triggering early termination of the commitment was allowed for over the 12-month horizon.

Chart 2
The Forward Rate Implied by Term Structure of OIS Rates



Notes: The forward rate (for the period between six and 12 months in the future) implied by the term structure of OIS rates (see text for explanation), for both the Canadian dollar and the U.S. dollar, over the course of 2009. The dotted vertical line marks the date of the announcement of the Bank of Canada's "conditional commitment." Daily data. Source: Bloomberg.

One also observes from Chart 2 that, during the first week of June the forward rate shot up again, to a level greater than 50 basis points (and higher than in the period before the "conditional commitment"). Since at this point in time, the period to which the commitment applied still included all of the next 12 months, one can only conclude that markets had developed more serious doubts about whether the policy rate would really remain at the floor through June 2010. These seem to have resulted from developments in the U.S.; the figure also shows that the corresponding U.S. forward rate shot up by an even larger amount. The spike in U.S. OIS rates occurred on June 5, 2009, in response to a better-than-expected U.S. Department of Labor report that "raised hopes" that the U.S. economy was "on the road to recovery," according to the *Financial Times* (Guha et al. 2009), and resulted in "the futures market pricing in at least one rate increase by the Fed by the end of the year," despite protests by Fed officials that such talk was premature. Traders in Canadian dollar OIS contracts were evidently either skeptical that the Bank of Canada would fail to follow such a move by the Fed, or expected that rapid improvement in the U.S. economy would bring similar consequences for the Canadian economy, and, hence, a change in the

outlook for Canadian inflation. In the latter case, they did not necessarily disbelieve the conditional commitment; but it became less the determinant of their interest-rate expectations, as the likelihood of the relevance of the escape clause increased.

The recent experiments of the Federal Reserve with announcements that the federal funds rate is expected to remain at its current floor for a stated period of time have similarly had measurable effects on market expectations of the future path of the funds rate, as illustrated for example by OIS rates. As I discuss further in the next section, these statements by the FOMC have had less of the character of an announcement of a policy *intention* than was true of the Bank of Canada's "conditional commitment"; instead, the FOMC has been careful only to offer a forecast of what is most likely to occur, given its current information. Nonetheless, these statements as well have clearly moved market expectations.

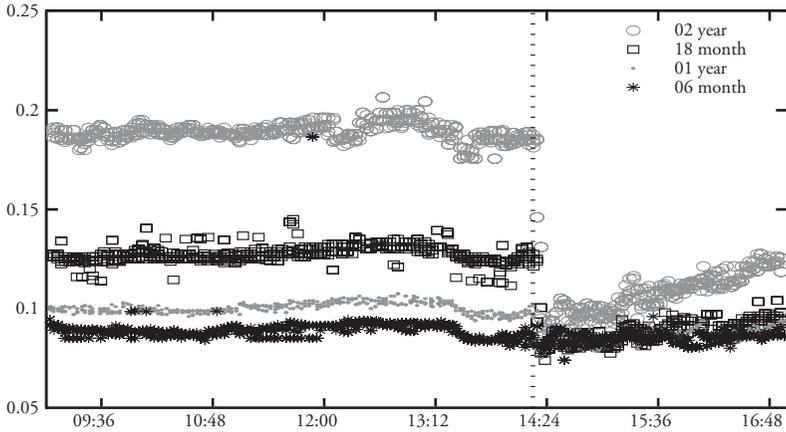
The FOMC began using forward guidance as soon as the zero lower bound was reached. In its post-meeting statement released on Dec. 16, 2008, it announced that the funds rate target was being cut to what has thus far been its lower bound, namely a band between zero and 25 basis points (with interest being paid on reserves at a rate of 25 basis points); but the same statement announced that this level of the target was expected to be maintained "for some time." In its statement of March 18, 2009, this declaration was strengthened (without any change in the target band), to state that conditions were likely to warrant a low funds rate "for an extended period." (These indications, not specifying an exact time period, were similar in style to the FOMC's reference, beginning in August 2003, to maintaining accommodation "for a considerable period," as an alternative to further cuts in the current funds rate target.¹⁰) A more aggressive form of forward guidance was first adopted in the statement of Aug. 9, 2011, in which the main news was the line: "The Committee currently anticipates that economic conditions...are likely to warrant exceptionally low levels of the federal funds rate at least through mid-2013." The forward guidance was further strengthened in the statement released Jan. 25, 2012, to say "...at least through late 2014."

Each of the four statements just mentioned led to a lower expected path for the federal funds rate, as indicated by the response of OIS rates at the time of the release. The “cleanest” tests of the effects of forward guidance were the last two instances; not only did these statements both include very precise specifications of a future funds rate path quite far into the future—that in each case made a stronger statement than the Committee had previously been willing to make, and came as something of a surprise—but in these cases, unlike the first two, the statement did not also contain important policy changes of any other sort at the same time.¹¹ Charts 3 and 4 show intraday data for U.S. dollar OIS contracts, on the days that these two statements were released. In each case, there is a clear, immediate effect on expectations of the future path of the funds rate: OIS rates fall, despite the fact that the current funds rate target remained unchanged.

Moreover, there is a clear flattening of the OIS yield curve in each case. In Chart 3, the six-month OIS rate is essentially unaffected (it continues to trade in the area of 9 basis points); this makes sense, given that the FOMC had already indicated that its existing target (which had resulted in a funds rate a little below 10 basis points) should be maintained “for an extended period” (evidently taken to mean at least six months). Longer-term OIS rates (especially the 18-month and two-year rates) immediately fall, however, to levels barely above 10 basis points; this is what one would expect if market participants believed that the FOMC would with high probability maintain its current target for two years into the future. In Chart 4, the one-year OIS rate (now trading just above 10 basis points) is barely affected; this makes sense, given that the FOMC’s existing forward guidance already extended more than a year into the future (“at least through mid-2013”). The two-year, three-year, and five-year rates instead immediately decline; these contracts all relate to periods that were not completely covered by the already existing forward guidance, so that the extension of the horizon through late 2014 should have mattered, if believed, for the pricing of these contracts.

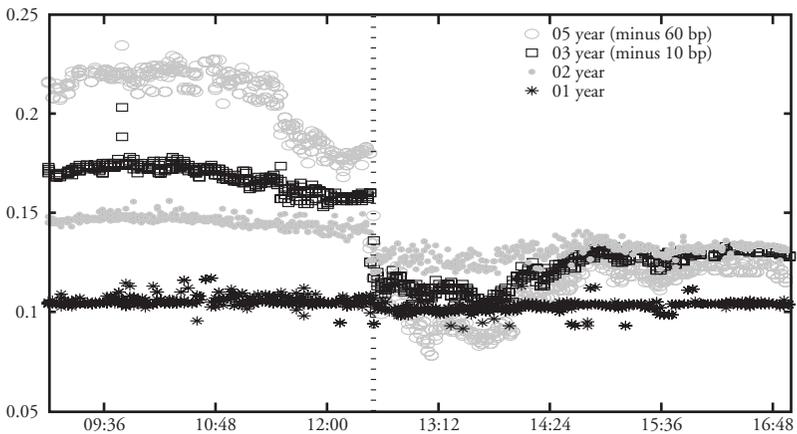
It is true that in Chart 4, the two-year and three-year OIS do not fall all the way to the level of the one-year rate,¹² despite the fact that the FOMC now announced that it anticipated maintaining its target

Chart 3
Intraday U.S. Dollar OIS Rates On Aug. 9, 2011



Note: The dotted vertical line indicates the time of release of the FOMC statement indicating an expectation that the funds rate target would remain unchanged “at least through mid-2013.”
Source: Bloomberg.

Chart 4
Intraday U.S. Dollar OIS Rates On Jan. 25, 2012



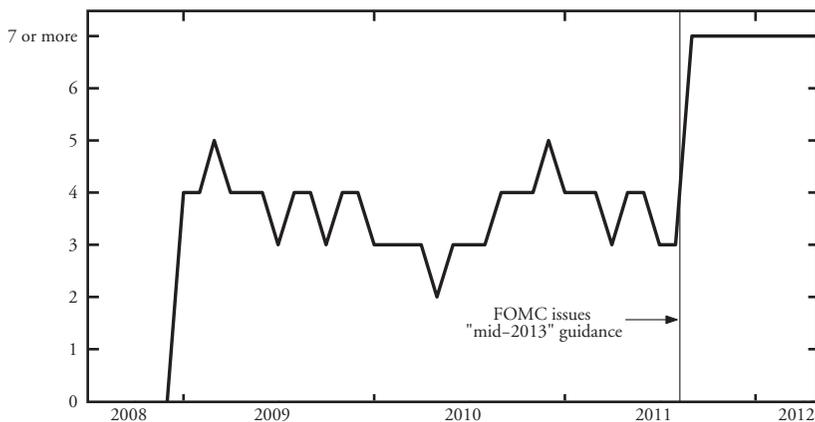
Note: The dotted vertical line indicates the time of release of the FOMC statement indicating an expectation that the funds rate target would remain unchanged “at least through late 2014.”
Source: Bloomberg.

unchanged for a period extending nearly three years into the future. Evidently market participants did not attach a 100-percent probability to maintenance of an unchanged target for that long. But as in the case of the Bank of Canada's forward guidance, one cannot really say that this shows that they did not believe *what they were told*, for the FOMC did not commit itself to maintain the target come what may for that period of time; it stated only that it *anticipated* conditions that would warrant such behavior. (There is a clear implication that not all conditions would.) The statement does seem to have had a definite impact on the expected forward path of the funds rate over a horizon extending years into the future, despite the fact that it was far from an unconditional commitment.

Additional evidence that the FOMC's statements influenced the beliefs of market participants about future policy can be found in the Blue Chip survey of professional forecasters, as noted by Swanson and Williams (2012). Chart 5 shows the median response of survey participants on successive survey rounds to a question about the number of quarters until the FOMC would first increase the federal funds rate target above 25 basis points. After December 2008, when the target was reduced to zero to 25 basis points and the FOMC announced that it should remain there "for some time," the median expectation of the length of time that the target should remain there jumped to four quarters, and it continued to fluctuate mainly between three and four quarters (and never outside the range of two to five quarters) for the next two-and-a-half years. After the FOMC's introduction of the "mid-2013" language in August 2011, instead, the median Blue Chip forecast of the length of time that the target would remain unchanged jumped to seven or more quarters, in accordance with the new FOMC prediction, and has continued at that level (a full year longer than the previous consensus) since then. This indicates a clear effect of the FOMC forward guidance, and suggests that outside forecasters accepted the validity of the FOMC's assessment as the best currently available forecast.

Swanson and Williams present additional interesting evidence of the credibility of the FOMC's explicit forward guidance. Using daily data on interest-rate options with a variety of strike prices and

Chart 5
Charting Blue Chip Survey Responses



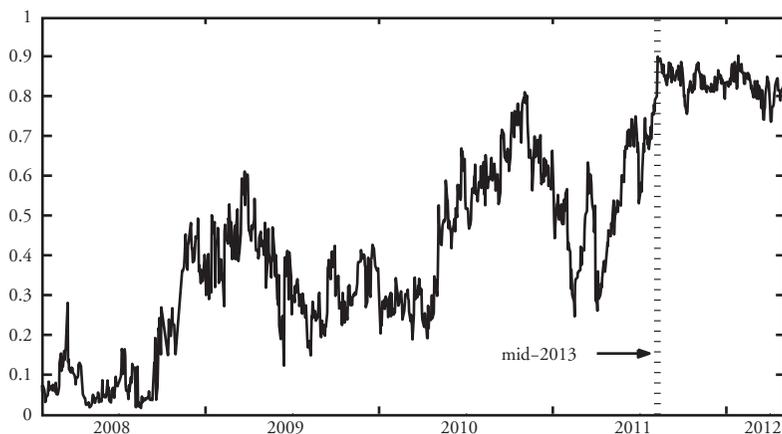
Notes: Median forecast of respondents in the Blue Chip Financial Forecasts survey, of the number of quarters until the federal funds rate target will exceed 25 basis points. Vertical line indicates the release of the first FOMC statement indicating continuing accommodation until “mid-2013.”

Source: Swanson and Williams (2012).

five quarters to maturity, they compute an implied market-expected probability distribution for the federal funds rate five months in the future, for each trading day. The implied probability of a funds rate below 50 basis points five quarters in the future is shown in Chart 6. The probability spikes up, and remains between 80 and 90 percent on most days, after the FOMC’s introduction of the “mid-2013” language, consistent with the consensus of the Blue Chip forecasters shown in Chart 5.

Swanson and Williams also measure the effects of surprises in various types of macroeconomic data releases on eurodollar futures prices. (These contracts settle based on the three-month term eurodollar rate at the date of expiration, and so the price at which such a contract currently trades can be viewed as providing a measure of market expectations of the average level of the funds rate over a three-month window a certain distance in the future.¹³) By looking at how an overall measure of the sensitivity of the futures prices to macroeconomic news varies over time (by plotting the regression coefficient obtained using a rolling window centered at each date),¹⁴ it is possible to observe the degree to which market participants believe that the level of future overnight interest rates will be state-contingent.

Chart 6
Fed Funds Rate



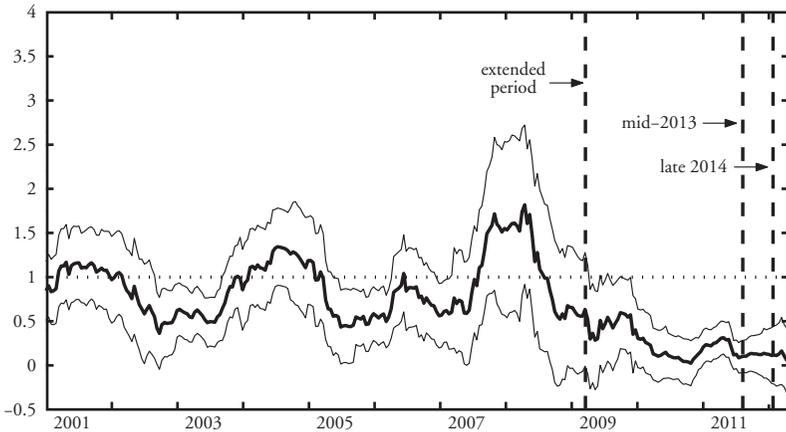
Note: Probability of a fed funds rate below 50 basis points, at a date five quarters in the future, as inferred from interest-rate options prices.

Source: Swanson and Williams (2012).

In periods when FOMC forward guidance forecasts a specific time-dependent path for the funds rate target, the degree of sensitivity of such expectations to news can provide a measure of the degree to which market participants are confident that the announced funds-rate path will actually be followed.

Based on the response of eurodollar futures, they conclude that market expectations regarding overnight rates over a three-month window beginning one to two quarters in the future became substantially less sensitive to macroeconomic news during the period in 2003 when the FOMC cut its funds rate target to unprecedentedly low levels (eventually as low as one percent), but showed no appetite for further cuts; became again about as sensitive as usual in 2004 as anticipation of rate increases (of an unknown timing and speed) grew; and became significantly less sensitive than usual again in 2005 and early 2006, when the FOMC steadily increased its target at the “measured pace” of 25 basis points per meeting. The sensitivity measure has fallen especially sharply during the recent period of increased forward guidance, and has been insignificantly different from zero since the introduction of the “mid-2013” language in August 2011.¹⁵ This suggests that FOMC forward guidance has shaped expectations about the path of the funds rate over the next few months in a way

Chart 7
Sensitivity Index of Eurodollar Futures Prices One to Two Quarters in the Future



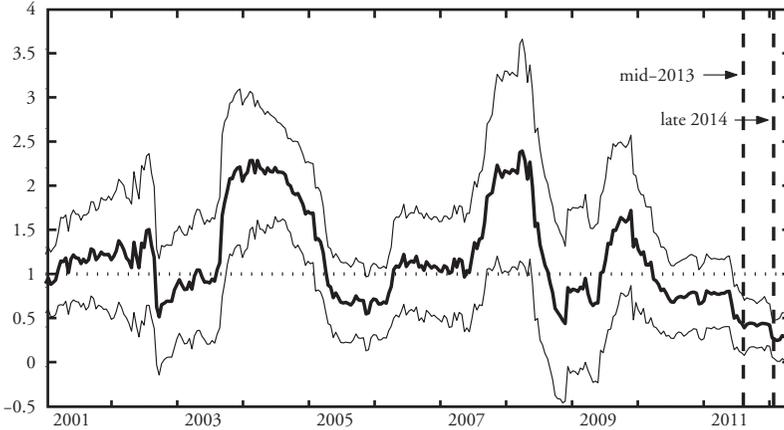
Notes: Index of the sensitivity of eurodollar futures prices to macroeconomic data surprises, in the case of a contract settling one to two quarters in the future. Here 1 (the horizontal dotted line) indicates the mean sensitivity over the period 1990-2000; the thin lines represent a 95-percent confidence interval. Vertical lines indicate the dates of changes in the FOMC's forward guidance.

Source: Swanson and Williams (2012).

that makes such expectations relatively insensitive to other macroeconomic developments.

Expectations about overnight rates farther in the future were instead apparently less affected by the kind of forward guidance used earlier in the decade; the corresponding sensitivity measure based on longer-horizon eurodollar futures is not significantly lower than its average value over the decade at any point during the period 2001-10 (see Chart 8 for an example). However, the sensitivity of expectations over a three-month window beginning four to five quarters in the future falls to a level significantly less than its average value after the introduction of the “mid-2013” language, as shown in Chart 8;¹⁶ and the sensitivity becomes even lower (a small fraction of the normal level, according to the point estimate, and only barely significantly different from zero) after the introduction of the “late 2014” language in January 2012. This suggests that the more explicit (and longer-horizon) form of forward guidance used by the FOMC more recently has been able to create definite expectations about the future path of the funds rate than was possible using its earlier approach to

Chart 8
Sensitivity Index of Eurodollar Futures Prices Four to Five Quarters in the Future



Notes: Index of the sensitivity of eurodollar futures prices to macroeconomic data surprises, in the case of a contract settling four to five quarters in the future. Here 1 (the horizontal dotted line) indicates the mean sensitivity over the period 1990-2000; the thin lines represent a 95-percent confidence interval. Vertical lines indicate the dates of changes in the FOMC's forward guidance.

Source: Swanson and Williams (2012).

forward guidance (to which it had essentially returned in 2009 and 2010).

Of course, one must note that these changes in FOMC forward guidance do not represent controlled experiments; the FOMC's willingness to experiment with stronger forms of forward guidance was a consequence of a continuing stream of discouraging macroeconomic news. Hence while market movements during a very short time window around an announcement can reasonably be attributed to news contained in the announcement, developments since 2010 of the kind shown in Charts 5, 6, 7 and 8 might alternatively be attributed simply to market participants' increasing doubts that conditions would warrant an increase in the funds rate target anytime soon, for reasons unrelated to the FOMC's statements.¹⁷ And even to the extent that one accepts that the timing of the changes in expectations suggests that the FOMC's changes in communication policy were an important part of the news, there remains the question whether what this conveyed was *news about the economic outlook or news about the FOMC's approach to the conduct of policy*. I return to this issue in Section I.iii below.

I.ii.c Consequences of Announcements of Central-Bank Policy-Rate Projections

Further evidence about the extent to which forward guidance not only can *affect beliefs*, but can more specifically cause people to believe *what the central bank says*, is provided by central banks that announce a forward path for their policy rate as a routine part of their communication about their policy decisions. The Reserve Bank of New Zealand (RBNZ) has announced its forecast of future short-term interest rates since 1997, much longer than any other central bank; there is consequently the greatest amount of data on the effects of such announcements in its case.

Moessner and Nelson (2008) test econometrically the degree to which the RBNZ's announcements affect market expectations, using futures contracts for 90-day bank bills (the money-market instrument for which the RBNZ forecasts future yields) deliverable at various future dates as proxies for market expectations. They estimate a regression of the form

$$f_{n,t} - f_{n,t-1} = \alpha + \beta(f_{n,t}^{cb} - E_{t-1}f_{n,t}^{cb}) + \epsilon_t,$$

where $f_{n,t}$ is the futures rate at the end of day t for a contract specifying delivery n quarters in the future; $f_{n,t}^{cb}$ is the RBNZ's forecast of the 90-day bank bill rate n quarters in the future, released on day t ; and $E_{t-1}f_{n,t}^{cb}$ is the market's expectation of what the RBNZ will forecast, the day before the release.¹⁸

Moessner and Nelson estimate a forecasting regression of this kind for values of n from one to six quarters in the future, and find highly significant positive values of β for all values between 2 and 6. (The coefficient is insignificantly different from zero when $n=1$.) However, even when $n>1$, the estimated coefficients are well below 1: they range between 0.17 and 0.22. Thus, while the RBNZ forecasts do seem to influence market expectations, market expectations do not simply jump to coincide perfectly with the Bank's forecast. (This result is consistent with the earlier study of Archer 2005.) There is, of course, no reason why market forecasts *should* coincide perfectly with the RBNZ's announcement, given that it is not announcing a *commitment* to target

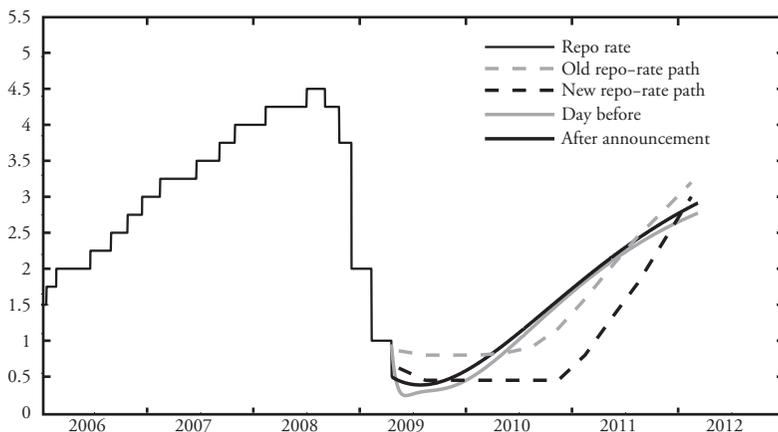
those particular rates at those future dates—only a forecast of what the rates will be, given its current projection of both how the economy is most likely to evolve and how it will conduct policy as a result. The RBNZ’s projections are evidently considered informative, but not dispositive as to what the optimal forecast must be.

More recently, a number of other central banks (led by the Norges Bank in 2005) have begun to regularly release forecasts of the future path of their policy rate. The experience of Sweden’s Riksbank, which has published such forecasts since February 2007, is of particular interest, because the Riksbank has also announced on more than one occasion that its policy rate would remain fixed for a specified period of time—the particular type of forward guidance of greatest relevance to the current discussion.

In a review of Sweden’s experience, Deputy Governor Lars Svensson (2010) argues that, through December 2008, the Riksbank had been relatively successful at “managing expectations” through its policy. Often, he notes, market expectations were already fairly close to the announced forward path for the repo rate [the Riksbank’s operating target for the overnight rate¹⁹] prior to the announcement, which he regards as an indication that the bank had succeeded in conducting a predictable policy and in making the systematic character of its policy evident to the public. “When there were some discrepancies,” he writes, “in most cases the market adjusted its expectations towards the [announced] policy-rate path after the announcement” (p. 48). However, as in the case of New Zealand, this does not mean that market expectations came to perfectly coincide with the path announced by the Riksbank—only that the forward curve that could be inferred from futures rates became closer to the Riksbank’s announced path than it had been.

The effects of the Riksbank’s more recent experiments with announcements of an anticipated duration for the current repo rate have been more mixed. On April 21, 2009 (a few hours before the Bank of Canada announcement discussed above), the Riksbank announced a cut of the repo rate to 50 basis points, together with a statement that “the repo rate is expected to remain at a low level until the beginning of 2011,” a date nearly two years in the future.

Chart 9
Forward Path of Repo Rate in Sweden



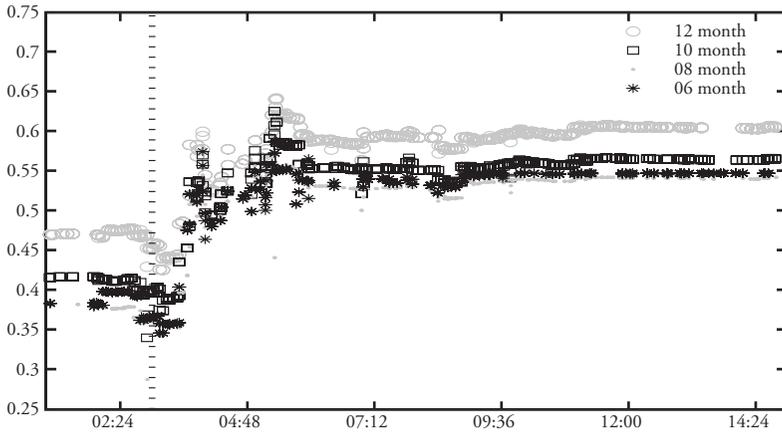
Note: Market expectations of the forward path of the repo rate in Sweden, before and after the Riksbank's press release on April 21, 2009, that indicated that the repo rate was "expected to remain at a low level until the beginning of 2011."

Source: Sveriges Riksbank.

The statement was accompanied by the release of a *Monetary Policy Update*, with a projected forward path that showed the repo rate at a constant level of 50 basis points through the end of 2010, as shown in Chart 9.

The chart shows the actual path of the repo rate as a solid black line (a step function); the projected forward path from April onward that was published on April 21; the market expected forward path, as inferred by the Riksbank on the basis of interest-rate forward and swap rates²⁰ the day before the announcement; and the corresponding market expected forward path after the announcement.²¹ Market participants evidently had expected an even larger cut in the repo rate than occurred, and for the repo rate to remain lower, at least for some months, than was indicated by the projected path. In response to the announcement, the market expected path rose, though still remaining lower than the path projected by the Riksbank, for the first few months after April. By early 2010, market participants had anticipated that the repo rate would already be rising above 50 basis points, whereas the Riksbank projected it to remain at 50 basis points for another year; but in response to the announcement, the market expected path for 2010 rose still further.

Chart 10
Intraday Swedish OIS Rates On April 21, 2009



Notes: Intraday Swedish OIS rates on April 21, 2009. The dotted vertical line indicates the time of the Riksbank's press release (9 a.m. in Sweden, or 3 a.m. EST).

Source: Bloomberg.

The result is that an announcement that was intended to shift down the anticipated forward path of rates, by announcing that a low rate would be maintained until the beginning of 2011, and so to immediately lower longer-term interest rates, had exactly the opposite effect: long rates rose, because the entire anticipated forward path of rates shifted *up*. What went wrong? While many things happened from one day to the next—as noted above, the Bank of Canada introduced its own “conditional commitment” six hours after the Riksbank’s announcement—it seems clear that it was the Riksbank’s announcement that moved market expectations. Chart 10 shows the intraday OIS rates for Sweden on April 21, with the time of the release of the *Monetary Policy Update* shown; the entire term structure of OIS rates moved up within two hours of the release, and well before any news from North America.

What seems to have happened is that market participants took on board *part* of the Riksbank’s forward guidance, and modified their own forecasts to conform more with it: the projection of a path that never fell below 50 points convinced many that (contrary to prior expectations) the Riksbank would not cut the repo rate below that level. This implied an increase in the projected path for the next two quarters. But since the news, as far as market participants were

concerned, was that the Riksbank was less inclined toward interest-rate cuts than they had supposed, the *entire* path was also shifted up.

In fact, the Riksbank's projected forward path contained *two* notable features: it was announced that the repo rate was projected to remain low for nearly two years into the future, *and*, quite remarkably relative to prior figures, it was projected to remain absolutely constant over that time—the only obvious reason for which would have to have been a decision to treat 50 basis points as the effective lower bound. It is true that the April *Monetary Policy Update* contained no announcement that this was a lower bound; it even referred to “some probability of further cuts in the future.” But as Svensson (2010) notes, it also emphasized that “the repo rate is now close to its lower limit,” and stated that “with a repo rate at this level, the traditional monetary policy has largely reached its lower limit.” Moreover, immediately after admitting the possibility in principle of further cuts, it cautioned: “But when the repo rate is at such low levels, one must consider the fact that this could have negative effects on the functioning of the financial markets.” It is easy enough to see how market participants could have read such remarks as indicating an intention by the Riksbank not to reduce the rate below 50 basis points (at least, under any but exceedingly dire circumstances). Such an announcement would, of course, be precisely the sort that should most affect market expectations: because it was interpreted as revealing something not previously known about the central bank's intentions with regard to policy, rather than the central bank's judgments about the economic outlook—and so, a matter about which the bank could undoubtedly be regarded as the most knowledgeable authority.²²

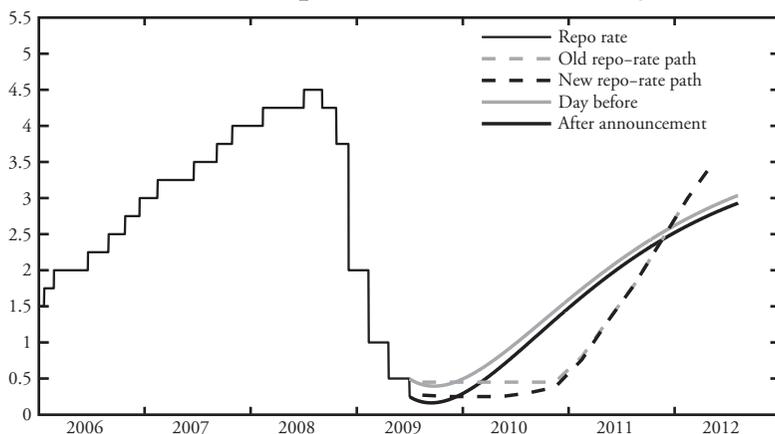
The Riksbank's other message—that it expected not to raise the repo rate before 2011—evidently made less of an impression. One reason might have been an assumption that this reflected the Riksbank's pessimism about the Swedish economy, and market participants might have been more optimistic, and so expected rate increases to be justified sooner than the bank anticipated. Svensson (2010) argues instead that survey data on traders' forecasts of inflation and growth indicate that they were no more optimistic than the Riksbank, and hence that market participants simply did not accept the Riksbank's forecasts about *its own future approach to policy*.

Why might this have been? It is notable that a large (and persistent) discrepancy between the forward paths announced by the Riksbank and those expected by market participants appeared only when the Riksbank began attempting to use projections of a policy rate that would remain fixed for an unusually long time, as a consequence of having reached its (self-imposed) lower bound. One may conjecture that the Riksbank sought, as an alternative to a deeper immediate interest-rate cut, to signal that rates would be kept low for a longer time than would ordinarily have been expected; and this supposition about future policy was incorporated into its projections. But this change in the assumption made about future policy was not credible to market participants, perhaps because no adequate explanation was given of how policy decisions would be made in the future. The mere fact that the Riksbank announced that it projected a low path for the repo rate until 2011 was not enough; market participants needed to have a view of how the Riksbank would make decisions in the future that would justify such a path (given the Bank's expectations regarding the economy's evolution), and evidently they were not provided with one.

Similar problems of credibility seem to have persisted since then. In July 2009, the Riksbank announced a further cut in the repo rate, to 25 basis points, but now only indicated that the target was expected to remain at its low level "until autumn 2010." (This might be considered to vindicate skeptics who had not believed the April projection of a low rate through the beginning of 2011.) As shown in Chart 11, this announcement did shift down market expectations of the forward path, but market participants continued to forecast that the repo rate would not remain at that level past the end of 2009, and expected it to be around 100 basis points by autumn 2010. (In fact, it was only raised to 50 basis points in July 2010 and to 75 basis points in September.) This apparent failure to credit the Riksbank's view of the length of time that the target would remain low made policy effectively tighter (in terms of its consequences for longer-term interest rates and hence for spending decisions) during 2009 than the Riksbank's projection assumed it would be.

Once the Riksbank began tightening policy again, market expectations continued to diverge from the Riksbank's announced forward

Chart 11
Forward Path of Repo Rate in Sweden On July 2, 2009



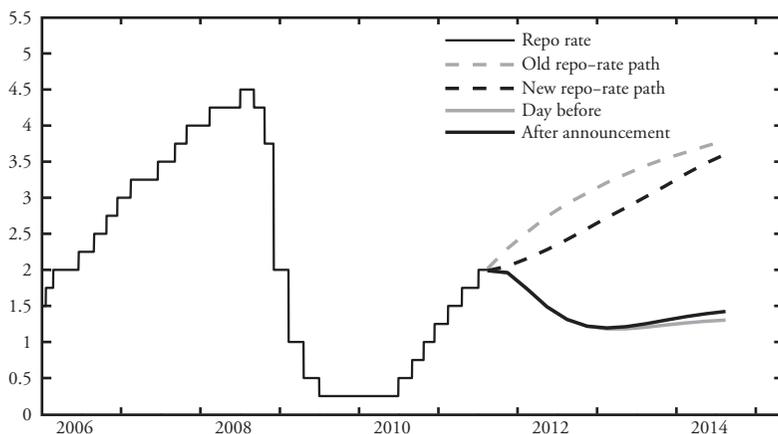
Note: Market expectations of the forward path of the repo rate in Sweden, before and after the Riksbank's press release on July 2, 2009, announcing an additional cut in the repo rate, and a shortening of the time that the low target was expected to be maintained.

Source: Sveriges Riksbank.

paths, but now in the direction of anticipating a *lower* future path for the repo rate than the Riksbank. For example, Chart 12 shows the market expected forward paths before and after the Riksbank's press release on Sept. 7, 2011. In this release, the Riksbank announced that the repo rate target would remain at 2 percent, rather than continuing to increase as it had previously projected,²³ owing to deterioration in global growth prospects. However, this was referred to as only a decision to "postpone continued increases somewhat"; the new, lower repo rate path continued to show the repo rate steadily rising over the next three years. Market expectations prior to the announcement had instead been for cuts in the repo rate to begin by later in the year and to continue through 2012; and the Riksbank's announcement had very little effect on those expectations, despite the reiteration of the Riksbank's expectation that the repo rate would continue on an upward path. In fact, there were no further target increases, and the timing of the first two target decreases (in December 2011 and February 2012) essentially followed the path anticipated by the markets back in September.

Svensson (2011) provides a variety of possible reasons for market expectations of a lower rate path than the one announced by the Riksbank.²⁴ These are all reasons why expectations about future economic

Chart 12
Forward Path of the Repo Rate in Sweden On Sept. 7, 2011



Note: Market expectations of the forward path of the repo rate in Sweden, before and after the Riksbank's press release on Sept. 7, 2011, announcing a "postponement" of further increases in the rate.

Source: Sveriges Riksbank.

conditions might plausibly have differed from the Riksbank's assumptions; for example, he notes that market expectations regarding the future path of U.S. interest rates indicated lower rates than the path assumed by the Riksbank in its projections. Under this interpretation, market participants may have accepted the Riksbank's forecast of how it would behave *if* conditions evolved as it assumed, but doubted that those conditions would be realized. But an alternative possibility is that market participants did not assign much weight to the Riksbank's assertions about its future intentions.²⁵ If so, it would seem that the attempt to use forward guidance more aggressively after April 2009 has been associated with a loss of market confidence in the informative value of the Riksbank's projections. Whether it will return once macroeconomic conditions have normalized remains to be seen.

I.iii What Kind of Forward Guidance Makes Sense?

The above review of recent experience with forward guidance suggests that central-bank statements about future policy can, at least under some circumstances, affect financial markets—and more specifically, that they can affect markets in ways that reflect a shift in beliefs about the future path of interest rates toward the one

announced by the central bank. This seems most clearly to have occurred when central banks that do not ordinarily make statements about policy rates very far in the future departed from their usual policy by stating that rates should remain low for an unusual length of time, owing to having reached their effective lower bound. Perhaps surprisingly, it is less clear how much influence on market expectations central banks have that routinely release detailed projections for the forward path of interest rates.

A possible explanation for this seeming paradox is that forward guidance outside the context of routine predictions about the future path of interest rates is more often interpreted as revealing central-bank *policy intentions*. Information about policy intentions is likely to affect the expectations of market participants more than information about the central bank's view of the economic outlook, because the way in which the bank intends to conduct policy is a matter about which the bank obviously knows more than do outsiders, no matter how closely they follow economic news. And a statement that is viewed as expressing a *commitment*, that by virtue of its having been stated should at least to some extent constrain future policy decisions, should be most informative of all.

The Bank of Canada's "conditional commitment" in April 2009 seems to have been one of the examples of forward guidance that most clearly changed market expectations, and this is also the case in which a central bank came closest to committing itself to a future course of action. The Bank of Canada did not shy away from using the word "commitment" in its press release, even if this was qualified by the word "conditional," and the nature of the conditionality was not fully spelled out. Other central banks, such as the Federal Reserve, have not gone as far; the FOMC's statements have referred only to what the Committee currently anticipates that future conditions will warrant. Yet even in these cases, observers may well have assumed that the unusual announcement made sense only if interpreted as a commitment, and indeed a good deal of commentary interpreted the FOMC's statements this way (and discussed whether the supposed promise was credible). To the extent that reasons are given for a commitment to make sense—as in the case of the Bank of Canada's

explicit reference to its desire to “influence rates” through “forward guidance”—the interpretation as a commitment is also more likely.

Releases of central-bank projections of the path of interest rates, in the context of a more general discussion of the central bank’s forecast of the economy’s evolution over the next few years, are less susceptible to interpretation as a commitment, or even as an expression of a definite intention about future policy that has already been formed. Apart from the fact that the central banks that use this communication strategy take pains to emphasize in the accompanying text that their projections for the policy rate are merely forecasts conditional on current information, the format in which the projections are presented also makes this evident. But to the extent that such projections are viewed simply as following from the bank’s forecast of the economy’s evolution, including a forecast of the evolution of the policy rate given how it is typically adjusted in response to varying economic conditions, then they provide news that should change other market observers’ forecasts of the future path of interest rates only to the extent to which they are thought to reflect superior information about the economic outlook that is available to the central bank. Other close observers of the economy may or may not believe this is true; and even when they do believe they can learn something from what the central bank reveals about its information, their own assessment of the best forecast will, in general, not put a weight of 100 percent on the central bank’s forecast.

I have remarked above that the degree to which market participants have regarded the Riksbank’s projected repo rate path as informative about the likely future path of the repo rate more than a few months into the future seems to have decreased since April 2009, when the target reached a level that the Riksbank was reluctant to go below, and a statement that the target should remain at that rate for a specific (fairly long) time was offered instead of a sharper immediate reduction. This may well have been interpreted as a departure from the bank’s previous practice in the way it produced its projections—but not, evidently, because the bank was now interpreted as making a commitment that it could be counted upon to fulfill.

A possible reason for the reduced credibility of the longer-horizon projections at this point is that this was the first occasion on which the announced path reflected a projection of future policy decisions that were history-dependent to any significant extent—that is, an assumption about future policy that differed from what one would expect that policy to be simply on the basis of conditions at the time. The reason why it would be desirable for policy to be expected to be history-dependent, under precisely the circumstances reached by the Riksbank in April 2009, has already been explained in Section I.i: the anticipation *at the time of the binding lower bound* of a lower subsequent repo rate than *would be desirable on purely forward-looking* grounds at the later date could have beneficial (stimulative) effects at the time of the binding constraint, albeit at the cost of less successful stabilization later. This may well be the sort of calculation that led the Riksbank to choose a repo rate path that indicated low rates so far into the future as it did. But in the absence of any intention to actually make policy decisions in a history-dependent way later—or at any rate, in the absence of an explanation of the procedures that would be followed in the future, that made it *credible* that future policy would be made in that way—there would be no reason for market expectations about the future conduct of policy to change.

The Riksbank's official description of its approach to monetary policy states that “in connection with every monetary policy decision, the Executive Board makes an assessment of the repo-rate path needed for monetary policy to be well-balanced” (Sveriges Riksbank 2010, p. 14). The document goes on to explain the competing considerations that must be taken into account in such an assessment; there is no suggestion that the exercise is anything but a purely forward-looking consideration, repeated afresh in each decision cycle, of which of the feasible forward paths for the economy from that date onward is most desirable, from the standpoint of a criterion that involves both the rate of inflation (and its distance from the official inflation target of 2 percent) and the level of real activity. Indeed, it stresses that the appropriate repo-rate path will be reassessed in each decision cycle, so that “the interest rate path is a forecast, not a promise” (p. 15).

If the model of the economy used in such an assessment of the possible forward paths at a given point in time incorporates forward-looking private-sector behavior—as the Riksbank’s RAMSES model (Adolfson et al. 2007) certainly does—and if the model is solved under the assumption that the projected forward path of the policy rate is anticipated by those forward-looking decisionmakers, then it might easily be concluded that the most desirable forward path at a given point in time is one which assumes history-dependent policy later. This is particularly likely to be the case when the current policy rate is constrained at its lower bound. But in such a case, repetition of the forward-looking exercise at the later date will not result in a decision to continue the interest-rate path previously projected, *even if there have been no surprise developments in the meantime*, for a forward-looking assessment of “well-balanced policy” at the later date will take no account of the effects of expected policy at that date on decisions expected to be taken in the private-sector earlier, according to the policy projections made at the earlier date.²⁶ A purely forward-looking forecast-targeting exercise of such a kind would accordingly be intertemporally inconsistent, as discussed in further detail in Woodford (2012). This means that there would be no reason for market participants to hold the expectations assumed in the projection exercise, even if they perfectly understand the central bank’s decision procedure. The problem might be that they understand it too well—that they have a more accurate forecast of the way that future policy will be made than the one assumed in the projection exercise.

I do not mean to imply that a time-consistent procedure, that assumes that future policy will be determined in a purely forward-looking way, would necessarily be superior. Such a targeting procedure would be intertemporally consistent, but the equilibrium implemented will generally be suboptimal, from the standpoint of the criterion used by the bank itself to rank possible forward paths. In particular, in a situation of the kind described in Section I.i, an inability to commit to a history-dependent policy would mean acceptance of a low-output trap, and of the fact that interest-rate policy can accomplish nothing more once the lower bound on the current overnight rate is reached. What is needed in order to achieve a better outcome, despite a correct understanding of the determinants of

future policy on the part of market participants, is for the central bank to adopt procedures under which it will indeed implement a history-dependent policy, and then to make its intentions clear to market participants. In fact, it does need to offer a “promise” and not merely a “forecast”—though the required form of promise need not be a commitment to a specific pre-announced path for the policy rate.

These comments should also not be taken to suggest that the form of forward guidance recently practiced instead by the Federal Reserve represents an ideal model. While the FOMC’s forward guidance has often been interpreted as making a commitment to keep the funds rate low for a specified period of time, in fact its communication about future policy—both through its post-meeting press releases and through the information about individual participants’ forecasts of the funds rate path in the quarterly Summary of Economic Projections—has taken only the form of predictions about the future path of the funds rate, given what can be known at present.

In particular, no indication of a decision to change the FOMC’s policy rule is ever given; it is, thus, always possible to interpret the FOMC’s announcements about future policy as simply reflecting changes in the FOMC’s view of likely future economic conditions, and hence the path of the funds rate that can be expected under their normal reaction function. For example, when the FOMC announced in January 2012 that “the Committee...currently anticipates that economic conditions...are likely to warrant exceptionally low levels for the federal funds rate at least through late 2014,” the headline of *The New York Times* online story about the announcement was “Fed Signals That a Full Recovery Is Years Away.” While the shift in the OIS yield curve indicates that market forecasts of the funds rate several years in the future fell after the announcement, as shown above, this might have been a response to expectations of a slower recovery rather than to any understanding that FOMC policy had changed.²⁷

Some will undoubtedly protest that a reference to the bank’s current forecasts is the only prudent form of forward guidance for a central bank to offer. If one supposes that the only alternative would have

been for the FOMC to offer an explicit promise to keep the funds rate target at zero to 25 basis points until late in 2014, then one might well think so; a non-state-contingent commitment extending three years into the future would surely have been unwise. The resort to a mere prediction might seem a clever way of allowing for state-contingency without having to explain all of the possible contingencies; we are saying what the path of the funds rate will be *if* things develop in the way that can be anticipated given what we now know, but we make it clear that this is only our *current anticipation*—policy may have to be different if unexpected developments arise.

It is certainly right that a desirable form of forward guidance—if it involves communication about anything but a fairly short horizon—would not make unconditional promises about the future path of the funds rate. Since Campbell et al. (2012) refer to the “late 2014” statement language as implementing “the policy recommendations of Eggertsson and Woodford (2003),” I should point out that Eggertsson and Woodford (2003) do *not* argue for the desirability of a commitment to keep the policy rate at zero for a fixed period of time. We argue for the desirability of a commitment to conduct policy in a different way than a discretionary central banker would wish to, *ex post*, and show that (in our New Keynesian model) the optimal commitment involves keeping the policy rate at zero for some time after the point at which a forward-looking inflation-targeting bank (or a bank following a forward-looking “Taylor rule”) would begin to raise interest rates. But the date T until which the policy rate should be kept at zero is not a date that can be announced with certainty at the time of the shock that causes the zero lower bound to bind; its optimal value depends on how the economy develops. (In the paper, we illustrate numerically how it should depend on the length of time for which the natural rate of interest remains abnormally low; and we give a more general analytical characterization of the optimal policy commitment that implies that T should depend on the evolution of cost-push disturbances as well.)

But this does not mean that mere communication of the forward path that the central bank currently forecasts is all that is likely to be

useful. Unfortunately, such an approach has a serious flaw, which is precisely that a given statement about the change in the anticipated forward path of the policy rate may be subject to multiple interpretations. If an announcement that the date T at which the policy rate will first rise above its lower bound has moved farther into the future is interpreted as meaning that the first date at which a standard (purely forward-looking) Taylor rule would require a policy rate above the floor has moved farther into the future (because of a weakening of the economic outlook)—without in any way challenging the expectation that the bank will, as always, follow such a rule—then the announcement (if also believed) should have a *contractionary* effect on aggregate demand, rather than an expansionary one. For rather than implying that, at a certain point in the future, interest rates will be held lower than one would have expected prior to the announcement (so that real incomes at that time will be greater than would previously have been expected, and likely inflation as well), the announcement would instead imply that real incomes at that time will be lower than would previously have been expected (and likely inflation as well)—which change in anticipations should reduce current willingness to spend rather than increasing it. “Forward guidance” of this kind would have a perverse effect, and be worse than not commenting on the outlook for future interest rates at all.

The only way to avoid this pitfall is to accompany any discussion of the forward path of interest rates with an explanation of the considerations behind it—in particular, of the policy commitments that the anticipated forward path reflects. Discussion of the forward path of interest rates implied by a central bank’s policy commitments may well be useful, for the reasons discussed in Section I.i. But this does not mean that presentation of the implied forward path for interest rates suffices as an explanation of the bank’s policy commitments.

I.iiia Which Criterion for “Liftoff” from the Lower Bound?

In the case of a central bank at the lower bound for its policy rate, it is important to discuss *what will determine* the date T at which “liftoff” from the floor should occur, and not simply the bank’s current estimate (or range of estimates) of that date. Eggertsson and

Woodford (2003) show that, in the context of their New Keynesian DSGE model, an optimal policy commitment can be expressed in terms of a commitment to maintain interest rates at their floor until a particular target is achieved. (After that, the bank should be expected to implement the kind of “flexible inflation targeting” regime characterized in studies of optimal monetary policy that abstract from the existence of a lower bound on interest rates.) The target specifies a path for an “output-gap adjusted” price level, (the log of) which is defined as the log of a general price index plus a positive multiple of the output gap; the coefficient multiplying the output gap in the optimal target criterion depends on the relative weight on output-gap stabilization (as opposed to inflation stabilization) in the bank’s objective. The policy rate should remain at its lower bound as long as even that degree of monetary stimulus results in a gap-adjusted price level below the target path. This means that even once financial conditions have normalized, so that it would be possible for the central bank to achieve both its inflation target and a zero output gap from then onward (at a normal level of the policy rate), it might be necessary to keep interest rates low for somewhat longer, in order to raise the gap-adjusted price level to the target path.

Under a fully optimal policy commitment, the target path for the gap-adjusted price level would not be deterministic (and so able to be fixed at the time of the shock that initially causes the lower bound to bind); instead, the target would be ratcheted up to a steadily higher level, the longer the target shortfalls required by the lower bound persist. Eggertsson and Woodford (2003) provide an explicit formula for the optimal adjustment, and show that it is independent of the nature of the shocks that hit the economy; the adjustment in each period depends only on the degree of shortfall of the actual gap-adjusted price level from the current target level. As a consequence, the existence of such adjustments need not undermine the verifiability of a central bank’s commitment to such a rule. Nonetheless, the adjustment formula would certainly complicate explanation of such a commitment to the public, and the numerical simulations presented by Eggertsson and Woodford suggest that nearly as good a stabilization outcome should be achieved under credible commitment to a much simpler criterion: one in which the target for the gap-adjusted price level

grows at a deterministic rate, given by the bank's long-run inflation target.²⁸ The gains from credible forward guidance depend mainly upon *not letting the target path shift down* in response to persistent target shortfalls during the period of the binding lower bound; the potential gains from actually shifting the path *up* in response to target shortfalls represent a comparatively minor refinement.²⁹

A commitment not to let the target path shift down means that, to the extent that the target path is undershot during the period of a binding lower bound for the policy rate, this automatically justifies anticipation of a (temporarily) more expansionary policy later, which anticipation should reduce the incentives for price cuts and spending cutbacks earlier, and so should tend to limit the degree of the undershooting. Such a commitment also avoids some of the common objections to the simple Krugman (1998) proposal that the central bank target a higher rate of inflation when the zero lower bound constrains policy.

For example, many central bankers are reluctant to consider announcing a higher inflation target on the ground that, while policy is constrained by the lower bound on interest rates, they may have no means by which to hit such a target; and they fear that announcing a target that they do not then achieve will only damage the credibility of any future announcements about their purported policy targets. But the proposal of Eggertsson and Woodford (2003) is not to commit to a target that one pretends can be hit over the next six months, or indeed over any pre-specified horizon; the role of the target is instead to specify how one can tell how far off track one has gotten, whether extremely low interest rates continue to be justified, and whether they are likely to continue to be justified for some time. The facts that the target is specified in terms of the level of a nominal quantity, rather than its rate of change, and that conformity with the target is measured in a backward-looking way (tracking the cumulative departure from the target path) rather than purely prospectively, both make the target meaningful as a commitment without an expectation that the target represents the intended actual outcome over some short horizon. Hence, the announcement of such a target, if properly explained, should create no issues of credibility—assuming,

that is, that the central bank means what it says, and does continue to refer to the target path in its decision process going forward.

Many central bankers also resist the advice to announce an ad hoc modification of their normal inflation target, even when constrained by the zero lower bound, for fear that market participants who see that the inflation target can be shifted on this occasion will thereafter wonder whether it cannot equally be shifted in many other circumstances as well, resulting in a loss of the benefits during normal times of well-anchored medium-run inflation expectations.³⁰ This is an understandable concern. But the “target path” proposal of Eggertsson and Woodford does not represent announcement of a new criterion for appropriate monetary policy that contradicts the way in which the central bank would have wanted people to expect it to act, prior to the disturbance that causes the lower bound to bind; for the proposed target path for the gap-adjusted price level is precisely the one that it would have been desirable for the central bank to commit to maintain, even in the absence of any expectation that the lower bound would prevent achievement of the target for a time. Thus it would have been possible (and according to the model, desirable) for a central bank to commit to an approach to monetary policy, prior to the occurrence of such a crisis, in which the behavior in question after the lower bound is reached would be precisely what the central bank had previously committed to do.

Moreover (and more to the point, under present circumstances), even in the absence of a such a prior commitment, it should be possible for a central bank to argue that a commitment to return to the target path, made explicit only after the lower bound becomes a binding constraint, is, nonetheless, consistent with the policy commitments that had shaped its policy previously—the proviso about what those commitments required in the event of a binding lower bound on interest rates (and consequent persistent undershooting of the target path) had simply never needed to be spelled out until such a situation arose.³¹ The same could not be said of a series of hypothetical future “temporary suspensions” of the inflation target each time unemployment was higher than desired. Hence, adoption of a commitment to reflation during a period at the interest-rate lower

bound, if properly explained, need not raise fears that excuses for continuing high inflation will easily be found.

Despite these conceptual advantages, one must admit that the notion of a target for a “gap-adjusted price level” would not easily be made to seem natural to a public not previously accustomed to discussion of economic policy in such terms. The fact that the “output-gap adjustment” would require reference to a debatable measure of potential output would further increase the possible grounds for suspicion and uncertainty about the policy’s implications. There might then be practical advantages to the formulation of one’s target criterion in terms more familiar to the public, and more easily verifiable, even at the expense of some departure from the theoretically optimal criterion under idealized assumptions about the public’s understanding.

An example of a target criterion that has received considerable attention within the Federal Reserve System is the “7/3 threshold rule” proposed by President Charles Evans of the Chicago Fed (Evans 2011) and analyzed in Campbell et al. (2012). Under this proposal the FOMC would pledge to maintain the funds rate target at its current low level as long as unemployment remains above 7 percent and the expected rate of inflation “over the medium term” remains below 3 percent per year, but would begin to raise the funds rate as soon as either threshold were breached. Adoption of such a commitment by the FOMC would be an important improvement upon current communication policy, in my view. It would emphasize the conditions for exit from the current extremely accommodative policy stance, rather than a date. And the stated conditions would involve both parts of the Fed’s dual legislative mandate, as in the case of the optimal criterion derived by Eggertsson and Woodford (2003). Reference to the unemployment rate might be considered a proxy for the theoretical concept of the output gap, and while it is not necessarily an ideal measure it has the advantages of being much better understood by the public, being widely recognized as a relevant measure of economic performance, and of not being a measure that the Fed itself must construct and might therefore be thought to manipulate.

Nonetheless, the Evans proposal fails to incorporate an important feature of the optimal policy commitment in the model of Eggertsson

and Woodford, which also characterizes optimal commitments in more general (and more realistic) New Keynesian models (for reasons discussed in Woodford 2011): the commitment to compensate subsequently for target misses due to the binding zero lower bound on interest-rate policy. Like a simple Taylor rule, the “7/3 threshold rule” is an example of a *purely forward-looking* criterion for policy: the appropriate policy at any time depends only on the paths for inflation and unemployment that can be achieved from that time onward, independently of the path by which the economy may have reached its current state. In the context of the simple macroeconomic model considered by Eggertsson and Woodford (where inflation and output determination are also purely forward-looking), such a rule will not imply any reason to delay immediately returning to the low-inflation steady state as soon as this is consistent with the zero lower bound on interest rates (i.e., as soon as the natural rate of interest returns to positive territory); it would not imply any commitment to keep the policy rate low for longer than would a strict inflation target or a purely contemporaneous Taylor rule. This means that a credible commitment to such a rule would do nothing to mitigate the problems created by the zero lower bound in the model of Eggertsson and Woodford.

In a model with more complex dynamics, such a commitment is not necessarily irrelevant. But the fact that it is a purely forward-looking criterion, that simply takes a temporarily different form than one that the central bank would be willing to follow under normal conditions, means that it must necessarily appear as an ad hoc departure, both from the policy that the bank had been expected to follow (and that it would have wanted to be expected to follow) prior to the crisis, and from the policy that it will again want the public to expect it to follow later. Adoption of such an arbitrary criterion, presented simply as a temporary suspension of the principles that ordinarily are expected to guide policy, would inevitably tend to reduce the credibility of the bank’s commitment to those principles at other times.

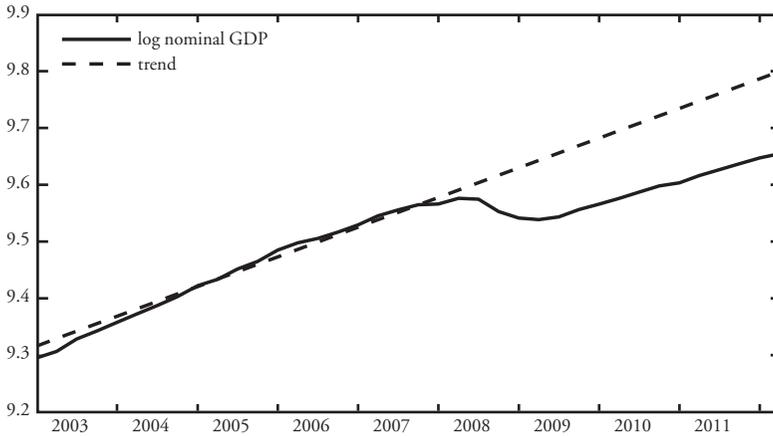
An alternative that I believe should be equally easy to explain to the general public, but that would preserve more of the advantages of the adjusted price-level target path, would be a criterion based on a *nominal GDP target path*, as proposed by Hatzius and Stehn (2011), Romer

(2011) and Sumner (2011) among others. Under this proposal, the FOMC would pledge to maintain the funds rate target at its lower bound as long as nominal GDP remains below a deterministic target path, representing the path that the FOMC would have kept it on (or near) if the interest-rate lower bound had not constrained policy since late 2008. Once nominal GDP again reaches the level of this path, it will be appropriate to raise nominal interest rates, to the level necessary to maintain a steady growth rate of nominal GDP thereafter.

Chart 13 shows the recent evolution of U.S. nominal GDP, with a log-linear trend line fit to the data between the first quarter of 1990 and the third quarter of 2008 (that is, the last quarter before the zero lower bound became a binding constraint).³² Nominal GDP is currently well below this trend line (15.6 percentage points below, as of the second quarter of 2012), and the gap continues to increase. Even if one regarded the average rate of nominal GDP growth over this period as too high to be consistent with a desirable inflation rate over a longer run, and had chosen instead to commit to a trend line with a moderately lower growth rate (say, only 4.5 percent per year), one would still conclude that nominal GDP today is more than 10 percentage points below a trend line extrapolated forward from the third quarter of 2008;³³ such a commitment would accordingly require pursuit of nominal GDP growth well above the intended long-run trend rate for a few years in order to close this gap. At the same time, such a commitment would clearly bound the amount of excess nominal income growth that would be allowed, at a level consistent with the Fed's announced long-run target for inflation.

In the theoretical analysis of Eggertsson and Woodford (2003), a simple nominal GDP target path would not achieve quite the full welfare gains associated with a credible commitment to the gap-adjusted price level target. (In particular, it is surely true—and not just in the special model of Eggertsson and Woodford—that if consensus could be reached about the path of potential output, it would be desirable in principle to adjust the target path for nominal GDP to account for variations over time in the growth of potential.) Nonetheless, such a proposal would retain several of the desirable characteristics of the gap-adjusted price level target that

Chart 13
U.S. Nominal GDP Growth



Note: U.S. nominal GDP growth compared to a log-linear trend line fit to the data between the first quarter of 1990 and the third quarter of 2008.

Source: Federal Reserve Bank of St. Louis.

have been stressed above, and these may well be the most robustly desirable features of that proposal.

Essentially, the nominal GDP target path represents a compromise between the aspiration to choose a target that would achieve an ideal equilibrium if correctly understood and the need to pick a target that can be widely understood and can be implemented in a way that allows for verification of the central bank's pursuit of its alleged target, in the spirit of Milton Friedman's celebrated proposal of a constant growth rate for a monetary aggregate. Indeed, it can be viewed as a modern version of Friedman's "k-percent rule" proposal, in which the variable that Friedman actually cared about stabilizing (the growth rate of nominal income³⁴) replaces the monetary aggregate that he proposed as a better proximate target, on the ground that the Fed had much more direct control over the money supply. On the one hand, the Fed's ability to directly control broad monetary aggregates (the ones more directly related to nominal income in the way that Friedman assumed) can no longer be taken for granted, under current conditions; and on the other hand, modern methods of forecast targeting make a commitment to the pursuit of a target defined in terms of variables that are not under the short-run control of the central bank more credible. Under these circumstances, a case

can be made that a nominal GDP target path would remain true to Friedman's fundamental concerns.³⁵

Would there be any role for releases of central-bank projections of the economy's likely future path under such an approach to forward guidance? Yes. I have argued that merely releasing projections (including projections of the expected path of the policy rate), without clarifying the target criterion that is assumed to shape future policy deliberations—or referring only to purely forward-looking criteria, that do not incorporate the kind of commitment to correction of target misses that a nominal GDP target path would imply—accomplishes little. But the success of the kind of forward guidance proposed here, based on commitment to a history-dependent target criterion as the basis for future policy, depends both on people's being able to *understand now* the future consequences of such a commitment, and on the *credibility* of the purported commitment on the part of the central bank. Both would be greatly facilitated by the adoption of a transparent forecast-targeting procedure as the basis for monetary policy deliberations and for communication with the public about the outcome of those deliberations.

II. Expanding the Supply of Bank Reserves

Much of the discussion of the possibilities for expansionary policy since the major central banks reached their interest-rate lower bounds around the end of 2008 has focused on the use of changes in the central bank's balance sheet—its overall size, the composition of its assets, and the share of its liabilities that are “monetary” in character—as additional dimensions of policy, apart from the central bank's influence over overnight money-market rates. Sometimes these additional dimensions of policy are proposed as measures that can usefully support a central bank's forward guidance, by allowing it to take *concrete actions* that may be viewed as underlining its intentions with regard to future interest-rate policy, instead of relying purely upon speech.

For others, they represent additional tools that should be used alongside forward guidance, given that, even under ideal circumstances (when the central bank's representations about future policy

are fully believed and their consequences fully understood), forward guidance alone cannot be expected to *completely eliminate* the distortions resulting from the interest-rate lower bound. To the extent that other means are regarded as having similar effects as a current reduction in short-term interest rates, it should still be useful to use such tools.

For yet others, they represent tools that should make forward guidance unnecessary. After all, in the model of Eggertsson and Woodford (2003), if a sufficient reduction in the current short-term interest rate were possible at all times, there would be no need for any commitment to history-dependent policy at all; so if balance-sheet policies can achieve the same outcome as an interest-rate cut without violating the lower bound, they should eliminate the need for any awkward public statements about future policy.

But what can one reasonably expect adjustments of the central bank's balance sheet to achieve? Can such policies properly be viewed as the equivalent of interest-rate reductions, achieved by other means? And if they are effective, do they therefore eliminate the need for explicit forward guidance?

In this section, I consider one particular type of balance-sheet policies, namely those aimed at expanding the monetary liabilities of the central bank (the monetary base), in practice by expanding the supply of bank reserves, either through extensions of central-bank credit or asset purchases of one kind or another. The implications of central bank purchases of particular types of assets (regardless of how this may be financed), and of changes in the composition (as opposed to the size) of the asset side of the balance sheet are deferred until Section III.

II.i Pure Quantitative Easing in Theory

The best-known doctrine according to which balance-sheet policies should remain effective at the interest-rate lower bound is the theory of "quantitative easing," put into practice in 2001-06 by the Bank of Japan (BOJ) (which originated the term). (Here I call this "*pure* quantitative easing," to distinguish the original theory of quantitative easing

from the way the term has come to be used in the press, to refer to a much broader class of policies and possible mechanisms.³⁶)

According to this theory, increases in the monetary base—which continue to be possible, and are completely under the control of the central bank, regardless of the level of overnight interest rates—should stimulate increased aggregate nominal expenditure, regardless of whether overnight interest rates change. The theory was urged upon the BOJ by outside critics after it had reduced its target for the call rate (the overnight interest rate which had been the BOJ’s policy rate) essentially to zero, but prices continued to fall while economic activity remained sluggish. For example, at a conference in 2000, Milton Friedman responded (as quoted by Beckworth 2011) to the suggestion that the possibilities for further monetary expansion had been exhausted:

Now, the Bank of Japan’s argument is, “Oh well, we’ve got the interest rate down to zero; what more can we do?” It’s very simple. They can buy long-term government securities, and they can keep buying them and providing high-powered money until the high-powered money starts getting the economy in an expansion.

Thus, Friedman argued that a further increase in the monetary base [which he calls “high-powered money”] would necessarily increase spending, and so end the deflationary slump.³⁷

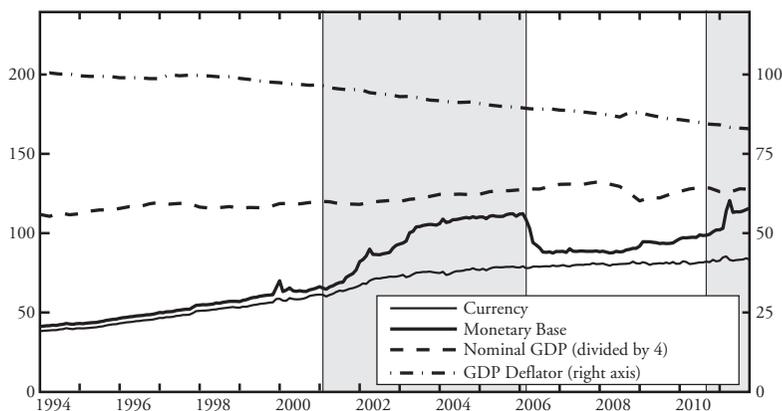
According to classic monetarist doctrine, what matters is the expansion of the central bank’s monetary liabilities, and not the nature of the assets acquired with the newly created base money; hence, it is most prudent (and involves the least unnecessary interference with market mechanisms) if the central bank restricts its purchases to safe government securities. What matters is the deliberate expansion of the monetary base to the extent necessary to support the desired level of aggregate nominal expenditure. The BOJ followed this advice in its policy of “quantitative easing,” under which it announced a series of progressively higher numerical targets for the current account balance (i.e., the supply of bank reserves), beginning in

March 2001, and undertook asset purchases (mainly of Japanese government bonds) so as to implement these targets. The resulting large increase in current account balances, and the consequent increase in the Japanese monetary base, are shown in Chart 14.

But while it is undeniable that a central bank can still further increase the monetary base after the supply of bank reserves is already great enough to drive short-term interest rates down to the level of the rate of interest paid on reserves (or to zero, if there is no interest on reserves), it is hardly obvious that this should have any effect on aggregate nominal expenditure. The doctrine that Friedman assumes in his advice to have been well-established empirically is one according to which an increase in the monetary base necessarily increases a broader monetary aggregate (M1 or M2), in proportion to the increase in base money if the “money multiplier” remains stable; and the increase in broad money necessarily increases aggregate nominal expenditure, in proportion to the increase in money if the “velocity of money” remains stable. The economic mechanism behind the causal chain is one according to which there should be a finite demand for real base money, proportional to the real volume of transactions in the economy, and a decreasing function of the opportunity cost of holding base money.³⁸ If the nominal size of the monetary base increases, one or more of these determinants of desired holdings must also change, to maintain equilibrium: a decline in the interest differential between short-term nonmonetary assets and bank reserves (to reduce the opportunity cost of holding reserves, and so increase demand for reserves); an increase in the real volume of transactions (which should proportionally increase the demand for real balances); or an increase in the price level (so that the real monetary base does not increase by as large a proportion as the nominal increase).

These familiar mechanisms may have resulted in a fairly reliable connection between expansions of the monetary base and increases in aggregate nominal expenditure under ordinary circumstances—under which a substantial opportunity cost of holding excess reserves exists—but there is no reason to expect them to work in the same way once the opportunity cost is eliminated, because money-market interest rates are no longer higher than the interest rate paid on

Chart 14
Evolution of the Monetary Base in Japan



Notes: The region between the curve labeled “monetary base” and the one labeled “currency” indicates the quantity of current account balances. The first grey region indicates the period of the BOJ’s policy of “quantitative easing”; the second grey region indicates the period of its policy of “comprehensive monetary easing.” Units (left scale): trillions of yen.

Sources: Bank of Japan and Federal Reserve Bank of St. Louis.

reserves. It should not be possible for the rate of interest at which banks are willing to hold short-term nonmonetary instruments to be no larger than the interest rate on reserves, unless reserves have *ceased to be scarce*, so that they no longer earn a “liquidity premium” owing to their special role in the payments system. And under the latter circumstance, the demand for reserves should become *infinitely elastic*, so that variations in the precise quantity of excess reserves (as opposed to other short-term, essentially riskless assets) that banks must hold will have no consequences for equilibrium determination. Indeed, it is this very fact that the demand for reserves becomes infinitely elastic at an opportunity cost of zero that explains why there is a lower bound on how low money-market interest rates can be driven through variation in the supply of reserves. The *same* mechanism implies that once that lower bound is reached, further expansion of the supply of reserves should *not* have any consequences for aggregate expenditure or the general level of prices (or for that matter, for broad monetary aggregates).³⁹

Monetarist authors often assert that the Keynesian notion of a “liquidity trap” depends on an overly narrow conception of the

monetary transmission mechanism, in which it is assumed that monetary policy can *only* influence spending through its effects on a short-term nominal interest rate, so that if that rate can no longer fall (owing to the zero lower bound) monetary policy must be impotent. To this monetarists reply that there are many other asset prices (including longer-term interest rates) that can still move in ways that would provide incentives for increased expenditure, even when the short-term nominal interest rate remains fixed; and that an excess supply of money will surely affect these other rates even if it can no longer decrease the short-term nominal interest rate. But this sort of reasoning treats the zero lower bound as if it were some additional institutional constraint (like the practice of not auctioning Treasury bills in a way that would allow the yield to be negative), and not a *consequence of the demand for bank reserves*. Because the demand for reserves becomes unbounded once the opportunity cost falls to zero, there is no “disequilibrium” that requires some other asset price to adjust in order to restore the balance of supply and demand for reserves. There is accordingly no reason for the equilibrium values of other asset prices to change, whether or not these other asset prices are also arguments of the demand for base money, and whether or not they are important determinants of aggregate expenditure.

It may be objected that the argument just given establishes only that once the interest-rate lower bound is reached, bank reserves and other very short-term riskless claims should become essentially perfect substitutes, so that increases in reserves that come about through central-bank purchases of riskless short-term assets—the sort of transactions that represented the preferred means of increasing or decreasing the supply of reserves under the “bills only” doctrine of the Federal Reserve System (Luckett 1960)—should have no effect. But this does not imply that the creation of reserves in order to purchase longer-term (or risky) assets should have no effect, and indeed it may be noticed that in the quotation above, Friedman refers to purchases of “*long-term* government securities.” Might QE not be effective as long as long-term government bonds are purchased?

The argument given above does not suffice to answer this question (which is considered further in the next section). However, I would

maintain that such a policy does not constitute “quantitative easing” in the pure sense. A policy of creating additional reserves in order to purchase long-term Treasury securities can equivalently be analyzed as the composition of two policies: one under which new reserves are created by purchasing short-term Treasuries (“pure QE”), and another under which the central bank sells short-term Treasuries to buy long-term Treasury securities, with no change in the size of its balance sheet (Operation Twist). To the extent that a reserve-financed purchase of long-term Treasuries has effects at the zero lower bound, on the price of Treasuries or anything else, these effects should be identical to the effects of Operation Twist alone.

It does not make sense, in my view, to consider such a policy a variety of “quantitative easing.” To the extent that it is effective, the mechanisms involved must be quite distinct from those invoked in the classic theory of quantitative easing, and have nothing to do with the supply of or demand for monetary liabilities of the central bank. (They instead involve only the asset side of the central bank’s balance sheet.) And, while such policies are among the actions that can be taken by a central bank, it is not obvious that they are properly considered as part of monetary policy. They belong to the sphere of *debt management*, and to be effective, should at the very least be coordinated with the Treasury, since it is the net change in the supply of securities of various types in the hands of the public that should matter, whether brought about through transactions by the Treasury or by the central bank. Some monetarists, such as Congdon (2011), argue that it would be best for such operations to be conducted *entirely* by the Treasury, in order to avoid the risk of accounting losses by the central bank. While I do not exclude such policies from consideration here, it makes sense to treat them as analytically distinct from pure QE policies.

Finally, some would argue that the well-established principle of the long-run neutrality of money implies that an expansion of the monetary base must eventually result in a proportional increase in the general level of prices (though with no real effects, in that long run), and, hence, must eventually be able to increase aggregate nominal expenditure, to whatever extent may be desired. But this argument concerns

the effects of a policy of *permanently* increasing the monetary base. The “irrelevance results” of Krugman (1998) and Eggertsson and Woodford (2003) instead pertain to a policy that increases the monetary base only during a period over which the zero lower bound prevents the central bank from achieving its usual targets, while the central bank is expected to return to its usual (purely forward-looking) approach to policy once it is no longer constrained—for example, by returning to the pursuit of its long-run inflation target, or by following a Taylor rule consistent with such a target. Policies of the kind assumed in these thought experiments do not imply that the increase in base money resulting from the QE policy will be permanent; eventually the lower bound ceases to be a binding constraint, and after that, the path of the monetary base is the one required by the inflation target or by the Taylor rule, independent of the amount of quantitative easing that has occurred previously.

If, instead, one were to assume a permanent increase in the size of the monetary base, and assume that it is immediately understood by everyone in the economy that such a permanent change in policy has occurred, then such a policy would be predicted to have an immediate positive effect on economic activity during the period in which the lower bound binds, in either the model of Krugman (1998) or Eggertsson and Woodford (2003). Indeed, this explains the apparently different result of Auerbach and Obstfeld (2005). These authors assume an initial policy defined by a deterministic path for the monetary base, in a situation in which a disturbance causes short-term nominal interest rates under this policy to fall to zero for a finite period of time, and then consider an alternative policy under which a permanently higher path for the monetary base is chosen; they find that even though the nominal interest rate is zero at the time of the increase in the monetary base, the change immediately increases both real activity and prices. The irrelevance result of Eggertsson and Woodford actually applies equally to the model of Auerbach and Obstfeld; the policy considered by Auerbach and Obstfeld is effective because in addition to the immediate increase in the monetary base, it *also* involves a commitment to a different policy *after* the zero lower bound ceases to bind—at which time it implies not just a

different path for the monetary base, but different *interest-rate policy* as well.

In fact, the effects of the policy proposed by Auerbach and Obstfeld follow *entirely* from the commitment to a different policy later; the effects would be the same if there were no open-market purchases of assets at all, until the end of the period in which the short-term nominal interest rate is zero under the original policy (i.e., the period in which there is already an excess supply of base money, relative to what is needed to keep the interest rate at zero). Moreover, these expansionary effects during the period of the “liquidity trap” depend entirely on the change in future policy already being able to be foreseen. The model is thus another example of the gains that are possible, in principle, through a commitment to more expansionary policy in the future.

Indeed, it is an example closely related to the proposal of commitment to a nominal GDP target path, discussed above. The model of Auerbach and Obstfeld assumes a cash-in-advance constraint, as a result of which the demand for base money is equal to aggregate nominal expenditure each period, except when the nominal interest rate falls to zero (in which case households are willing to hold any quantity of base money at least equal to their planned nominal spending). Hence, the commitment to a deterministic path for the monetary base is equivalent to a commitment to a deterministic target path for aggregate nominal expenditure, together with a commitment to use monetary policy to keep nominal expenditure equal to the target at all times, unless expenditure undershoots the target even when the money supply is already large enough to drive the nominal interest rate to zero. The demonstration by Auerbach and Obstfeld that welfare can be increased by permanently increasing the supply of base money could alternatively be used to show that welfare could be increased by committing to keep the nominal interest rate at zero until it is possible to hit a certain deterministic target path for nominal GDP, and then use monetary policy to keep nominal GDP growing at a steady rate thereafter. The inferior initial equilibrium is instead one in which nominal GDP is allowed to follow a permanently lower path, albeit with the same long-run growth rate.

Might one, nonetheless, conclude from the Auerbach and Obstfeld example that money-financed asset purchases (as they describe the policy that they analyze) can be an effective substitute for forward guidance? No, for the effects that they analyze occur only under the assumption that the implications for the long-run path of base money are immediately apparent to everyone as a result of the policy change. One can only maintain that such effects would occur even in the absence of explicit forward guidance—an announcement of the central bank's commitment to a different kind of future monetary policy than had previously been expected—if one supposes that the mere fact of the current expansion of the monetary base would give people a reason to anticipate a correspondingly higher long-run monetary base, without a need for any accompanying explanation.

And there would be no reason to expect that. Even if, in the past (in the absence of a binding lower bound constraint), fluctuations in the monetary base have been extremely persistent, there would be no reason for people to expect that an extraordinary increase in the supply of bank reserves as part of a “quantitative easing” policy would have the same kind of consequences for the long-run monetary base as past variations in the supply of reserves. There is certainly nothing about such an increase in the supply of bank reserves that is, for mechanical reasons, difficult to reverse. Indeed, Japan's experience with quantitative easing illustrates this.

II.ii Quantitative Easing in Practice

As noted above, the term “quantitative easing” was introduced by the BOJ in March 2001, to describe a new policy that replaced its previous operating target for an overnight interest rate (the call rate, which had been near zero much of the time for a few years at that point) by quantity targets for the supply of bank reserves (current account balances). While the complete set of policy measures undertaken was (not surprisingly) fairly complex, the basic thrust of the policy was fairly close to providing an illustration of the kind of policy to which the irrelevance results of Krugman (1998) and Eggertsson and Woodford (2003) should apply.⁴⁰

The explicit intention of the policy was to increase the supply of bank reserves beyond the level required to keep the policy rate near zero, and policy announcements focused on the BOJ's quantity targets for current account balances. Importantly, there was no commitment to maintain the increased supply of reserves permanently, and indeed they were promptly withdrawn once the justification for a special regime was considered to have passed. As Chart 14 shows, most of the increase in current account balances was reversed, in the space of a few months, after the policy was suspended in March 2006, as a result of CPI inflation that had been measured to be slightly above zero. The Japanese monetary base resumed a path that was close to a continuation of its trend prior to the QE period; hence, market participants who had continued to hold expectations about the long-run Japanese monetary base that were unchanged as a result of the QE policy would not have been far off in their prediction.

And as the theoretical models would predict—but contrary to the quantity-theoretic reasoning that had provided the basis for the policy proposal—there was little effect of the policy on aggregate nominal expenditure. As shown in Chart 14, the increased supply of bank reserves raised the total monetary base by 60 percent over the first two years of the policy, and eventually by nearly 75 percent. Yet there was no corresponding increase in aggregate nominal expenditure: nominal GDP was only 6 percent higher after five years of QE than it had been in the first quarter of 2001, despite a massive increase in the monetary base. And as the chart also shows, deflation (here measured by the GDP deflator) continued unabated.

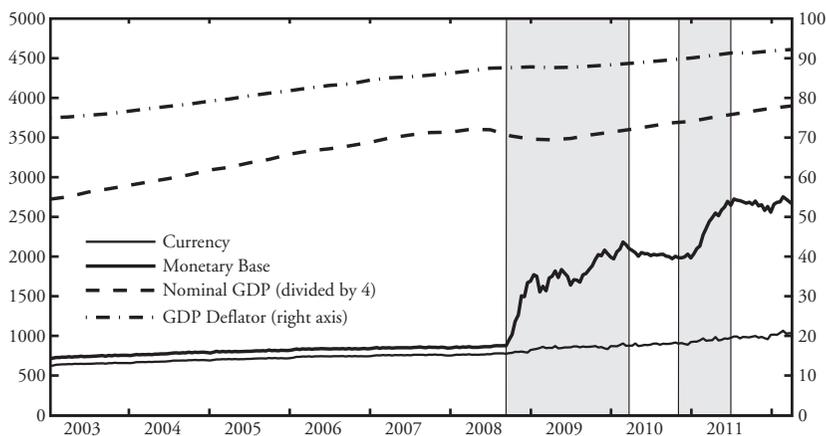
A number of studies of the effects of this experiment attribute some reduction in longer-term bond yields to it, though without measurable consequences for aggregate demand or the rate of deflation. But it is not obvious that even these asset-price effects should be attributed to the “pure quantitative easing” aspect of the BOJ's policy. The introduction of the new policy in March 2001 was also accompanied by a new and stronger form of forward guidance: the BOJ committed to maintain reserve balances large enough to keep the overnight interest rate near zero until CPI inflation became zero or higher on a sustained basis.⁴¹

A number of scholars conclude that this new commitment as to the duration of the zero-interest-rate policy had substantial effects on market expectations regarding future short-term interest rates, and through this expectation channel on longer-term interest rates (Okina and Shiratsuka 2004, Oda and Ueda 2007, Ugai, 2007). Some studies also conclude that the composition of the BOJ's asset purchases affected yields on the particular type of assets purchased; in particular, the term premia on longer-term Japanese government bonds may have been reduced by the BOJ's purchases. But few authors find any evidence of effects of variation in the BOJ's quantity targets for current account balances, and, hence, for effects that can be attributed to pure quantitative easing (Ueda 2012a, 2012b).

The BOJ itself appears no longer to put great stock in pure QE as a policy. Its "comprehensive monetary easing" policy, introduced in October 2010 in response to continuing concerns about the economic outlook, has again resulted in a significant increase in the size of the BOJ's balance sheet, but does not involve quantitative targets for current account balances. Instead, its main elements are a "virtually zero" target for the call rate, and forward guidance about the conditions that will determine how long the low target will be maintained; targeted purchases of a variety of specific types of assets; and a credit facility that supplies funds for relatively long terms at a fixed rate, against specific types of collateral, as a relatively direct way of lowering longer-term interest rates (Lam 2011). These latter programs appear to have had some effects on asset prices, but as Chart 14 shows, there is little evidence of any immediate effect of the associated expansion of the monetary base on aggregate nominal expenditure, of the kind posited by quantity theorists.

The even more massive increases in the monetary base by the Federal Reserve's unconventional policies since fall 2008 have similarly had little evident effect on aggregate nominal expenditure, as shown in Chart 15. The grey regions in this figure mark two periods over which the FOMC has substantially increased the size of the Fed's balance sheet.⁴² While neither of these policy changes was officially described as a program of pure quantitative easing—and Chairman Bernanke, in a speech in January 2009, explicitly denied that the

Chart 15
Evolution of the Monetary Base in the U.S.



Notes: The region between the curve labeled “monetary base” and the one labeled “currency” indicates the quantity of current account balances. The gray regions identify two periods of unusual balance-sheet expansion discussed further in the text. Units (left scale): billions of dollars.

Source: Federal Reserve Bank of St. Louis.

Fed’s first planned program of asset purchases constituted “quantitative easing” (Bernanke 2009)—proponents of that theory might suppose that the real effect of these programs resulted from the way that they increased the supply of base money, regardless of its official justification. But there has been little sign of the effects that that theory would predict. As shown in the chart, the U.S. monetary base has more than tripled since September 2008, but the growth in nominal GDP from the fourth quarter of 2008 to the second quarter of 2012 has been less than 11 percent. There has similarly been as yet little sign of any acceleration of inflation, despite the warnings of some monetarists. Thus such effects as the programs have had (which are discussed further below) do not seem to support the theory of pure quantitative easing.

III. Targeted Asset Purchases

I turn now to the question of the way in which it should be expected to matter what the central bank holds on the asset side of its balance sheet. As noted above, the argument already given for the irrelevance of purchases of short-term Treasury bills by creating additional bank reserves also implies that there should be no difference

between purchases of longer-term assets that are financed by creating additional bank reserves and purchases of longer-term assets that are financed by selling short-term Treasury bills previously held by the central bank (as under the Fed's maturity extension program). But what effect should one expect there to be of an increase in central-bank purchases of longer-term assets, under either of those scenarios? Once again, I begin with a brief review of the theoretical literature and then turn to the lessons that may be gleaned from recent experience, focusing on the several recent asset-purchase programs of the Federal Reserve.

III.i Effects of Targeted Asset Purchases in Theory

It is often supposed that open-market purchases of securities by the central bank must inevitably affect the market prices of those securities (and hence other prices and quantities as well), through what is called a "portfolio-balance effect": if the central bank holds less of certain assets and more of others, then the private sector is forced (as a requirement for equilibrium) to hold more of the former and less of the latter, and a change in the relative prices of the assets will almost always be required to induce the private parties to change the portfolios that they prefer. In order for such an effect to exist, it is thought to suffice that private parties not be perfectly indifferent between the two types of assets; and there are all sorts of reasons why differences in the risky payoffs associated with different assets should make them not perfect substitutes, even in a world with frictionless financial markets. Thus, while many authors would agree that central-bank exchanges of very short-maturity Treasury bills for overnight balances at the central bank should have little consequence, once the interest rate on the Treasury bills has fallen to essentially the level of the interest rate paid on reserves—on the ground that in this case, the instruments being exchanged are close to being perfect substitutes—they assume that this should not be equally true of central-bank purchases of other types of assets, including longer-maturity Treasury securities.

But it is important to note that such "portfolio-balance effects" do not exist in a modern, general-equilibrium theory of asset prices—in which assets are assumed to be valued for their state-contingent

payoffs in different states of the world, and investors are assumed to correctly anticipate the consequences of their portfolio choices for their wealth in different future states—at least to the extent that financial markets are modeled as frictionless. It is clearly inconsistent with a representative-household asset pricing theory (even though the argument sketched above, and many classic expositions of portfolio-balance theory, make no reference to any heterogeneity on the part of private investors). In the representative-household theory, the market price of any asset should be determined by the present value of the random returns to which it is a claim, where the present value is calculated using an asset pricing kernel (stochastic discount factor) derived from the representative household's marginal utility of income in different future states of the world. Insofar as a mere reshuffling of assets between the central bank and the private sector should not change the real quantity of resources available for consumption in each state of the world, the representative household's marginal utility of income in different states of the world should not change. Hence, the pricing kernel should not change, and the market price of one unit of a given asset should not change, either, assuming that the risky returns to which the asset represents a claim have not changed.

How does 1950s-vintage “portfolio-balance” theory obtain a different result, even when the private sector is represented by a representative mean-variance investor? It assumes that if the private sector is forced to hold a portfolio that includes more exposure to a particular risk—say, a low return in the event of a real-estate crash—then private investors' willingness to hold that particular risk will be reduced: investors will anticipate a higher marginal utility of income in the state in which the real-estate crash occurs, and so will pay less than before for securities that have especially low returns in that state. But the fact that the central bank takes the real-estate risk onto its own balance sheet, and allows the representative household to hold only securities that pay as much in the event of a crash as in other states, does not make the risk disappear from the economy. The central bank's earnings on its portfolio will be lower in the crash state as a result of the asset exchange, and this will mean lower earnings distributed to the Treasury, which will in turn mean that higher taxes

will have to be collected by the government from the private sector in that state; so the representative household's after-tax income will be just as dependent on the real-estate risk as before. This is why the asset pricing kernel in a modern representative-household asset-pricing model does not change, and why asset prices are unaffected by the open-market operation.⁴³

The irrelevance result is easiest to derive in the context of a representative-household model, but in fact it does not depend on the existence of a representative household, nor upon the existence of a complete set of financial markets. All that one needs for the argument are the assumptions that (i) the assets in question *are valued only for their pecuniary returns*—they may not be perfect substitutes from the standpoint of investors, owing to different risk characteristics, but not for any other reason—and that (ii) *all investors can purchase arbitrary quantities of the same assets at the same (market) prices*, with no binding constraints on the positions that any investor can take, other than her overall budget constraint.

Under these assumptions, the irrelevance of central-bank open-market operations is essentially a Modigliani-Miller result, as noted by Wallace (1981). If the central bank buys more of asset x by selling shares of asset y , private investors should wish purchase more of asset y and divest themselves of asset x , by exactly the amounts that undo the effects of the central bank's trades. The reason that they optimally choose to do this is in order to hedge the additional tax/transfer income risk that they take on as a result of the change in the central bank's portfolio. If share θ_b of the returns on the central bank's portfolio are distributed to household b , where the $\{\theta_b\}$ are a set of weights that sum to 1, then household b should choose a trade that cancels exactly fraction θ_b of the central bank's trade, in order to afford exactly the same state-contingent consumption stream as before. Summing over all households, the private sector chooses trades that in aggregate precisely cancel the central bank's trade. The result obtains even if different households have very different attitudes toward risk, different time profiles of income, different types of nontradable income risk that they need to hedge, and so on, and regardless of how large or small the set of marketed securities may be. One can easily

introduce heterogeneity of the kind that is often invoked as an explanation of time-varying risk premia without this implying that any “portfolio-balance” effects of central-bank transactions should exist.

Many readers of Wallace (1981) are likely to have found his result paradoxical, and doubted the practical relevance of the entire line of reasoning. For the result seemed to imply not only that exchanges of Treasuries for mortgage-backed securities (MBS) by the Federal Reserve, holding fixed the overall size of the Fed’s balance sheet, should have no effect, but also that increases in the supply of bank reserves as a result of open-market purchases of Treasuries should have no effect. Yet the latter kind of open-market operation had long been routinely used by the Fed to bring about desired changes in the federal funds rate, as every undergraduate learns. The theory seemed patently inapplicable to the operations of actual central banks in actual market economies.

Moreover, it is clear that overnight balances at the Fed have often been held despite being dominated in rate of return; until October 2008, these balances earned a zero nominal return, while other overnight interest rates (such as the federal funds rate) were invariably higher, for reasons that cannot be attributed purely to default risk. A natural (and thoroughly conventional) inference is that this particular asset is (or at least, has often been) held for reasons beyond its pecuniary return alone; we may suppose that reserves at the Fed (and base money more generally) supply transactions services, by relaxing constraints that would otherwise restrict the transactions in which the holders of the asset can engage. The existence of these nonpecuniary returns—which may be modeled using any of a variety of familiar devices—will invalidate the Wallace (1981) neutrality result, at least insofar as open-market purchases of securities that increase the supply of reserves are concerned.

But one can introduce a transactions role for reserves, or for monetary liabilities of the central bank more generally, however, while still entertaining the hypothesis that with regard to *all assets other than monetary liabilities of the central bank*, the two postulates still hold: assets other than “money” are valued only for their pecuniary returns, and all investors can purchase arbitrary quantities of any of these assets at the same (market) prices. In this case, a weaker version of the irrelevance

result for central-bank trades still applies. No open-market operation that changes the composition of the central bank's asset portfolio, while keeping unchanged the outstanding volume of the monetary liabilities of the central bank, should have any effects on asset prices, goods prices, or the allocation of resources.⁴⁴ Again, the argument is essentially a Modigliani-Miller theorem, and holds despite an arbitrary degree of heterogeneity in the situations of different households, and regardless of the size of the set of traded securities.

The result in this case validates the classic monetarist position: the supply of monetary liabilities by the central bank matters for macroeconomic equilibrium, but it does not matter at all what kinds of assets might “back” those liabilities on the other side of the central bank's balance sheet, or how the base money gets to be in circulation. Hence, a generation or two of texts in monetary economics have found it convenient to analyze monetary policy using models in which there is no central-bank balance sheet—merely a government printing press which creates additional “money” at a greater or lesser rate, which is then put in the hands of private parties, perhaps by dropping it from helicopters. Again, the omission is completely justifiable, if financial markets function efficiently enough for the two postulates mentioned above to hold, except for the qualification regarding the special properties of “money.”

Under this view, there would be still be no ground for viewing targeted asset purchases as a relevant dimension of central-bank policy, though variations in the supply of monetary central-bank liabilities would matter. This might seem to support the doctrine of pure quantitative easing discussed in the previous section: the aim of policy should simply be to achieve a sufficient expansion in the monetary base, with the particular assets purchased being irrelevant to the consequences that should be expected for aggregate demand. But the kind of model just sketched would provide no support for a policy of pure quantitative easing, either. Expansion of the supply of bank reserves stimulates aggregate demand under normal circumstances, as it ordinarily implies a reduction in the short-term riskless rate of interest; yet once the supply of reserves is sufficient to drive the short-term riskless rate to zero (or to the rate of interest

paid on reserves), there is no reason to expect further increases in the supply of reserves to increase aggregate demand any further, as argued in the previous section. Once banks are no longer foregoing any otherwise available pecuniary return in order to hold reserves, there is no reason to believe that reserves continue to supply any liquidity services at the margin; and if they do not, the Modigliani-Miller reasoning applies once again to open market operations that increase the supply of reserves, just as in the model of Wallace.

It is possible, of course, that assets other than just the monetary liabilities of the central bank may be valued for their role in facilitating transactions, and not merely for their state-contingent pecuniary returns. For example, Krishnamurthy and Vissing-Jorgensen (2012a) present evidence that variations in the supply of U.S. Treasury debt influence market yields on Treasury securities, and interpret this as evidence of a “safety premium” associated with uses in the financial sector for certain especially safe assets such as Treasuries (for example, as collateral in repo transactions), that increases the market value of these securities beyond what would follow from their state-contingent pecuniary returns alone. In the presence of such a premium, even the generalized irrelevance result stated above would not necessarily be valid, in the case of central-bank transactions involving sales or purchases of long-term Treasuries.

One’s conclusions depend on how the “safety premium” is modeled. If one supposes that there is one kind of convenience yield, supplied by both Fed liabilities (base money) and Treasury debt to differing extents, then while the existence of a convenience yield in connection with Treasury securities would require modification of the restricted irrelevance proposition stated above (it would not, in general, be enough that the central bank’s trades leave fixed the supply of base money), it would not imply any new opportunities for effective open-market operations at the zero lower bound. For once there is a sufficient supply of bank reserves for the shadow value of the convenience yield to be zero (as shown by the disappearance of a differential between the overnight rate and the interest rate on reserves), the contribution of the safety premium to the market value of

Treasuries should also be zero, and continue to be regardless of modest variations in the quantity of safe assets in the hands of the public.

Alternatively, one might suppose that Treasuries supply a convenience yield of a *different sort* than is provided by bank reserves, so that the fact that the liquidity premium for bank reserves has fallen to zero would not necessarily imply that there could not still be a positive safety premium for Treasuries. This possibility is suggested by the findings of Krishnamurthy and Vissing-Jorgensen (2011) with regard to the effects of Fed asset purchases even with the federal funds rate at its lower bound, discussed in the next section. If so, it should still be possible for central-bank purchases or sales of long-term Treasury securities to affect the size of the safety premium, and so to affect long-term yields even in the absence of any change in the expected path of short rates.

But even in this case, such a model would not provide a justification for central-bank purchases of long-term Treasury securities as a way of stimulating the economy when constrained by the interest-rate lower bound. First of all, even though purchases of long-term Treasuries could raise the price of (and so lower the yield on) Treasuries, this would not necessarily imply any reduction in *other* long-term interest rates, since the increase in the price of Treasuries would reflect an increase in the safety premium, and not necessarily any increase in their price *apart from* the safety premium (and, hence, not necessarily any reduction in the discount rate that the market uses to value future payments). This means that while the U.S. Treasury would then be able to finance itself more cheaply at the margin, there would not necessarily be any such benefit for private borrowers, and hence any stimulus to aggregate expenditure.

Moreover, as Krishnamurthy and Vissing-Jorgensen (2012b) note, an increase in the safety premium obtained by making “safe assets” (in the relevant sense) more scarce would in itself be *welfare-reducing*. If Treasuries provide a convenience yield not available from other assets (including bank reserves), then reducing the quantity of Treasuries in the hands of the public reduces the benefits obtained from this service flow; and supplying more reserves instead would not replace the convenience yield from Treasuries with a different convenience

yield of equal value, since (under the hypothesis of short-term interest rates at the lower bound) additional reserves would supply a marginal convenience yield of zero (given that reserves are already supplied beyond the satiation level). This points to an important lesson of greater generality: it is not reasonable to consider policy successful simply because it can raise the price of some asset by whatever means. It matters which asset prices are affected, and it matters which distortions are created in order to affect the equilibrium valuations of those assets.

Another reason for the irrelevance result stated above not to hold in practice can be the existence of binding constraints on participation in particular markets or on the positions that particular traders can take in those markets.⁴⁵ Suppose, for example, that certain assets are only purchased (or only purchased at low cost) by “specialists,” who have an advantage over other investors in this particular activity. The Modigliani-Miller theorem no longer holds if the central bank purchases assets of this kind, if not all of the addition to the central bank’s state-contingent portfolio earnings is distributed to the “specialists”; for other investors will not reduce their holdings of the assets in question, even if the change in their state-contingent tax liabilities gives them a hedging motive for doing so, owing to their nonparticipation in that market. Moreover, under this mechanism, the central bank’s purchases have a disproportionate effect on demand for the particular assets that it purchases, and the hypothesis of market segmentation makes it particularly likely that a substantial change in the price of those particular assets will be required for market clearing.

Market segmentation of this kind is particularly likely to be important (and potentially a source of significant inefficiencies) during times of financial turmoil like that following the failure of Lehman Brothers, and the kind of theory just sketched provides a plausible account of the relevance of various special credit programs of the Fed in that period. For example, Woodford (2011b) discusses how the effects of the Fed’s Commercial Paper Funding Facility (CPFF), introduced in October 2008, on spreads between commercial-paper yields and other money-market rates might be understood in these

terms.⁴⁶ In this application, the sharp increase in these spreads after the run on money market mutual funds provides a plausible case for segmentation of this market, at least at that point in time, so it is not surprising that the introduction of the Fed's facility would make a considerable difference. Such an example shows that at least under some circumstances, central-bank purchases can be expected to influence asset prices. But two caveats are necessary about the generality of the lesson that should be drawn from it.

First, the example relates to a particularly severe moment of market disruption. The fact of market segmentation at such a time does not imply that it is equally significant under more routine conditions. By February 2010 the Fed had closed the CPFF, along with most of the other special credit facilities created in response to the financial crisis, on the ground that markets could again function relatively well on their own; to the extent that that judgment was correct, one should not expect that targeted purchases (or provision of financing for such a specific category of assets) would have similar effects today.

Second, the existence of market segmentation makes it possible for central-bank purchases to affect the price of an asset, but at the same time limits the generality of the effects of a change in that particular asset price on the rest of the economy. In order for the policy to be judged effective, it is necessary that influencing that particular asset price can be expected to achieve an important aim. In the case of the CPFF, this presumably was the case—only the financing costs of a particular narrow class of borrowers were affected, rather than financial conditions more generally, but the program achieved a specific goal that motivated its creation. One cannot, however, point to such a program as evidence that purchasing any kind of assets eases financial conditions generally. Instead, to the extent that market segmentation is relied upon as the basis for a policy's effectiveness, one should expect the effects to be relatively local, and the composition of the asset purchases needs to be tailored to the desired effect.

III.ii Effects of the Fed's Large-Scale Asset-Purchase Programs

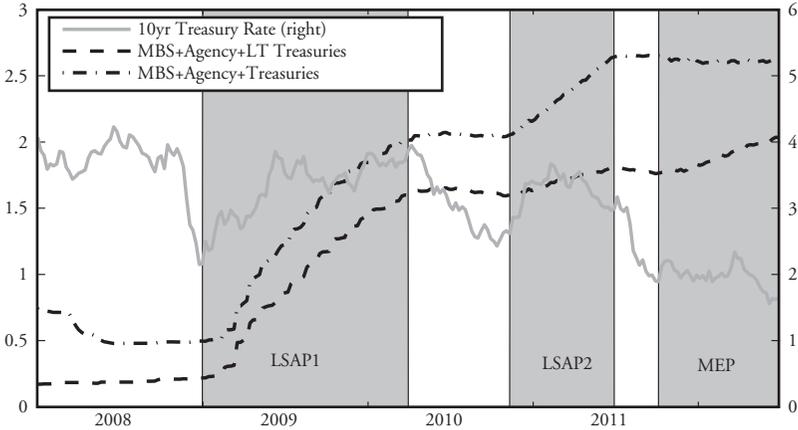
Since reaching the zero lower bound for its federal funds rate operating target late in 2008, the Federal Reserve has undertaken a

number of targeted asset-purchase programs. These include the Fed's first Large-Scale Asset Purchase program (LSAP1), announced in November 2008, and eventually involving purchases of \$1.75 trillion of longer-term securities (\$1.25 trillion of which were MBS), over a period through March 2010; a second program (LSAP2), hinted at in a speech by Chairman Bernanke in August 2010 and made definite in November 2010, under which another \$600 billion of longer-term Treasuries were purchased, over a period through June 2011; and the Fed's maturity extension program (MEP), announced in September 2011 and then extended in June 2012, which involves purchases of long-term Treasuries coupled with sales of an equal quantity of short-term Treasuries, in an amount that is envisioned to reach more than \$650 billion by the end of 2012. (The timing of these programs, and their consequences for the asset side of the Fed's balance sheet, are shown in Chart 16.⁴⁷) These programs have resulted in a substantial increase in the size of the Fed's balance sheet, but even more notably, they have resulted in a great increase in the degree to which the Fed holds longer-term securities, rather than relatively short-dated Treasury bills, on its balance sheet.

III.iii Have They Lowered Long-Term Interest Rates?

The declared intention of the programs has been to lower the market yields (and, hence, to raise the prices) of longer-term bonds (not necessarily limited to the particular types purchased by the Fed), with a view to easing the terms on which credit is available to both households and firms in the United States. Their effectiveness in this regard is a matter of considerable debate. As shown in Chart 16, the yield on 10-year Treasuries has shown a general downward trend since late 2008. But this should not necessarily be attributed solely to the Fed's purchases of longer-term securities over this period; the period is one in which a continuing series of bad news has progressively increased the likelihood that market participants are likely to attach to the possibility of a protracted period of feeble economic growth and low inflation (or even deflation), and of course the FOMC has progressively extended farther into the future the length of the period for which it anticipates keeping its funds rate target in a band just above zero. Hence, it is plausible to suppose that expectations

Chart 16
Evolution of Federal Reserve Securities
Holdings vs. 10-Year Treasuries



Notes: The gray regions identify three programs under which the Fed has increased its holdings of longer-term securities, discussed in the text. Units: (left scale) trillions of dollars; (right scale) percent per annum.

Source: Federal Reserve Bank of St. Louis.

regarding the length of time that short-term interest rates are likely to remain low have generally increased since the fall of 2008 (when it was not yet even obvious that the zero lower bound would be reached).

Indeed, a comparison of the timing of the increases in the Fed's holdings of long-term bonds with the timing of the declines in the 10-year yield does not obviously support a portfolio-balance interpretation of the overall decline in long-term interest rates; the 10-year yield actually rose over the course of the LSAP1 and LSAP2 programs, rather than declining. There was a substantial decline in the 10-year yield around the time LSAP1 was announced, but this might well be attributed to the dramatically worsening economic conditions (the very conditions that motivated the FOMC to consider so bold an experiment), rather than to the announcement of the FOMC's intended purchases. Long-term bond yields then began rising again as concerns about immediate economic collapse began to abate over the course of 2009, despite the fact that the Fed was increasing its bond purchases. The most important subsequent declines in long-term bond yields might similarly have had more to do with the deteriorating conditions that triggered the adoption of each

of the next two asset-purchase programs, than with anticipation that further purchases would occur.⁴⁸

A closer analysis of market movements around the time of particular announcements connected with the asset-purchase programs has, nonetheless, suggested that these announcements have had important effects on the market pricing of long-term bonds. In one particularly influential study, Gagnon et al. (2011) consider the market responses to eight specific official Fed communications that contained new information about the LSAP1 program (listed in the Appendix below), and look at changes in bond yields over a one-day window around each announcement. They argue that with one small exception,⁴⁹ the direction of the change in bond yields each day is the one that would be expected, under a portfolio-balance theory of the effects of Fed asset purchases, given the news in the statement—statements indicating larger asset purchases than had previously been announced were associated with reduced bond yields, and statements indicating smaller asset purchases with increased bond yields. The cumulative change in the level of long-term interest rates obtained by summing the one-day changes on the announcement days was substantial: a 91-basis-point decline in the yield on 10-year Treasuries, a 113-basis-point decline in the yield on agency MBS, and a 156-basis-point decline in the yield on 10-year agency debt. Moreover, there was even a cumulative decline of 67 basis points in the yields on Baa-rated corporate bonds, which the authors took as evidence that the LSAP program had a broad-based effect on the costs of borrowing, not limited to its effects on the prices of the particular types of assets purchased by the Fed.

Of course, taking the sum of the market movements on these announcement days only as a measure of the cumulative effect of the program as a whole—rather than, say, the cumulative change in long-term interest rates over the entire period of the program, shown in Chart 16—depends on believing that the program should only influence bond prices at the times when there is news that changes the expected size of the program, and that the effects of news are (nearly) immediate and permanent. The latter assumptions are familiar ones in event studies in financial economics, of course; but it

is important to recall that the justification for this familiar methodology is an assumption that *securities markets are “efficient,”* so that expected returns looking forward from any point in time are essentially constant (which requires that the effects of news on prices be realized instantaneously and not be subsequently reversed). Such an assumption is not obviously consistent with the existence of effects of Fed purchases on the prices of securities that the study is intended to demonstrate; for if the quantity purchased influences the price of a security, then the market is not efficient in the Samuelson-Fama sense. It is true enough that, as Gagnon et al. point out, there are many other factors that might be affecting the cumulative change in long-term bond prices over the entire LSAP1 period, so that one cannot regard that cumulative change as a measure of the effect of the program; but it is not clear that their announcement-days-only measure should be regarded as correct, either.

III.iib How Much of the Effect Is Due to News About Future Policy?

An even more serious question about the interpretation of such event studies is whether one should believe that the news about the likely size and character of the central bank’s asset purchases is the only thing that should move financial markets on that day. The one-day window is narrow enough that, in most cases, one can plausibly argue that the FOMC’s statement was the only big news affecting fixed-income markets that day; but it is less obvious that the only news *in the statement* was information about the likely size of the asset-purchase program. In particular, if the statement also contained information that *changed expectations about the future path of the federal funds rate*, then bond yields should have changed on those days, even in a world where there are no portfolio-balance effects.

And of course, because the FOMC concentrates its main official communications at particular points in time, it would hardly be coincidental for there to be news that should affect forecasts of the funds rate on the same days as there was news about the asset-purchase program. It is true that there were no changes in the target range for the funds rate after December 2008, but as explained above, expectations about how long the funds rate would remain at that level and how fast it might eventually rise did not remain constant, and should have

been critical determinants of long-term bond yields. As it happens, the two dates considered by Gagnon et al. on which there were the largest declines in long bond yields—accounting for 73 basis points out of the cumulative 91-basis-point decline that they report—were both dates on which there were very important statements about the funds rate target. These were Dec. 16, 2008 (on which, in addition to a dramatic immediate cut in the funds rate target from 1 percent to the zero to 25 basis point target band, the FOMC announced that it expected to maintain the low level “for some time”) and March 18, 2009 (on which the FOMC first announced that it expected to maintain the low level of the funds rate “for an extended period”). Attributing *all* of the declines in long-term bond yields on these days to the LSAP news is surely an exaggeration.

Even on dates when the FOMC statement contains no change in the funds rate target or in any explicit forward guidance with regard to the funds rate, it does not follow that expectations about the future path of the funds rate should not have changed as a result of the announcement. Each statement always contains a summary of the FOMC’s view of the outlook for real activity and inflation, and a statement indicating greater perceived downside risk or less worry about inflation on the horizon could be a reason to reduce the probability assigned to an increase in the funds rate anytime soon. (It would not be coincidental if there were more language of that sort in statements that also reveal that the Committee has decided to expand asset purchases.) In addition, the LSAP announcement itself might be taken to have implications for future interest-rate policy: a decision to increase asset purchases might be thought to reveal something about the shifting balance of influence among different opinions on the Committee, or about how troubling the FOMC’s information about the economic outlook really is.

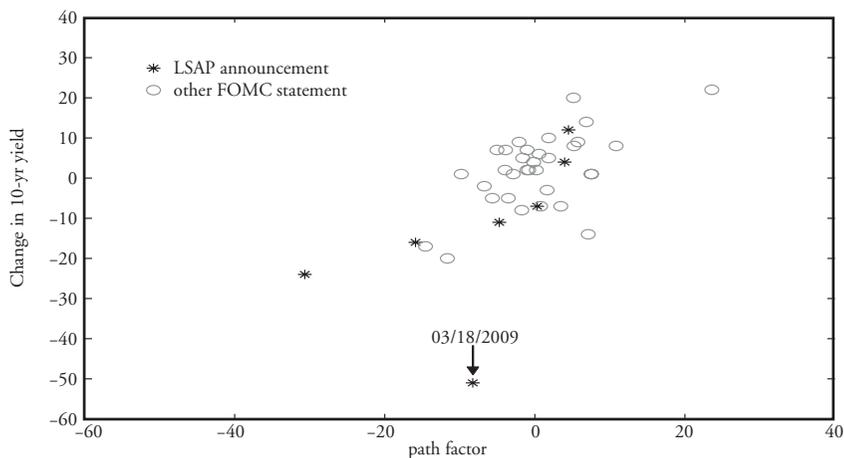
In fact, market-based measures of expectations regarding the future path of the funds rate indicate important shifts in expectations not only on the two dates just mentioned, but on other dates on which FOMC announcements contained news about asset purchases as well. One measure of the extent to which FOMC statements contain news about the future funds-rate path (as opposed to the

current funds rate) is the “path factor” computed using the method of Gürkaynak et al. (2005), on the basis of changes in fed funds futures prices around announcements (discussed above in Section I.ii.a). Campbell et al. (2012) compare this measure of the news about future interest-rate policy on announcement days with the change in the yield on 10-year Treasury notes on the same day.

As shown in Chart 17, not only is there a fairly strong positive correlation between the news about interest-rate policy and the change in the long-term bond yield, but, except for one outlier,⁵⁰ the points in the figure corresponding to important announcement dates connected with the LSAP1 and LSAP2 programs⁵¹ display a similar relationship between the two changes as the one that is found in the case of other FOMC statement dates on which there were no announcements about asset purchases. (The complete sample of dates considered in the figure includes all of the FOMC meetings between the beginning of the financial crisis in August 2007 and December 2011, plus the Nov. 25, 2008, press release announcing LSAP1.) The positive association between the “path factor” and changes in long-term bond yields on the dates of FOMC post-meeting statements confirms the result obtained by Gürkaynak et al. (2005) for an earlier period, that they had presented as evidence that the information about future interest-rate policy contained in FOMC statements affected long-term bond yields. The fact that one also observes on the LSAP announcement dates changes in bond yields reasonably similar to what one would expect, given this relationship, from the news in that announcement about future interest-rate policy (as measured by the “path factor”) suggests that the effect of the statement on bond yields may occur mainly as a result of the information conveyed about the likely path of the funds rate.

Another common market-based measure of expectations of the forward path of the funds rate is given by OIS rates, as also discussed above. Chart 18 compares the one-day changes in the 10-year zero-coupon Treasury yield⁵² on the dates of important FOMC announcements connected with the LSAP1, LSAP2 and MEP programs with the one-day change in the 10-year OIS rate on the same day. (Here, the eight LSAP1 announcement dates are the same ones as in the study of Gagnon

Chart 17
One-Day Changes in Yield on a 10-Year Treasury Note



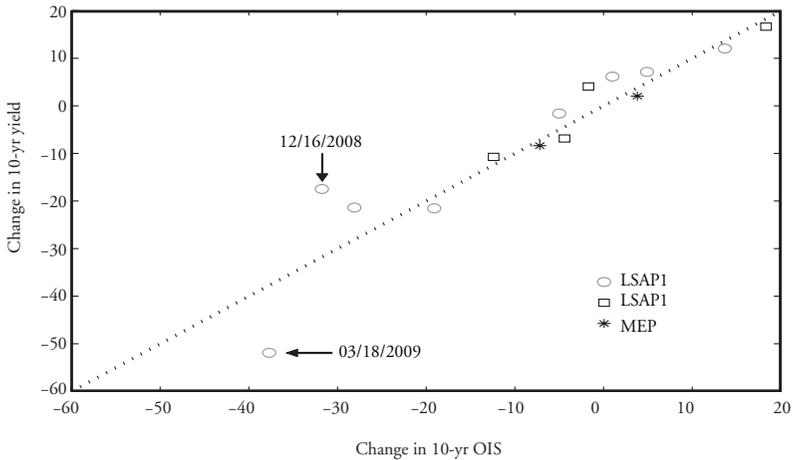
Notes: One-day changes in the yield on a 10-year Treasury note compared to the “path factor” measured from changes in federal funds rate futures on that day, for each of the FOMC announcement dates between August 2007 and December 2011. Dates of important announcements about LSAP programs are distinguished from other FOMC statements. Units (vertical axis): basis points.

Sources: Campbell et al. (2012) and author’s calculations.

et al. 2011, but additional dates relating to the LSAP2 and MEP programs are added for purposes of comparison.⁵³) As noted by Bauer and Rudebusch (2011), who make a similar comparison in the case of the eight LSAP1 dates, there is a great deal of similarity in the two changes. The figure shows that this is equally true of the announcements connected with all three of the FOMC’s major programs of long-term asset purchases.⁵⁴

If we take the change in the 10-year OIS rate as due solely to a change in expectations about the path of the funds rate over the next 10 years, this figure would suggest that much of the response of long-term Treasury yields on LSAP announcement dates can be attributed to the change in those expectations as a result of news in the FOMC’s statement. It would be too simplistic to insist that this is the only possible interpretation of the change in the OIS rate; while shorter-maturity OIS rates can reasonably be considered mainly to reflect the expected path of the funds rate over the relevant horizon (as discussed in Section I.iib), the 10-year OIS rate might well involve a nontrivial term premium. Bauer and Rudebusch (2011) suggest that one might nonetheless take the correlation shown in the figure

Chart 18
One-Day Changes in 10-Year Zero Coupon Treasury Yield



Notes: One-day changes in the 10-year zero-coupon Treasury yield compared to the change in the 10-year OIS rate on that day, for each of the dates of major FOMC announcements about asset-purchase programs listed in the Appendix. Units (both axes): basis points.

Sources: Bloomberg, Federal Reserve Board.

to indicate that portfolio-balance effects are small, under an identifying assumption according to which such effects should be specific to the particular security purchased by the central bank—and, hence, should affect the yields on Treasuries but not the swap rates. This would reflect one possible view of the type of market segmentation that could lead to substantial price effects of central-bank purchases, but it is not the only possibility. What the figure does show is that at any rate there seems to be fairly effective arbitrage between the OIS market and the market for Treasuries. It does not in itself prove that there is also effective arbitrage between the markets for bonds of different maturities, but as an indication that there were still traders looking out for arbitrage opportunities in the fixed-income markets, even in late 2008 and in 2009, it might incline one to find that hypothesis more likely.

Gagnon et al. (2011) consider the issue of the extent to which the changes in long-term bond yields that they measure may be due to changed expectations of the future path of short-term interest rates, using an estimated arbitrage-free dynamic term-structure model (DTSM) due to Kim and Wright (2005) to decompose changes in

bond yields into a part reflecting the change in the rationally-anticipated average level of short rates and a part reflecting a change in the size of the term premium. Using this approach, they conclude that 71 basis points of the cumulative decline in the 10-year bond yield represented a reduction in the 10-year term premium, rather than the expectations component.⁵⁵ They then propose this term premium effect as a measure of the effect on the bond yield due to Fed asset purchases, rather than to effects of the FOMC announcements on the expected path of short rates. Their estimate of the cumulative effect of the \$1.75 trillion of purchases under LSAP1 on the 10-year term premium is in fact the result that Gagnon et al. emphasize, and propose as the basis for an estimate of the further reductions in long-term interest rates that could be expected from additional purchases of long-term bonds.⁵⁶

This conclusion, however, should not necessarily be taken at face value. Bauer and Rudebusch (2011) argue, based on the work of Bauer et al. (2012), that the approach to DTSM estimation used by Kim and Wright results in biased coefficient estimates that exaggerate the degree of mean-reversion of the short-rate process, and consequently attribute too high a share of the movement in long rates to changes in term premia. They find that an LSAP1 event study focusing on the same eight announcement dates as Gagnon et al., but using term premia estimates based on a DTSM with reduced bias, finds a much larger role for changes in the expected path of short rates as an explanation for the declines in bond yields on those dates. Their point estimates suggest that about half of the cumulative decline in 10-year bond yields can be attributed to the expectations channel, though they stress that with a relatively unrestricted DTSM specification, the confidence intervals for a decomposition of this kind are wide, so that neither the “all-expectations” nor the “all-term-premia” hypothesis can be rejected.

This emphasis on the fragility of any conclusions obtained from DTSM estimation is surely correct. Apart from the general sensitivity of such estimation results to the details of model specification, one must particularly doubt the validity of any conclusions obtained by assuming that a model of interest-rate dynamics estimated using data prior to

the financial crisis will continue to be valid in a period when the federal funds rate is constrained by the zero lower bound.⁵⁷ Moreover, even if one were to accept the conclusion of Gagnon et al. (2011), according to which much of the decline in long-term bond yields on LSAP1 announcement days was due to declines in term premia on those days, rather than to changes in the expectations component of bond yields, this would not necessarily indicate the existence of portfolio-balance effects. After all, changes in the expected evolution of short rates should also effect term premia, and not just the expectations component. Suppose that FOMC forward guidance were to convince market participants that there was *no possibility of a* funds rate above zero for the next 10 years. In such a case, an absence of arbitrage opportunities would require the 10-year zero-coupon yield to fall to zero as well—meaning that *both* the expectations component and the term premium would be reduced to zero by such a change in expectations about the short-rate process. While the example is extreme, it illustrates a general principle: term premia are affected by expectations about the short-rate process (in particular, the degree of uncertainty about future short rates).

III.iic Further Channels for Policy Effects

Krishnamurthy and Vissing-Jorgensen (2011) instead seek to extend the event-study methodology of Gagnon et al. (2011), both by increasing the set of announcement dates studied to include FOMC announcements relating to the LSAP2 program,⁵⁸ and by looking at market movements on announcement days of a larger set of financial instruments, so as to be able to discriminate among a more finely-differentiated set of possible hypotheses about the channels through which central-bank purchases affect financial markets. For example, they consider a larger set of ways in which financial markets may have moved as a result of changed expectations about future macroeconomic conditions. In addition to the hypothesis of a change in the expected path of short-term interest rates also considered by Gagnon et al. (and called by these authors the signaling channel), they also consider the possibility of changes in the expected future rate of inflation, and changes in the degree of uncertainty about future inflation.

Like Bauer and Rudebusch, they find evidence that market yields were affected through the signaling channel in the case of LSAP1 announcements. Krishnamurthy and Vissing-Jorgensen base their conclusion on the shift downward in the expected future path of the federal funds rate indicated by fed funds futures prices on LSAP1 announcement days. They interpret the shift in the implied forward curve as a shift farther into the future of the funds-rate increases that were expected eventually to occur, and conclude that their five LSAP1 announcement days resulted in a cumulative shift of the expected timing of funds-rate increases to be later by more than six months. They similarly find evidence for the signaling channel in the case of the LSAP2 announcements, though the effects of the LSAP2 announcements were smaller (both through this channel and overall).

They also argue, by looking at movements in inflation-protected bonds as well as nominal bonds, and at options prices, that LSAP announcements both increased the expected rate of inflation and reduced inflation uncertainty (presumably by lowering the probability attached to deflationary “tail risk”); again they find some evidence of both inflation channels in the case of both LSAP1 and LSAP2 announcements. Like their conclusions about the signaling channel, these results indicate that financial markets were affected through expectational channels that are more indirect than the simple expectation of a mechanical effect of asset purchases on the market price of the assets purchased. One cannot say from these results alone why market participants changed their beliefs about future inflation. One interpretation could be that they expected the eventual Fed asset purchases to stimulate the economy through channels unrelated to the path of the federal funds rate; but another would be that the expected path of inflation changed because of what the announcement signaled about *future interest-rate policy*. In the latter case, while the effect on asset prices would not occur purely as a consequence of the expectations theory of the term structure, these “channels” would also reflect the information that the FOMC statements contained about likely future interest-rate policy, and not necessarily any portfolio-balance effects of asset purchases.

Krishnamurthy and Vissing-Jorgensen also use their comparisons of the effects of the announcements on different financial instruments to look for effects on a variety of possible types of risk premia that one would expect to change if portfolio-balance effects were important. They find little or no evidence of a “duration-risk channel” of the kind emphasized by Gagnon et al. (2011), according to which purchases of long-term bonds (of whatever type) should reduce a duration-risk premium common to all assets of a given duration. This is the theory under which Fed purchases of long-term Treasuries would be expected to reduce yields not only on those Treasury securities but on MBS and corporate bonds of similar duration as well—and, hence, the primary justification that has been offered for expecting desirable macroeconomic effects from programs like LSAP2 or MEP. They instead find that the effects of the LSAP1 announcements provide some evidence of a “prepayment risk channel” (under which Fed purchases of MBS reduced a risk premium that is specific to mortgage-related securities) and a “safety channel” (under which Fed purchases of safe assets, such as Treasuries or agency debt, increases the scarcity of such assets, lowering the yield on them relative to less-safe assets such as MBS or lower-rated corporate bonds).

Effects of the latter type are most naturally interpreted as portfolio-balance effects; for example, the “prepayment risk channel” could be expected to be important if the MBS market were segmented, with a class of specialized arbitrageurs who operate predominantly in the MBS market being the main ones to determine the size of the prepayment risk premium. But the authors find no role for a “prepayment risk channel” in the case of the LSAP2 announcements; in these cases, they find that the effects on MBS yields can be fully explained by the change in expectations of the future path of the federal funds rate. This might simply reflect the fact that LSAP2 did not involve further purchases of MBS; but it leaves open the possibility that the “prepayment risk channel” was important in the case of LSAP1 only because of the unusual disruption of the markets involved in mortgage securitization in 2008 and 2009. The degree of market segmentation required for such effects to be important might not characterize periods in which financial markets function more normally.

While Krishnamurthy and Vissing-Jorgensen again find evidence of a “safety channel” in the case of the LSAP2 announcements, they note that this channel is of little relevance for “the nominal rates that are most important for households and many corporations—mortgage rates and rates on lower-grade corporate bonds” (p. 255). They conclude that the effects of LSAP2 on these interest rates occurred mainly through the signaling channel or the two inflation channels—which is to say, through expectational channels that may simply have reflected changes in forecasts of future interest-rate policy.

The finding that an important part of the effects of the FOMC’s LSAP announcement came through expectational channels implies that a central bank would do well to give careful consideration to the signal about future policy that its announcement is intended to convey. The expectational effects were evidently not solely the consequence of explicit statements by the FOMC about future interest rates, and it may be (though it is hard to be certain) that the fact of the intended purchases was itself an important signal. Nonetheless, it would make sense to be explicit about the implications of such a decision for future policy, if one wants to have some degree of control over the inferences that are drawn, rather than leaving it to market participants to draw what inferences they may.

Furthermore, findings of Krishnamurthy and Vissing-Jorgensen suggest that, to the extent that there are portfolio-balance effects of Fed asset purchases, they are unlikely to be of the type relied upon as justification for the FOMC’s LSAP2 and MEP programs. There seems little reason to believe that purchases of long-term Treasuries should be an effective way of lowering the kind of longer-term interest rates that matter most for stimulating economic activity.⁵⁹ If one’s goal is to lower the yields on mortgage-backed securities in order to increase the availability of credit in the mortgage market, purchases of MBS are more likely to be effective, though even purchases of that kind would not necessarily be as effective under current conditions as they were under the unusual circumstances of 2009.⁶⁰

IV. Concluding Reflections

I have reviewed some of the leading approaches to the provision of additional monetary accommodation after the interest-rate lower bound is reached, both in the theoretical literature and among recent policy experiments. While it is difficult to be certain about the effects of such policies without a greater body of experience than is so far available, some provisional conclusions may be possible. The results that we have do not imply that the task of a central banker under current conditions is an easy one; there seem to me to be fewer options that are likely to be effective, and that are likely to be attractive on other grounds, than central bankers sometimes suggest when seeking to reassure the public.

Central bankers confronting the problem of the interest-rate lower bound have tended to be especially attracted to proposals that offer the prospect of additional monetary stimulus while (i) not requiring the central bank to commit itself with regard to future policy decisions, and (ii) purporting to alter general financial conditions in a way that should affect all parts of the economy relatively uniformly, so that the central bank can avoid involving itself in decisions about the allocation of credit. Unfortunately, the belief that methods exist that can be effective while satisfying these two *desiderata* seems to depend to a great extent on wishful thinking.

On the one hand, a number of central banks have sought to give signals about the likely future evolution of their policy rates, in order to influence current longer-term interest rates and hence, hopefully, current spending decisions; but this has typically been done in a way that avoids actually making any promises about how future policy decisions will be made, and, in particular, without giving listeners any reason to suppose that future decisions will be made on anything but a purely forward-looking basis. It is unclear why announcements of this form should have the desired effect, and on at least some occasions they seem to have little effect. Moreover, simply presenting a forecast that the policy rate will remain lower for longer than had previously been expected, in the absence of any reason to believe that future policy decisions will be made in a different way, runs

the risk of being interpreted as simply an announcement that the future is likely to involve lower real income growth and/or lower inflation than had previously been anticipated—information that, if believed, should have a contractionary rather than an expansionary effect (though it might lead long-term interest rates and indicators of interest-rate expectations to decline).

To the extent that central banks have been willing to make explicit commitments about the future conduct of policy, it has most often been to underline their commitment not to allow higher inflation than would correspond to their normal, long-run inflation target, even temporarily. Thus, in the case of the Federal Reserve, the introduction of more explicit forms of forward guidance and aggressive expansions of the Fed balance sheet have been accompanied by assurances that these policies should in no way suggest that there will be any relaxation of the FOMC's vigilance when it comes to preventing any increase in inflation. Such assurances tend to contradict precisely the kind of signals that one would want such policies to send in order for them to be effective in providing people a reason to spend more. They imply that simultaneous forecasts of low nominal interest rates for a longer time must indeed reflect pessimism about the speed of the economy's recovery rather than any change in the criterion that will be used to determine the appropriate timing of the interest-rate liftoff; and they work to reduce the extent to which asset purchases can affect the economy through either the signaling channel or the inflation risk channel found to be important by Krishnamurthy and Vissing-Jorgensen (2011).

A more useful form of forward guidance, I believe, would be one that emphasizes the target criterion that will be used to determine when it is appropriate to raise the federal funds rate target above its current level, rather than estimates of the liftoff date. If such an explicit criterion made it clear that short-term interest rates will not immediately be increased as soon as a Taylor rule descriptive of past FOMC behavior would justify a funds rate above 25 basis points, this would provide a reason for market participants to expect easier future monetary and financial conditions than they may currently be anticipating, and that should both ease current financial conditions

and provide an incentive for increased spending. An example of a suitable target criterion would be a commitment to return nominal GDP to the trend path that it had been on up until fall 2008. This would both make it clear that policy will have to remain looser in the near term than a purely forward-looking Taylor rule would imply, and, at the same time, provide assurance that the unusually stimulative current policy stance does not imply any intention to tolerate continuing inflation above the Fed's declared long-run inflation target—that in fact, it will not lead to a future level of nominal income any higher than what people had reason to anticipate at the time that they acquired their existing nominal assets and undertook their existing nominal obligations.

An obvious question about such a proposal remains: is there nothing the central bank can do to make people believe that it can and will achieve such a target, apart from declaring the target? Here I think it is important to recognize that a nominal GDP level target—unlike a target for the near-term rate of GDP growth, the near-term rate of inflation, or any other target tied to a specific horizon—need not lose credibility simply because the announcement fails to shift behavior immediately; for a lack of immediate progress in closing the nominal GDP gap would not be a reason to doubt that easier policy later would still be used to the extent necessary to close the gap—it would simply move the date by which that is expected to occur farther into the future. Nonetheless, it would undeniably be useful to be able to complement the announcement of the target with measures that would create immediate movement toward the target, without this movement depending entirely on the credibility of the achievement of the target—precisely in order to make the achievement of the target credible, whether it initially is or not.

Some argue that a vigorous program of quantitative easing is the obvious way to show that a central bank can and will act immediately, rather than simply waiting for expectations to change as a result of its announcements. But this argument presumes that central-bank asset purchases can stimulate additional spending, in ways not solely reliant upon expectational channels. Yet while central banks like the Federal Reserve and the Bank of England have been willing to

expand the size of their balance sheets rather dramatically in response to the recent crisis, they have often preferred to do this by purchasing extremely safe Treasury securities (the only kind of purchases undertaken under the FOMC's most recent two asset purchase programs, the LSAP2 and MEP programs, for example), in the hope that such purchases will improve general financial conditions without more direct involvement by the central bank in extending credit to particular sectors of the economy. Unfortunately, neither of the theories typically relied upon to explain why that should be the case—the quantity-theoretic doctrine that expansion of the monetary base must inevitably lead to increased aggregate nominal spending, or the particular kind of preferred-habitat model of the term structure that would imply the existence of a duration-risk channel—has a robust theoretical basis (in the sense of following from hypotheses that seem likely to be true, rather than relying upon special assumptions that might nonetheless conceivably be true) or finds much support from experience thus far.

It might, nonetheless, be argued that such purchases can be helpful as ways of changing expectations about future policy—essentially, as a type of signaling that can usefully supplement purely verbal forms of forward guidance. And one possible interpretation of the findings of Krishnamurthy and Vissing-Jorgensen (2011) that the Fed's LSAP programs have affected bond yields through expectational channels is that these purchases have been interpreted as signals about future policy. But if a central bank's intention in announcing such purchases is to send such a signal, the signal would seem more likely to have the desired effect if accompanied by explicit forward guidance, rather than regarded as a substitute for it.

Asset purchases also seem more likely to change expectations about future policy in a desired way if the purchases are rationally related to the future policy that they are intended to signal. For example, is there really a reason for a large expansion in the supply of bank reserves now to indicate that policy will also be more expansionary after the lower bound ceases to be a binding constraint? In the thought experiment of Auerbach and Obstfeld (2005), one could argue that such a connection exists: the immediate expansion of the monetary

base could be taken as a signal that the future monetary base will permanently be higher, because the immediate increase is of the same size as the permanent increase that one wants people to anticipate. It is true that there is no logical necessity for the long-run increase in the monetary base to be the same as the immediate increase; but it is at least a fairly simple and intuitively plausible idea to get across. (An announcement that the current increase in base money is intended to be permanent is more likely to be understood and believed than an announcement that the monetary base will be permanently increased at some future date, where both the date of the future action and the motive for the changed behavior at that time would likely seem obscure.) But the kinds of large increases in the monetary base associated with quantitative easing in Japan or with the Fed's recent programs do not suggest particular expectations about future policy in the same way: the expansions have been much too large for any plausible suggestion that they are intended to be permanent, nor is the size of the expansion tied in any obvious way to any aspect of the central bank's future targets that one might be trying to signal.

Clouse et al. (2000) suggest that purchases of long-term bonds might be a way of signaling a commitment to keep short-term interest rates low in the future, on the ground that the change in the central bank's balance sheet would then leave it vulnerable to capital losses in the event that it were to raise short-term rates sooner than it had wished people to believe it would. But it would be hard to defend the use of such a policy as a signal in order *not* to have to make any verbal commitments about policy. Compared to a verbal commitment to a history-dependent criterion for making later policy decisions, such as a nominal GDP target path, a don't-talk-but-buy-assets plan runs both a greater risk of tying the central bank's hands in a way that turns out to be awkward *ex post*—because it cannot allow an interest-rate commitment that is contingent on how the economy subsequently evolves—and a greater risk of failing to influence expectations in the desired way, because it relies upon market participants to correctly analyze the central bank's future incentives rather than directly stating its intentions. The proposal is somewhat more appealing if intended as a complement to a verbal commitment, since in that case the intended signal should be clearer. But it remains an

awkward and possibly costly form of signal, because the thing that the central bank should wish to signal is not a commitment to keep interest rates low for a fixed calendar period, but rather a commitment to maintain policy accommodation until the nominal GDP target path is reached.

A more logical policy would rely on a combination of commitment to a clear target criterion to guide future decisions about interest-rate policy with immediate policy actions that should *stimulate spending immediately* without relying too much on expectational channels. Neither a program of expanding the supply of bank reserves nor a program of expanding the central bank's holdings of longer-term Treasury securities is a good example of the latter kind of policy. Additional purchases of MBS by the Fed might instead still be useful as a way of reducing the cost of mortgage borrowing, though it is hard to be certain that additional purchases now would reduce MBS yields by the amount that the Fed's purchases under LSAP1 apparently did, given the less perilous situation of private financial intermediaries now, and it is hard to be certain that reductions in MBS yields would be passed on to mortgage rates. A kind of policy more certain to expand mortgage lending would be one like the Funding for Lending Scheme (FLS) recently announced by the Bank of England and the U.K. Treasury, which subsidizes lenders for increasing the amount of loans that they make.

Of course, it is not necessarily up to the central bank alone to institute policies of that kind, that can more directly influence private-sector decisions, for such actions are more properly viewed as part of fiscal policy. It is probably no accident that the FLS is a joint project of the Bank of England and the Treasury. And indeed, more generally, the most obvious recipe for success is one that requires coordination between the monetary and fiscal authorities. The most obvious source of a boost to current aggregate demand that would not depend solely on expectational channels is fiscal stimulus—whether through an increase in government purchases, tax incentives for current expenditure such as an investment tax credit, or subsidies for lending like the FLS. At the same time, commitment to a nominal GDP target path by the central bank would increase the bang for the buck

from fiscal stimulus, by assuring people that premature interest-rate increases in response to rising economic activity and prices would not crowd out other types of spending than those directly affected by fiscal policy. And the existence of the central bank's declared nominal GDP target path should also limit the degree of alarm that might arise about risks of unbridled inflation when special fiscal stimulus measures are introduced. For those who worry that fiscal stimulus always comes too late and goes too far, there would be the central bank's commitment to revert to a policy of active control of aggregate demand through monetary policy once the nominal GDP target path is reached.

Is the coordination between monetary and fiscal authorities required for such a solution in the realm of the possible? The U.K.'s FLS is at least a demonstration that coordination can occur. And while central bankers cannot bring about such coordination on their own, agreement among themselves about what needs to be done is surely a first prerequisite for such cooperation to be imaginable.

I would like to thank James Bullard, Vasco Cúrdia, Charles Evans, Gauti Eggertsson, Jonas Fisher, Narayana Kocherlakota, Edward Nelson, Adam Posen, Argia Sbordone, Lars Svensson, Eric Swanson and John Williams for helpful discussions, while absolving them from responsibility for the arguments presented. I also thank Kyle Jurado for research assistance, and the National Science Foundation for supporting my research on this issue. The opinions expressed are those of the author alone, and do not represent the views of the Federal Reserve Bank of New York, the Federal Reserve System, or Sveriges Riksbank.

A Appendix: FOMC Announcement Dates for Asset-Purchase Programs

This Appendix lists the announcement dates used in the various event studies discussed in Section III.ii. All 14 dates listed below are used in Chart 18. The first eight dates (the LSAP1 dates) are used in the event studies of Gagnon et al. (2011) and Bauer and Rudebusch (2011). The event study of Krishnamurthy and Vissing-Jorgensen (2011) [KVJ] uses data for both the LSAP1 and LSAP2 dates below, with the exception of those dates indicated in the footnotes. (Their study is thus based on a subset of only five of the LSAP1 dates and only two of the LSAP2 dates.) Chart 17, taken from Campbell et al. (2012), marks as “LSAP dates” each of the seven dates used by KVJ, except for Dec. 1, 2008 (which does not appear in the chart as it was not an official announcement by the FOMC), and also marks as an “LSAP date” the additional date Nov. 3, 2010, which KVJ discuss but end up excluding from their event study.

A.1 LSAP1 Dates

Date	Announcement
11/25/2008	The Federal Reserve will purchase “up to \$100 billion in GSE direct obligations,” and “up to \$500 billion in MBS.”
12/01/2008	In a speech, Chairman Bernanke states that the Federal Reserve “could purchase longer-term Treasury or agency securities...in substantial quantities.” ^a
12/16/2008	The FOMC “anticipates...exceptionally low levels of the federal funds rate for some time.” It also “stands ready to expand its purchases of agency debt and mortgage-backed securities...[and] is also evaluating the potential benefits of purchasing longer-term Treasury securities.”
01/28/2009	The FOMC “is prepared to purchase longer-term Treasury securities.”
03/18/2009	The FOMC “anticipates...exceptionally low levels of the federal funds rate for an extended period.” It will also purchase “up to an additional \$750 billion of agency mortgage-backed securities,” “up to \$100 billion” in agency debt, and “up to \$300 billion of longer-term Treasury securities over the next six months.”
08/12/2009	The FOMC “decided to gradually slow the pace” of Treasury purchases (“up to” language with reference to Treasury purchases is also removed). ^b
09/23/2009	The FOMC “will gradually slow the pace of” of agency MBS purchases (“up to” language with reference to agency MBS purchases is also removed). ^b
11/04/2009	The FOMC “will purchase...about \$175 billion of agency debt” (“up to” language with reference to agency debt is also removed). ^b

A.2 LSAP2 Dates

Date	Announcement
08/10/2010	The FOMC will reinvest “principal payments from agency debt and agency mortgage-backed securities in longer-term Treasury securities.”
08/27/2010	In a speech, Chairman Bernanke announces that “additional purchases of longer-term securities...would be effective in further easing financial conditions.” ^c
09/21/2010	The FOMC “is prepared to provide additional accommodation if needed.”
11/03/2010	The FOMC “intends to purchase a further \$600 billion of longer-term Treasury securities by the end of the second quarter of 2011, a pace of about \$75 billion per month.” ^d

^aExcluded by Campbell et al. (2012) as not an official announcement of the FOMC, though included by KVJ.

^bExcluded by KVJ as a minor announcement.

^cExcluded by KVJ as not an official announcement of the FOMC.

^dExcluded by KVJ as the announcement had been widely anticipated, and the news may actually have been smaller purchases than anticipated, but included by Campbell et al. (2012).

A.3 MEP Dates

Date	Announcement
09/21/2011	The FOMC “intends to purchase, by the end of June 2012, \$400 billion of Treasury securities with remaining maturities of 6 years to 30 years and to sell an equal amount of Treasury securities with remaining maturities of 3 years or less.”
06/20/2012	The FOMC “decided to continue through the end of the year its program to extend the average maturity of its holdings of securities.” An accompanying statement by the Federal Reserve Bank of New York clarifies that this continuation will “result in the purchase, as well as the sale and redemption, of about \$267 billion in Treasury securities by the end of 2012.”

^dExcluded by KVJ as the announcement had been widely anticipated, and the news may actually have been smaller purchases than anticipated, but included by Campbell et al. (2012).

Endnotes

¹In the case of the U.S., the federal funds rate has generally been trading 10-15 basis points below the rate of interest paid on bank reserves (IOR) held at the Fed (25 basis points). The IOR has not provided an absolute floor because some institutions with accounts at the Fed (notably the “government-sponsored enterprises”) cannot earn interest on them, and so are willing to lend overnight at a rate below the IOR, and evidently institutions that can earn the IOR are either sufficiently unwilling to borrow further, even to earn a sure return, or have sufficient monopsony power, to not have completely competed away this arbitrage opportunity (Bech and Klee 2011). Nonetheless, the spread remains small, despite a massive increase in the supply of reserves (as shown in Chart 16); so it is unlikely that the Fed would be able to push the funds rate much farther below the IOR, simply by further increasing the supply of reserves.

²In its response to the global financial crisis, the BOJ has again substantially increased the supply of bank reserves, but unlike the situation in the 200-06 period of “quantitative easing,” this has resulted in a reduction of the overnight rate only to 10 basis points, rather than to zero, because the BOJ has instituted an IOR of 10 basis points, for reasons similar to those cited by the Fed for maintaining a positive IOR. The fact that overnight rates were pushed to zero in the earlier period, when no interest was paid on reserves, indicates that this would be technically feasible.

³Even pure quantitative easing—adoption of a target for the supply of bank reserves beyond the level required to reduce overnight interest rates to the floor established by the the rate of interest paid on reserves—could in principle be a relevant dimension of policy away from the lower bound, if it were considered desirable to maintain a high degree of liquidity in the banking system, for reasons unrelated to the control of short-term interest rates, while using a variable IOR to implement desired variations in the policy rate. Such an approach to the implementation of interest-rate policy is recommended, for example, by Goodfriend (2002).

⁴The issues reviewed in this section are discussed in greater detail in Woodford (2005).

⁵See, for example, Woodford (2003, Chapter 4) for analysis of the mechanisms giving rise to this result.

⁶This is the reason why, in the numerical simulations of Eggertsson and Woodford (2003), even the expectation of a modest inflationary boom immediately following the return of the natural rate of interest to its normal level has a dramatic effect on the severity of both the economic contraction and the deflation that occur during the period of the negative natural rate.

⁷See also the discussion of these results in Bernanke et al. (2004), who develop their implications for the usefulness of forward guidance when policy is constrained by the zero lower bound.

⁸This is an oversimplification, as it neglects the consequences of interest-rate risk over that horizon.

⁹The word “commitment” is used in the title of the press release, as well as in the text.

¹⁰See Woodford (2005) for discussion of this earlier episode.

¹¹The statement on Dec. 16, 2008, had announced, among other things, a substantial cut in the current funds rate target; abandoned the FOMC’s previous practice of announcing a point target, in favor of a band; and announced that the Fed would “purchase large quantities of agency debt and mortgage-backed securities.” The statement on March 18, 2009, had not announced any change in the funds rate target, but specified the amounts of various types of long-term securities that would be purchased.

¹²Note that in the chart, the three-year rate has been shifted down by 10 basis points in order to show the several series on a single graph. This contract continues to trade at a rate above 20 basis points, contrary to how the figure may appear.

¹³See Gürkaynak et al. (2007a) for analysis of the usefulness of eurodollar futures prices as market-based forecasts of future FOMC policy.

¹⁴See Swanson and Williams (2012) for details of how the sensitivity index plotted in Charts 7 and 8 is constructed.

¹⁵After this point, the confidence interval includes a zero value for the sensitivity index. Note that negative values of the sensitivity index are possible, indicating responses to macroeconomic data surprises with a sign opposite to the usual one. Note also that in the chart, the sensitivity index is estimated to fall to a value insignificantly greater than zero slightly before the date of the FOMC’s new forward guidance. This may, however, simply reflect the fact that Swanson and Williams estimate the coefficient for each date using a one-year centered rolling window, so that reduced sensitivity after the August 2011 announcement is also reflected in the estimated sensitivity coefficients over a period up to six months prior to the announcement.

¹⁶Again, the drop in the sensitivity coefficient appears in the chart to occur slightly before the timing of the FOMC statement; but this is probably due to the use of a centered rolling window to estimate the time-varying coefficient.

¹⁷For example, one notes that the reduced sensitivity of near-term interest-rate expectations to macroeconomic news in 2003, shown in Chart 7, actually begins well in advance of the FOMC’s introduction of explicit forward guidance in August (when it announced that accommodation was expected to be maintained “for a considerable period”). The timing is more consistent with a view that market participants (correctly) expected short-term interest rates to be pinned at a low

level, with little room to vary, as concerns about possible deflationary risks began to grow.

¹⁸There is no direct measure of market expectations of the forecast. As a proxy, the authors use a weighted average of the previous day's futures rate, $f_{n,t-1}$, and the RBNZ's previous forecast (its forecast, a quarter earlier, of the 90-day rate $n + 1$ quarters in the future), with the relative weights on the two proxies determined to maximize the fraction of the variance of the changes in the futures rate that is explained by the regression.

¹⁹It is called "the repo rate" because at one time the bank's policy was implemented through lending at that rate under repurchase agreements, though this is not currently the case. It now defines the center of a corridor for the overnight rate, 20 basis points in width, maintained by the Riksbank.

²⁰See Svensson (2010, footnote 7) for more details. The implied forward rates include corrections for credit risk and maturity premia.

²¹The chart also shows the Riksbank's previously announced repo-rate path, from February, so as to show to what extent the new path represented a change from the bank's own most recent forecast.

²²Nonetheless, the Riksbank did cut the rate further at its July meeting.

²³The dashed gray line in the chart shows the repo rate path that had been projected in July, showing a steady series of small increases continuing into 2014.

²⁴His discussion refers to an earlier stage in the Riksbank's series of repo rate increases in 2010, when market expectations consistently failed to extrapolate a series of rate increases continuing to as high a level as the path projected by the Riksbank.

²⁵An awareness of divisions within the Executive Board may have contributed to such skepticism. Deputy Governors Karolina Ekholm and Lars Svensson have repeatedly dissented from the policy decisions of the majority in favor of lower repo-rate paths, since July 2010.

²⁶In discussing this pitfall of a forecast-targeting approach to monetary policy, I do not mean to assert that the approach described is necessarily that of Riksbank. At least some members of the Riksbank's Executive Board clearly understand the analytical point made here, and approaches to forecast targeting that would institutionalize history-dependence are discussed, for example, in Svensson and Woodford (2005) and Svensson (2005). It is not clear, however, that current Riksbank policy institutionalizes history-dependence of this sort, and still less that market participants have been given a reason to expect this.

²⁷The Fed itself took some pains to deny that it was attempting to use forward guidance as a tool of policy at all. Unlike the Bank of Canada, which explained its "conditional commitment" as an attempt to "provide more explicit guidance"

in order to “influence rates,” Chairman Bernanke was quoted as saying during the press conference following the release of the FOMC’s statement, “I wouldn’t overstate the Fed’s ability to massively change expectations through its statements” (*The New York Times* 2012).

²⁸In the analysis of Eggertsson and Woodford (2003), the long-run inflation target is assumed to be zero, so the simple rule that is analyzed is one in which the target for the gap-adjusted price level is constant, rather than growing at a constant rate; but all of the paper’s results directly generalize to the case of a long-run inflation target different from zero.

²⁹Levin et al. (2010) consider the issue under different numerical assumptions, and conclude that a constant gap-adjusted price level target would not be as close an approximation to fully optimal policy as in the results of Eggertsson and Woodford; in the case considered by these authors, the degree to which it would be optimal to ratchet up the price-level target during the period of the binding lower bound is greater. Nonetheless, the results of Levin et al. again indicate that a credible commitment to an adjusted price-level target would substantially improve upon the equilibrium that results from an inability to commit to any history-dependent policy. These authors stress that, under their numerical assumptions, the outcomes achievable through forward guidance are still significantly worse than what would be possible if the lower bound could be circumvented; hence they argue for the desirability of seeking to use additional instruments as well. Their work does not, however, show that the effects of forward guidance alone would be insignificant.

³⁰For examples of expression of this concern by Federal Reserve officials, see Kohn (2009) and Bernanke (2010).

³¹More technically, the adoption of such a policy commitment, even if ad hoc in the sense of only being introduced once the central bank finds itself in unexpected circumstances, is not an example of disregard for prior expectations of the kind that occurs under a purely discretionary approach to the conduct of policy, if the commitment is chosen to be optimal “from a timeless perspective” (as is the Eggertsson-Woodford rule). See Woodford (2011) for further discussion of the selection of policy commitments from a timeless perspective.

³²Most of the period to which the trend line is fit is not shown in the chart in order to highlight the recent shortfall in nominal GDP growth.

³³See, for example, the calculations of Hatzius and Stehn (2011), who extrapolate forward from the fourth quarter of 2007.

³⁴On the consistency of nominal income stabilization with the positions of Friedman and other monetarist authors, see, for example, McCallum and Nelson (1998), especially their footnote 4.

³⁵See, for example, Beckworth (2011) for an argument to this effect. Beckworth notes that Friedman (2003) praised the accuracy of “the Fed’s thermostat,” for

having reduced M2 growth during the period of increasing “velocity” in 1988-97, and then increased M2 growth by several percentage points during a period of decreasing velocity in 1997-2003. One might conclude that Friedman valued successful stabilization of nominal GDP growth more than strict fidelity to a “k-percent rule.”

³⁶Shiratsuka (2009) and Ueda (2012b) define “pure quantitative easing” in the way that the term is used here, though both point out that the actual unconventional policies of the BOJ during the 2001-06 period included other elements, in addition to the ones that they would classify as pure QE. Ueda (2012a) simply calls this kind of policy “quantitative easing,” distinguishing QE from “targeted asset purchases.”

³⁷After the BOJ began its “quantitative easing” policy, Friedman praised the action on the ground that “the Japanese central bank has finally started...buying long-term government securities and increasing high-powered money” (UPI 2001). On Friedman’s support for quantitative easing, see further Nelson (2012).

³⁸In a complete exposition, one should distinguish between the demand for currency and the demand for bank reserves. The present discussion has mainly to do only with the demand for bank reserves, since policies like the BOJ’s QE policy have little effect on the supply of currency.

³⁹The logic of equilibrium determination in such a “liquidity trap” is displayed in a simple two-period analysis by Krugman (1998), and expounded in the context of a complete intertemporal general equilibrium model by Eggertsson and Woodford (2003).

⁴⁰See Shiratsuka (2009) and Ueda (2012a, 2012b) for discussion of the BOJ’s policy in this period.

⁴¹Note that this was the type of forward guidance that, according to the argument made above, should be particularly likely to matter: one which stated a promise about the criteria that would determine future interest-rate policy, rather than simply offering a forecast of the forward path of the policy rate. However, the commitment was less ambitious than the kind recommended by the discussion above: because it focused only on the rate of inflation going forward, rather than committing to make up for past price-level declines, it was not a commitment that, even if fully credible, should have done much to eliminate the distortions resulting from the zero lower bound. In the model of Eggertsson and Woodford (2003), a fully credible commitment of this form would have no effect: the severely contractionary and deflationary equilibrium that they display is associated with policy expectations under which the policy rate is never raised above zero until it is possible to maintain a non-negative rate of inflation from then onward.

⁴²The first is the period beginning immediately after the failure of Lehman Brothers (in the third week of September 2008), when a number of special liquidity facilities and credit programs were either introduced or greatly expanded, and concluding with the end of the Fed's first series of large-scale asset purchases (LSAP1) in March 2010; the second is the Fed's second program of large-scale asset purchases (LSAP2), beginning in November 2010 and ending in June 2011.

⁴³Eggertsson and Woodford (2003) find in the context of a representative-household model that it does not matter which assets a central bank purchases in its open-market operations, precisely because of this reasoning.

⁴⁴This is the result obtained by Eggertsson and Woodford (2003), in the context of a representative-household model with transactions services represented by money in the utility function.

⁴⁵Examples of general-equilibrium analyses in which central-bank asset purchases are shown to affect both asset prices and the allocation of resources because of the existence of such constraints include Cúrdia and Woodford (2011) and Araújo et al. (2011).

⁴⁶For further discussion of the crisis in the commercial paper market, this Fed program and its effects, see Adrian *et al.* (2010) and Kacperczyk and Schnabi (2010).

⁴⁷The shaded gray region for each program indicates the period of time over which purchases occurred. LSAP1 is accordingly dated from Jan. 5, 2009, through March 2010; LSAP2 is dated from Nov. 12, 2010, through June 2011; and MEP is dated from Oct. 3, 2011, to the present. Some effects of the programs on bond yields might, of course, have occurred before the start of the gray region, to the extent the programs were announced prior to these start dates. Thus, LSAP1 might alternatively be dated from the FOMC's announcement of the program on Nov. 25, 2008, while LSAP2 might be dated from Chairman Bernanke's speech on Aug. 27, 2010, indicating support for such a program. See the Appendix for a chronology of important official statements relating to these programs.

⁴⁸The fact that the stock market has also declined in value during each of the three periods when there were sustained declines in long-term bond yields supports this interpretation.

⁴⁹On Sept. 23, 2009, long-term interest rates fell slightly despite FOMC language that might have been expected to reduce the size of the Fed's asset purchases, but the decline was modest: yields on 10-year U.S. Treasuries and 10-year agency debt each fell by 3 basis points that day. It is worth noting that the 10-year OIS rate fell by 5 basis points that day, which means that this date was not at all an outlier with respect to the pattern shown in Chart 18. A possible interpretation is that changes in the statement's implications for the likely future path of the funds rate were more important than its implications for the quantity of asset purchases.

⁵⁰This is the statement of March 18, 2009, described further in the Appendix, which contained both important news about the LSAP1 program and additional forward guidance with regard to interest-rate policy. The size of the decline in the 10-year Treasury yield that day was much larger than would normally accompany a negative “path factor” of the size measured that day, suggesting that the news that the Fed would purchase a large quantity of longer-term Treasuries did reduce Treasury yields, beyond what was justified by the reduction in expected future short-term interest rates. In the figure of Campbell et al. (2012), this point is an even larger outlier, as they compute a large positive, rather than a negative “path factor” for that day, but this reflects a doubtful measure of the change in futures prices that day. The measures used by both Krishnamurthy and Vissing-Jorgensen (2011, Table 2) and Bauer and Rudebusch (2011, Table 3) indicate that the expected federal funds rate declined at all future horizons, rather than rising sharply at longer horizons as implied by the calculations of Campbell et al. In Chart 17, the path factor for March 18, 2009, is calculated using the factor loadings of Campbell *et al.*, but using estimates of the futures rate changes for horizons longer than six months computed using the method of Bauer and Rudebusch (2011), which produces numbers fairly similar to those reported by Krishnamurthy and Vissing-Jorgensen. The path factors computed by Campbell et al. are plotted for all other announcement dates.

⁵¹See the Appendix for a list of these dates.

⁵²As in Bauer and Rudebusch (2011), these are constructed as in Gurkaynak *et al.* (2007b), and available online at <http://www.federalreserve.gov/econresdata/researchdata.htm> (accessed Aug. 10, 2012).

⁵³The complete list of announcement dates is in the Appendix.

⁵⁴As indicated in the chart, the two occasions on which there has been the largest discrepancy between the change in the OIS rate and the change in the bond yield are Dec. 16, 2008, and March 18, 2009—two dates that both involved important changes in the FOMC’s explicit forward guidance (as discussed in Section I.iib) in addition to announcements about the LSAP1 program. These were also the dates of the two largest changes in the OIS rate in this sample, and one might well suppose that the nature of the new information about likely future policy on these days differed from the other days in ways that go beyond merely being larger in magnitude. While the anomalous outcome on March 18, 2009, is consistent with the view that the FOMC’s announcement of an intention to purchase longer-term Treasuries increased the market price of those securities over and above the valuation that would be implied by expectations of the path of short rates (as was also true of the anomaly in Chart 17 connected with that day), the anomaly on Dec. 16, 2008, is of the opposite kind: the decline in the 10-year yield is smaller than the decline in the OIS rate, so that, to the extent that the discrepancy is thought to reflect a “portfolio-balance effect,” it would have the opposite sign to the one expected from the theory.

⁵⁵As Bauer and Rudebusch discuss, the cumulative decline in the fitted 10-year bond yield using the Kim-Wright model is actually 102 basis points (and not the 91-basis-point decline of the actual bond yields), so that this calculation still attributes 33 basis points of decline in the 10-year bond yield to the expectations component.

⁵⁶Rather than emphasizing the 71-basis-point estimate alone, they present results from a variety of possible approaches, leading to a range of estimates of the cumulative reduction in the term premium, and conclude that it “appears to be somewhere between 30 and 100 basis points, with most estimates in the lower and middle thirds of this range” (p. 38).

⁵⁷Swanson and Williams (2012) present a variety of types of interesting evidence showing that interest-rate dynamics have been different since late 2008, some of which has been discussed in Section I.iib.

⁵⁸At the same time, they reduce the number of LSAP1 announcement dates in their study, on the ground that including questionable dates could bias their conclusions about the relative importance of the different channels through which LSAP announcements affected markets, while excluding valid dates should only reduce the power of their hypothesis tests. See the Appendix for the five LSAP1 dates and two LSAP2 dates included in their study.

⁵⁹The effects of Treasuries purchases on the market yields of Treasuries can furthermore be taken as an indication that such purchases actually reduce welfare, as discussed in Section III.i and in Krishnamurthy and Vissing-Jorgensen (2012b).

⁶⁰There is also some question as to how well reductions in MBS yields are currently being passed through to mortgage loan rates (Bauer 2012).

References

- Adolfson, Malin, Stefan Laséen, Jesper Lindé, and Mattias Villani, "RAMSES: A New General Equilibrium Model for Monetary Policy Analysis," *Sveriges Riksbank Economic Review* 2007(2): 5-40.
- Adrian, Tobias, Karin Kimbrough, and Dina Marchioni, "The Federal Reserve's Commercial Paper Funding Facility," Federal Reserve Bank of New York Staff Report no. 423, January 2010.
- Araújo, Aloísio, Susan Schommer, and Michael Woodford, "Conventional and Unconventional Monetary Policy with Endogenous Collateral Constraints," unpublished, IMPA, December 2011.
- Archer, David, "Central Bank Communication and the Publication of Interest-Rate Projections," unpublished, Bank for International Settlements, June 2005.
- Auerbach, Alan J., and Maurice Obstfeld, "The Case for Open-Market Purchases in a Liquidity Trap," *American Economic Review*, 95: 110-137 (2005).
- Bank of Canada, "Bank of Canada Lowers Overnight Rate Target by ¼ Percentage Point to ¼ Per Cent and, Conditional on the Inflation Outlook, Commits to Hold Current Policy Rate Until the End of the Second Quarter of 2010," press release, April 21, 2009. (Available online at: <http://www.bankofcanada.ca/2009/04/press-releases/fad-press-release2009-04-21/>.)
- Bauer, Michael D., "Fed Asset Buying and Private Borrowing Rates," Federal Reserve Bank of San Francisco, *Economic Letter*, 2012-16, May 21, 2012.
- Bauer, Michael D., and Glenn D. Rudebusch, "The Signaling Channel for Federal Reserve Bond Purchases," Federal Reserve Bank of San Francisco, working paper no. 2011-21, December 2011.
- Bauer, Michael D., Glenn D. Rudebusch, and Jing Cynthia Wu, "Correcting Estimation Bias in Dynamic Term Structure Models," *Journal of Business and Economic Statistics*, 30: xxxx (2012).
- Bech, Morten L., and Elizabeth Klee, "The Mechanics of a Graceful Exit: Interest on Reserves and Segmentation in the Federal Funds Market," *Journal of Monetary Economics*, 58: 415-431 (2011).
- Beckworth, David, "Brad DeLong, Jim Grant, and Milton Friedman," posted June 17, 2011. (Available online at: <http://macromarketmusings.blogspot.com/2011/06/brad-delong-jim-grant-and-milton.html>)
- Bernanke, Ben S., "The Crisis and the Policy Response," Stamp Lecture, London School of Economics, Jan. 13, 2009. (Available online at: <http://www.federalreserve.gov/newsevents/speech/bernanke20090113a.htm>)

- Bernanke, Ben S., "The Economic Outlook and Monetary Policy," speech at the Federal Reserve Bank of Kansas City Symposium on Macroeconomic Challenges: The Decade Ahead, Jackson Hole, Wyo., Aug. 28, 2010.
- Bernanke, Ben S., Vincent R. Reinhart, and Brian P. Sack, "Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment," *Brookings Papers on Economic Activity*, 2004(2): 1-78.
- Campbell, Jeffrey R., Charles L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano, "Macroeconomic Effects of FOMC Forward Guidance," unpublished, Federal Reserve Bank of Chicago, May 2012.
- Chehal, Puneet, and Bharat Trehan, "Talking About Tomorrow's Monetary Policy Today," Federal Reserve Bank of San Francisco, *Economic Letter*, 2009-35, Nov. 9, 2009.
- Clouse, James, Dale Henderson, Athanasios Orphanides, David Small, and Peter Tinsley, "Monetary Policy When the Nominal Short-Term Interest Rate is Zero," *BE Press Journal of Macroeconomics: Topics in Macroeconomics*, 3(1): article 12 (2000).
- Congdon, Tim, "Keynes, Bernanke, and Krugman, and the Pathologies of Capitalism," in *Money in a Free Society*, New York: Encounter Books, 2011.
- Cúrdia, Vasco, and Michael Woodford, "The Central-Bank Balance Sheet as an Instrument of Monetary Policy," *Journal of Monetary Economics*, 58: 54-79 (2011).
- Eggertsson, Gauti B., and Michael Woodford, "The Zero Bound on Interest Rates and Optimal Monetary Policy," *Brookings Papers on Economic Activity*, 2003(1): 139-211.
- Evans, Charles L., "The Fed's Dual Mandate Responsibilities and Challenges Facing US Monetary Policy," remarks delivered at the European Economics and Financial Centre Distinguished Speaker Seminar, London, Sept. 7, 2011. (Text available online at: http://www.chicagofed.org/webpages/publications/speeches/2011/09_07_dual_mandate.cfm)
- Friedman, Milton, "The Fed's Thermostat," *The Wall Street Journal*, Aug. 19, 2003.
- Gagnon, Joseph, Matthew Raskin, Julie Remache and Brian Sack, "The Financial Market Effects of the Federal Reserve's Large-Scale Asset Purchases," *International Journal of Central Banking* 7(1):3-43 (2011).
- Goodfriend, Marvin, "Interest on Reserves and Monetary Policy," Federal Reserve Bank of New York, *Economic Policy Review*, May 2002, pp. 77-84.
- Goodfriend, Marvin, "Central Banking in the Credit Turmoil: An Assessment of Federal Reserve Practice," *Journal of Monetary Economics* 58: 1-12 (2011).

- Guha, Krishna, Aline van Duyn, and Anuj Gangahar, "Rate of US Job Losses Falls Sharply," *Financial Times*, June 6, 2009.
- Gürkaynak, Refet S., Brian P. Sack, and Eric T. Swanson, "Do Actions Speak Louder than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 1(1): 55-93 (2005).
- Gürkaynak, Refet S., Brian P. Sack, and Eric T. Swanson, "Market-Based Measures of Monetary Policy Expectations," *Journal of Business and Economic Statistics*, 25: 201-212 (2007a).
- Gürkaynak, Refet S., Brian P. Sack, and Jonathan H. Wright, "The U.S. Treasury Yield Curve: 1961 to the Present," *Journal of Monetary Economics*, 54: 2291-2304 (2007b).
- Hatzius, Jan, and Sven Jari Stehn, "The Case for a Nominal GDP Level Target," Goldman Sachs US Economics Analyst no. 11/41, Oct. 14, 2011.
- Kacperczyk, Marcin, and Philipp Schnabl, "When Safe Proved Risky: Commercial Paper during the Financial Crisis of 2007-2009," *Journal of Economic Perspectives*, 24(1): 2950 (2010).
- Kim, Don H., and Jonathan H. Wright, "An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates," Federal Reserve Board FEDS discussion paper no. 2005-33, September 2005.
- Kohn, Donald L., "Monetary Policy Research and the Financial Crisis: Strengths and Shortcomings," speech at the Conference on Key Developments in Monetary Economics, Washington, D.C., Oct. 9, 2009.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy," *Brookings Papers on Economic Activity* 2011(2): 215-265.
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, "The Aggregate Demand for Treasury Debt," *Journal of Political Economy*, 120: 233-267 (2012a).
- Krishnamurthy, Arvind, and Annette Vissing-Jorgensen, "Why an MBS-Treasury Swap is Better Policy than the Treasury Twist," unpublished, Kellogg School of Management, July 2012b.
- Krugman, Paul R., "It's Baaack: Japan's Slump and the Return of the Liquidity Trap," *Brookings Papers on Economic Activity*, 1998(2): 137-206.
- Kydland, Finn E., and Edward C. Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85: 473-491 (1977).

- Lam, W. Raphael, "Bank of Japan's Monetary Easing Measures: Are They Powerful and Comprehensive?" IMF Working Paper 11-264, November 2011.
- Levin, Andrew, David López-Salido, Edward Nelson, and Tack Yun, "Limitations on the Effectiveness of Forward Guidance at the Zero Lower Bound," *International Journal of Central Banking*, 6(1): 143-189 (2010).
- Luckett, Dudley G., " 'Bills Only': A Critical Appraisal," *Review of Economics and Statistics*, 42: 301-306 (1960).
- McCallum, Bennett T., and Edward Nelson, "Nominal Income Targeting in an Open Economy Optimizing Model," NBER Working Paper no. 6675, August 1998.
- Moessner, Richhild, and William R. Nelson, "Central Bank Policy Rate Guidance and Financial Market Functioning," *International Journal of Central Banking*, 4: 204-226 (2008).
- Nelson, Edward, "Friedman's Monetary Economics in Practice," unpublished, Federal Reserve Board, April 2012.
- The New York Times*, "Fed Signals That a Full Recovery Is Years Away," Jan. 25, 2012.
- Oda, Nobuyuki, and Kazuo Ueda, "The Effects of the Bank of Japan's Zero Interest Rate Commitment and Quantitative Monetary Easing on the Yield Curve: A Macro-Finance Approach," *The Japanese Economic Review*, 58: 303-328 (2007).
- Okina, Kunio, and Shigenori Shiratsuka, "Policy Commitment and Expectation Formation: Japan's Experience Under Zero Interest Rates," *North American Journal of Economics and Finance*, 15: 75-100 (2004).
- Romer, Christina D., "Dear Ben: It's Time For Your Volcker Moment," *The New York Times*, Oct. 29, 2011.
- Shiratsuka, Shigenori, "Size and Composition of the Central Bank Balance Sheet: Revisiting Japan's Experience of the Quantitative Easing Policy," Bank of Japan IMES Discussion Paper no. 2009-E-25, November 2009.
- Sumner, Scott, "Re-Targeting the Fed," *National Affairs*, Fall 2011, pp. 79-96.
- Svensson, Lars E.O., "Monetary Policy with Judgment: Forecast Targeting," *International Journal of Central Banking*, 1(1): 1-54 (2005).
- Svensson, Lars E.O., "Policy Expectations and Policy Evaluations: The Role of Transparency and Communication," Sveriges Riksbank, *Economic Review*, 2010(1): 43-78.
- Svensson, Lars E.O., "Practical Monetary Policy: Examples from Sweden and the United States," *Brookings Papers on Economic Activity*, 2011(2): 289-332.

- Svensson, Lars E.O., and Michael Woodford, "Implementing Optimal Policy through Inflation-Forecast Targeting," in B.S. Bernanke and M. Woodford, eds., *The Inflation Targeting Debate*, Chicago: University of Chicago Press, 2005.
- Sveriges Riksbank, Monetary Policy in Sweden, June 3, 2010. (Available online at: http://www.riksbank.se/Upload/Dokument/riksbank/Kat_publicerat/Rapporter/2010/Monetary_policy_2010.pdf)
- Swanson, Eric T., and John C. Williams, "Measuring the Effect of the Zero Lower Bound on Medium-and Longer-Term Interest Rates," Federal Reserve Bank of San Francisco Working Paper no, 2012-02, May 2012.
- Ueda, Kazuo, "The Effectiveness of Nontraditional Monetary Policy Measures: The Case of the Bank of Japan," *Japanese Economic Review*, 63: 1-22 (2012a).
- Ueda, Kazuo, "Deleveraging and Monetary Policy: Japan Since the 1990s and the United States Since 2007," *Journal of Economic Perspectives*, 26(3): 177-202 (2012b).
- Ugai, Hiroshi, "Effects of Quantitative Easing Policy: A Survey of Empirical Analyses," *Bank of Japan Monetary and Economic Studies*, 25: 1-47 (2007).
- United Press International, "Friedman Opposes Stimulus Package," Oct. 9, 2001.
- Wallace, Neil, "A Modigliani-Miller Theorem for Open-Market Operations," *American Economic Review*, 71: 267-274 (1981).
- Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton: Princeton University Press, 2003.
- Woodford, Michael, "Central-Bank Communication and Policy Effectiveness," in *The Greenspan Era: Lessons for the Future*, Federal Reserve Bank of Kansas City, 2005.
- Woodford, Michael, "Optimal Monetary Stabilization Policy," in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, vol. 3B, Amsterdam: Elsevier, 2011.
- Woodford, Michael, "Forecast Targeting as a Monetary Policy Strategy: Policy Rules in Practice," in E.F. Koenig, R. Leeson, and J.B. Taylor, eds., *The Taylor Rule and the Transformation of Monetary Policy*, Stanford, CA: Hoover Institution Press, 2012.