

Inflation Targeting and Financial Stability*

Michael Woodford
Columbia University

November 10, 2011

*Revised text of a talk given as a keynote address at a conference, “The Future of Central Banking,” at the Einaudi Institute for Economics and Finance, Rome, September 30, 2010. I would like to thank Jean Boivin, Steve Cecchetti, Vasco Cúrdia, Stefan Ingves, Jean-Pierre Landau, Florencio Lopez-de-Silanes, Rick Mishkin, Benoit Mojon, and Lars Svensson for helpful discussions, Claes Berg for editorial comments on an earlier draft, and the (U.S.) National Science Foundation for research support.

A number of commentators have suggested that central banks should reconsider the desirability of inflation targeting in the light of the global financial crisis. Early on, Paul DeGrauwe (2007) asserted that the crisis had “unveiled the fallacy” of the consensus view in favor of inflation targeting as an approach; a little later, Axel Leijonhufvud (2008) argued that inflation targeting “has failed” as a strategy, and that “the problems we now face are in large part due to this policy failure”; and more recently, Francesco Giavazzi and Alberto Giovannini (2010) have proposed that inflation targeting, as conventionally practiced, “can ... increase the likelihood of a financial crisis.”

How seriously should inflation-targeting central banks take these charges? I think it is important to distinguish between inflation targeting as such and the more specific doctrine — enunciated by some prominent proponents of inflation targeting, but not, in my view, a defining feature of this approach to the conduct of monetary policy — according to which central banks need not pay attention to asset prices, or more generally to concerns relating to financial stability, when making monetary policy decisions.

I do not believe that the central claims that were made by proponents of inflation targeting on behalf of this approach are challenged in any direct way by the events of the crisis. It is worth recalling what inflation targeting was intended to achieve. It was expected, above all, to serve to stabilize medium-term inflation expectations. This, it was asserted, would allow monetary policy to be used more aggressively for purposes of stabilization of the real economy, without so much sacrifice of price stability as would be required in the absence of such well-anchored inflation expectations. It was expected to eliminate a particular source of macroeconomic instability, namely, the possibility of wage-price spirals triggered by commodity-price shocks, of the kind that had been problematic in the 1970s. And it was expected to allow countries to avoid the possibility of a deflationary trap of the kind experienced by many countries in the 1930s, in which expectations of deflation, once entrenched, become self-fulfilling.

Failures of any of these central claims to be born out in practice would give one serious reason to reconsider the basic theory of inflation targeting. But thus far they have held up quite well. Rather than discrediting inflation targeting, one could argue that the events of the last several years have provided further vindication for it. Despite a serious disruption

of the world financial system, that some have compared in magnitude to that suffered in the 1930s, this time none of the major economies fell into deflationary spirals. And despite large swings in oil prices, the effects on the dynamics of wages and prices this time have been modest. These comparatively benign outcomes are surely due in large part to the fact that inflation expectations in most of the major economies have remained quite well anchored in the face of these substantial disturbances. And it is arguable that the credibility with regard to control of the rate of inflation that the leading central banks have achieved over the past twenty years deserves a great deal of the credit for this stability.

Of course, the global financial crisis has done great damage, and this has understandably led to questions as to whether the disaster might have been avoided, or its severity reduced, had policies been different. The aspects of policy that have most obviously been called into question have to do with the regulation of the financial system. But it is also worth asking whether alternative monetary policies might have made a difference.

In particular, the crisis does justify reconsideration of at least one aspect of the inflation targeting doctrine that had developed over the previous two decades. This is the thesis that a central bank with an inflation target need not pay attention to financial developments — such as a credit-financed real estate boom — except to the extent that such developments affect the outlook for inflation (or perhaps, either for inflation or for real activity). While this thesis is not, in my view, a central, definitional aspect of an inflation targeting regime, it was undoubtedly a common view among proponents of inflation targeting prior to the crisis. It is therefore important to reconsider both the extent to which such a view is defensible, and the extent to which it is a necessary element of a coherent approach to inflation targeting. Can this previously conventional view still be maintained, after recent experience? And if not, would this require abandonment of inflation targeting as well?

I shall begin by reviewing some common arguments that have been offered for setting aside the question of financial stability in the conduct of monetary policy. I conclude that there is a persuasive case for taking this issue into account, as at least one factor, when making decisions about interest rates. But I shall also argue that it is possible to do this in a way that represents a natural extension of “flexible inflation targeting,” as that concept has been developed in the literature prior to the crisis. It should thus be possible to adapt

the framework used to structure monetary policy deliberations in a way that takes account of legitimate concerns raised by the recent crisis, but without having to discard what was learned from the previous quarter century of experience with and analysis of methods of inflation stabilization.

1 Is Financial Stability Relevant to Monetary Policy Deliberations?

Prior to the global financial crisis, many (though certainly not all) central bankers took the view that considerations of financial stability should play no role at all when making decisions about monetary policy. A variety of arguments were offered in defense of this view, and it is worth discussing them briefly before proposing my own view of the matter.

One of the simplest arguments was that, however desirable it might be to act to head off financial crises were one able to do so, such crises are simply *not predictable enough* for there to be any point in trying to “lean against” developing financial-sector risks. This view gained particular credence when the issue was cast as one of using monetary policy to lean against (or even to “prick”) asset “bubbles,” which were in turn defined as situations in which the market price of some asset was significantly higher than its fundamental value. How, it was asked, should central banks expect to know the correct valuation of assets if the correct value was not sufficiently obvious for market participants to have gotten it right. Because “bubbles” are, by their nature, situations that are difficult to identify until after they have burst, it was argued to be more practical for a central bank to simply plan to “mop up” after the crash of the bubble than to try to prevent it from occurring.

But complacency about the ease of “mopping up” after a financial crisis is much more difficult after the recent global crisis; despite unprecedented and heroic efforts on the part of a number of central banks, it was not possible to prevent a very sharp contraction of world trade and economic activity, and even years later many economies are still struggling with the after-effects of the crisis. And the excuse that crises are unpredictable is not as compelling as it might at first seem. After all, in order for it to be useful to adjust policy in

order to reduce the risk of financial crisis, one needn't be able to predict exactly when crises will occur; it suffices that one be able to identify circumstances under which the *risk* of a crisis increases (and that there be policies that can affect these risks). It is true enough that our understanding of how to measure such risks is much more incomplete than we should want. But there are indicators that have been found to have predictive value (e.g., Borio and Drehmann, 2009), and it is hard to justify not trying to improve our ability to measure financial crisis risks.

It is important, I believe, to realize that the real issue is not identifying whether one type of asset or another is currently overvalued. Instead, what central banks (and potentially other “macro-prudential” regulators) need to be able to monitor is the degree to which the positions taken by leveraged institutions pose a *risk to financial stability*. Of course, a belief that multiple institutions have each borrowed in order to invest in an asset the value of which is likely to collapse, because its current price is far above its true value, is one possible reason to believe that there is a substantial risk of a systemic crisis; but a central bank need not be able to identify asset over-valuations in order to recognize situations in which the probability of simultaneous financial distress at several institutions is non-trivial. The typical case against which the central bank should be on guard is not one in which the *mean* of the distribution of possible future net worths for the institutions is too low, but rather one in which the *lower tail* of the distribution is too large. Moreover, the question of greatest concern is not even the size of the lower tail of outcomes for individual institutions, but the probability of a bad *joint* outcome. This question of systemic risk is not one with which individual institutions may have much concern in their financial decisions, and so the belief that is useful for the central bank or other regulators to assess systemic risk does not depend on a belief that the regulators are able to forecast better than private institutions can.

A second ground for skepticism about the relevance of financial stability concerns in monetary policymaking is based on doubts about how much monetary policy can do to influence the buildup of risks to financial stability, even granting that it might be possible to identify such risks in real time. Adjustment of the short-term interest rates controlled by central banks will have little effect on stock-market or real-estate “bubbles,” it is often argued — if short rates are relevant to such valuations at all, the change in monetary policy

required to make a difference would be very severe, and, given the unpredictability of the evolution of such bubbles, the effects of the sudden, sharp change in monetary policy would be difficult to predict. Again it is often concluded that it should be easier to “mop up” afterwards than to try to contain a bubble as it develops, on the ground that it is clearer what monetary policy can do to help once the problem becomes a shortage of liquidity.

But once again, I think that many discussions of this point dismiss the potential relevance of monetary policy too easily, by posing the question as one of using interest-rate policy to control “bubbles” in asset prices. The real issue, I would argue, should not be one of controlling the possible mis-pricing of assets in the marketplace — where the central bank has good reason to doubt whether its judgments should be more reliable than those of market participants — but rather, one of seeking to deter extreme levels of leverage and of maturity transformation in the financial sector. Once the problem is recast in this way, the relevance of interest-rate policy decisions — whether to exacerbate the problem or to mitigate it — is more obvious. Even modest changes in short-term rates can have a significant effect on firms’ incentives to seek high degrees of leverage or excessively short-term sources of funding.¹ Again, this is something that we need to understand better than we currently do; acceptance that monetary policy deliberations should take account of the consequences of the policy decision for financial stability will require a sustained research effort, to develop the quantitative models that will be needed as a basis for such a discussion. But there is certainly no ground, on the basis of current economic knowledge, to assert that interest-rate policy is likely to be irrelevant.

A third ground for skepticism would assert that, even if one grants that monetary policy might be able to influence the risk of occurrence of a financial crisis, there are surely better tools available for this purpose. It is best to assign only one goal to each available policy instrument, in accordance with what is sometimes called the “Tinbergen principle,” and in that case monetary policy is not the right tool to use to ensure financial stability. That should instead be the task of supervisory policy, of regulatory policy, or perhaps of new

¹Woodford (2011b) provides an example of an explicit model in which monetary policy decisions affect the endogenous capital structure decisions of intermediaries, and as a consequence, the severity of the “fire sale” externalities associated with a crisis state.

instruments of “macro-prudential policy” such as countercyclical capital requirements.

It is indeed true that these other aspects of policy should have an important role in maintaining financial stability. While I have just argued that it is plausible to believe that monetary policy has an effect on the risk of financial crisis, it hardly follows from this that the interest-rate policy of the central bank can or should provide a *complete solution* to the problem. That would be true only if one believed not only that interest-rate policy can be a very effective tool to deal with the problem, but that there are no costs to subordinating interest-rate policy to that end. The latter is surely not the case, as the model sketched in the next section is intended to illustrate. Hence acceptance of the proposition that monetary policy is relevant to financial stability is no excuse for failing to improve bank regulation, tighten capital requirements, or develop additional tools of macro-prudential policy.²

But by the same token, the existence of other instruments that can help to reduce the risk of a financial crisis does not, in general, justify complete neglect of the issue of financial stability in monetary policy deliberations. That would be true only if one could count on the other policy instruments to *completely eliminate* the problem of financial instability, and without other costs of having to resort to those instruments. This is unlikely, and at any rate, it is certainly not the situation in which central banks already find themselves. The recent crisis points up the weakness of the existing regulatory and supervisory regimes in many countries, and while many reforms are currently under discussion, it is too soon to be certain how much will change and how effective the new structures will be at controlling risk-taking in the financial sector. Central banks should certainly applaud the development of other tools that can help to minimize the risks to financial stability, as this can only make their own task simpler and more effective; but until it is clear that the problem has genuinely been solved by those other means, it would be prudent for them to also develop analytical capability for thinking about the impact of their own actions on financial stability.

Still, one might reasonably ask, will there not be a conflict between the use of monetary

²Woodford (2011b) illustrates how the creation of an additional macro-prudential policy instrument, such as variation in the interest rate paid on reserves, can improve both financial stability and the central bank’s ability to achieve its traditional stabilization objectives of price stability and full utilization of productive capacity.

policy to control risks to financial stability, and the use of it to maintain price stability and stable real activity? Yes, I think there will almost inevitably be a tension between these alternative objectives, as the model in the next section illustrates. But I wish to argue that this tension is no different, in principle, than the conflict between inflation stabilization and output-gap stabilization, in the conventional theory of “flexible inflation targeting.” Proponents of inflation targeting generally admit that the interest-rate policy required to maintain complete stability of prices will not always be the same one that would best stabilize aggregate output around its efficient level. And yet, in mainstream accounts of inflation targeting — certainly in the view of it espoused by the theorists of inflation targeting, such as Mervyn King, Ben Bernanke, and Lars Svensson, who are actually involved in the conduct of monetary policy — it does not follow that one must therefore set aside all concern with the effects of interest-rate policy on the real economy.³ Rather, it is argued that a sound approach will seek to balance a concern for the effects of policy on real activity with a concern for its effects on inflation; and it is furthermore argued that it should be possible to use policy to mitigate short-run instability of the output gap without any substantial sacrifice of the stability of medium-run inflation expectations.

The view that I wish to propose of the place of financial stability concerns in monetary policy deliberations is a similar one. I think that central banks should admit that monetary policy may well have consequences for financial stability, rather than pretending that the issue should not be their responsibility because they have no influence over it; and that they should recognize that it would require considerable luck for the policy that best serves their traditional stabilization objectives to turn out always to coincide perfectly with the one that is best from the standpoint of financial stability. Accordingly, I believe that it is appropriate for a “flexible inflation targeting” central bank to endeavor to balance financial stability objectives against both its price stability objective and its concern for output-gap stabilization, when choosing among alternative short-run paths for the economy at a given conjuncture. Nonetheless, I shall argue that it is possible to do this through a straightforward adaptation of the way that inflation-targeting central banks already think about the short-run tradeoff between price stability and output stability, and that once again an allowance

³See, e.g., King (1997), Bernanke *et al.* (1999), and Svensson (2011).

for other objectives in choosing among short-run transition paths should not require any compromise of the primacy of price stability as the central bank's longer-run objective.

2 A Simple Model of Optimal Stabilization with Endogenous Financial Crises

A simple model may be helpful in clarifying the way in which an inflation targeting regime could be modified to incorporate concerns for the effects of monetary policy decisions on financial stability. In order to address the concerns raised above, it is essential that the occurrence of crises that disrupt financial intermediation not be treated as purely exogenous, as it is in analyses such as those of Cúrdia and Woodford (2009, 2011), Del Negro *et al.* (2010), or Gertler and Karadi (2011), that treat only the question of how central-bank policy can mitigate the effects of a crisis in the event that one occurs. Here I shall not propose any sophisticated model of the endogenous mechanisms that give rise to a crisis — a complex topic that is the subject of much ongoing research — but will instead simply postulate a reduced-form model of the way in which endogenous state variables affect the probability of a crisis, and consider how allowance for such a relationship would change the standard theory of optimal monetary stabilization policy.⁴ Analysis on the basis of such a crude hypothesis can at best be regarded as suggestive, rather than prescriptive. Nonetheless, if one believes that a relationship of this general type is important, even though a correct specification would be more complex, a simple analysis of the kind offered here may be more useful than an analysis that assumes there are no such effects at all.

2.1 Sketch of the Model

Let us consider a simplified version of the model of the macroeconomic effects of credit frictions developed in Cúrdia and Woodford (2009). The most important effect of the credit frictions in that model is to modify the relation that would otherwise exist between aggregate real expenditure and the path of interest rates. The model is one in which households

⁴That standard theory, abstracting from financial frictions, is reviewed in some detail in Woodford (2011a).

are heterogeneous, and at a given point in time, some are credit-constrained while others are not, and the marginal utilities of income of the two types differ as a result. (With frictionless financial intermediation, the marginal utilities of the different households would co-move perfectly, despite the differences in their incomes and spending opportunities.) A key additional state variable of the model is Ω_t , a measure of the gap that exists at any time t between the marginal utilities of income of the two types. This variable measures the distortion of the allocation of expenditure due to credit frictions — a larger value of Ω_t means that the marginal utility of borrowers exceeds that of savers to a greater extent, which means that the spending by borrowers is inefficiently low to a greater extent — and hence is a useful measure of the severity of credit frictions. In the Cúrdia-Woodford model, this variable also corresponds to a *credit spread* between two different long-term bond yields: the spread between the equilibrium yield on long-term bonds (of a particular duration) issued by risky private borrowers on the one hand and those issued by the government on the other.⁵ An empirical correlate of this state variable would therefore be an average spread between yields on risky corporate bonds and those on Treasury securities of a comparable maturity.

The reason this variable is important for the positive predictions of the model is that variations in Ω_t shift the predicted relation between aggregate real expenditure and the average marginal utility of income. In a representative-household model (or a model without financial frictions), the marginal utility of income should be a decreasing function of aggregate expenditure; this structural relationship can be shifted by exogenous changes in government purchases, household impatience to consume, or the marginal efficiency of investment opportunities, among other factors (the various sources of “IS disturbances” in a standard New Keynesian model). In the Cúrdia-Woodford model, this relation is also shifted by changes

⁵In the paper, we show that, to a log-linear approximation, the variable Ω_t (actually denoted $\hat{\Omega}_t$ in the paper) will be a forward-looking moving average of the short-term credit spread (denoted by $\hat{\omega}_t$), where the short-term spread is the differential between the one-period interest rate at which private non-financial borrowers can borrow and the one-period interest rate on government liabilities. Hence Ω_t can alternatively be expressed as the difference between forward-looking moving averages of those two different short-term interest rates, which would correspond to the spread between the yields on certain long-term bonds issued by the two types of borrowers. The hypothetical bonds for which this would be exactly the credit spread would be claims to a stream of future payments that are exponentially declining at a certain rate.

in Ω_t . Under the calibration proposed there as most realistic, a higher value of Ω_t will lower the marginal utility of income associated with a given level of aggregate expenditure, as a consequence of the less efficient composition of expenditure; an increase in Ω_t thus has effects similar to those of a reduction in government purchases or a reduction in the attractiveness of current private spending opportunities.

Because of this modification, the “intertemporal IS equation” of the basic (three-equation) New Keynesian model takes the more general form

$$y_t - g_t + \chi\Omega_t = E_t[y_{t+1} - g_{t+1} + \chi\Omega_{t+1}] - \sigma[i_t - E_t\pi_{t+1}], \quad (2.1)$$

where y_t denotes the output gap (i.e., the amount by which the log of aggregate real expenditure exceeds the currently efficient level, which latter quantity is assumed to depend solely on exogenous factors), g_t is a composite both of the various exogenous factors that shift the relation between the marginal utility of income and aggregate expenditure even in the absence of credit frictions⁶ and of those that shift the efficient level of aggregate output, i_t is a short-term nominal interest rate, π_{t+1} is the rate of inflation between periods t and $t+1$, and all variables denote deviations from their steady-state values (so that constants are omitted). Under the proposed calibration, the coefficients satisfy $\chi, \sigma > 0$. In the presence of credit frictions, the variable i_t (a weighted average of the interest rates that are relevant for borrowers and savers respectively) is no longer identical with the central bank’s policy rate, and this introduces an additional term if the IS equation is instead to be written in terms of the policy rate, as in Cúrdia and Woodford (2009). Here I omit that complication, as I am not interested in deriving a rule for the particular instrument adjustment required to achieve particular macroeconomic targets; for purposes of the present discussion, it suffices that the i_t in (2.1) is a variable that the central bank can influence (even if the influence is not quite so direct as in the case of the policy rate).

Under this calibration, real aggregate demand depends not only on exogenous factors (such as the evolution of government purchases) and the expected path of (average) real interest rates, but also on the magnitude of the distortions indicated by credit spreads; other things equal, a larger value of Ω_t will depress aggregate demand in period t . Thus the

⁶For example, an increase in government purchases increases the value of the term g_t .

additional Ω_t terms in (2.1) can be thought of as representing what are sometimes called “financial headwinds.”

For similar reasons, the model’s aggregate supply relation must be modified relative to the familiar “New Keynesian Phillips Curve” specification, taking now the form

$$\pi_t = \kappa_y y_t + \kappa_\Omega \Omega_t + \beta E_t \pi_{t+1} + u_t, \quad (2.2)$$

where the coefficients satisfy $\kappa_y, \kappa_\Omega > 0, 0 < \beta < 1$, and u_t is a composite term representing various possible exogenous “cost-push” factors. The credit frictions change this relationship only through the appearance of the Ω_t term, again reflecting the way that changes in Ω_t shift the relationship between aggregate real expenditure and the marginal utility of income. One of the reasons for an upward-sloping short-run aggregate supply curve is that higher real activity is associated with a lower marginal utility of income, which increases real wage demands, and hence the real marginal cost of supplying goods. Since larger credit frictions also reduce the average marginal utility of income, for a given level of real activity, they also increase the real marginal cost and hence the inflationary pressure resulting from a given level of real activity.

The crucial new element that I wish to consider here is some degree of endogeneity of the evolution of the financial distortion measure $\{\Omega_t\}$.⁷ I shall simplify by assuming that Ω_t is always in one of two states: either it takes a low value $\underline{\Omega}$ (the “normal” state) or a high value $\bar{\Omega}$. I shall furthermore suppose that the probability each period of transition from the crisis state back to the normal state (conditional on being in the crisis state) is $0 < \delta < 1$, while the probability γ_t of transition from the normal state to the crisis state (conditional on being in the normal state) is time-varying, and moreover (at least possibly) a function of endogenous macroeconomic conditions. It is this potential endogeneity of the probability γ_t of occurrence of a crisis that raises the question of the implications of monetary policy

⁷Cúrdia and Woodford (2009) already allow for one specific type of endogeneity of financial distortions: in their model, Ω_t is a forward-looking moving average of the short-run credit spread ω_t , which is allowed to depend on the current volume of privately intermediated credit, in addition to various exogenous factors. This endogenous dependence of spreads on the volume of credit can be thought of as movement along a “supply curve for intermediation” of the kind proposed in Woodford (2009); shifts in the location of the supply curve, however, are purely exogenous disturbances in the model of Cúrdia and Woodford (2009).

decisions for financial stability.

The assumption that financial conditions jump between two discrete states — in one of which credit spreads are low, and in the other of which they are high — is obviously an oversimplification, but it captures something important about financial crises of the kind with which we are here concerned: that they are typically characterized by sudden, substantial increases in credit spreads that are instead relatively stable under normal circumstances. A regime-switching model is a parsimoniously parameterized way of capturing this episodic character of periods of financial stress, as in the empirical model of Davig and Hakkio (2010).⁸ An advantage of this approach is that it responds to a common complaint about policy analyses using DSGE models, namely, that the use of local perturbation methods necessarily abstracts from the possibility of occasional excursions far from the normal range of variation in the state variables as a result of nonlinearities — which extreme outcomes are precisely the ones that one must be concerned about in an analysis of risks to financial stability. A regime-switching model allows for a non-trivial probability of occasional excursions far from normal conditions, and allows the probability of such excursions to be endogenous (the critical issue for the present discussion). It does not seek to model the nonlinear mechanisms that actually allow a relatively abrupt transition to another part of the state space to occur, instead contenting itself with a reduced-form model of the probability of such an excursion occurring and a specification of the conditions that result from one on average. But this is about as specific a model as we can expect to parameterize on the basis of available empirical evidence, anyway, given the heterogeneity and relative infrequency of crises. And it allows us to use local perturbation methods to analyze the linkages between the various endogenous variables of such a model — including the transition probabilities and the values of endogenous variables *conditional* on the regime that one is in — without this requiring any assumption that crises do not involve large changes in the values of many variables.

For purposes of illustration, I shall here assume one very simple kind of endogeneity of the transition probability, Suppose that $\gamma_t = \gamma_t(L_t)$, where L_t is a measure of the degree of

⁸In their model, the two states are characterized by different mean levels of the Kansas City Fed Financial Stress Index, many elements of which are credit spreads (Hakkio and Keeton, 2009).

leverage in the financial sector, and $\gamma_t(\cdot)$ is a function satisfying

$$\gamma_t(L), \gamma'_t(L), \gamma''_t(L) > 0.$$

(The time subscript on the function means that there can also be exogenous shifts in this function over time.) The idea of the positive dependence on leverage is that the more highly levered financial institutions are, the smaller the unexpected decline in asset values required to tip institutions into insolvency — or into a situation where there may be *doubts* about their solvency — and hence the smaller the exogenous shock required to trigger a crisis. Given some distribution function for the exogenous shocks, the lower the threshold for a shock to trigger a crisis, the larger the probability that a crisis will occur over a given time interval. Moreover, not only does greater leverage increase the probability that any given bank will come to be in financial difficulty as a result of an exogenous shock, it also increases the probability that a given bank's financial distress will tip others into distress as well, so that the chance of a chain reaction of significant magnitude. If the overall crisis state represented by $\Omega_t = \bar{\Omega}$ occurs only when such a chain reaction occurs, then the probability is likely to be sharply increasing in the degree of leverage beyond some point, though it might well remain relatively constant (and low) for all degrees of leverage below some threshold.

Of course, both the risk of individual banks' insolvency and the risk of a chain reaction occurring depend on more than just the banks' leverage ratios: for example, the degree of maturity mismatch and liquidity mismatch between their assets and their liabilities is highly important as well. Here I shall use a single variable L_t to stand for a variety of changes of this kind in financing arrangements that increase the risk of financial instability, and all of which tend to increase in periods when there is less risk avoidance by financial institutions. The use of a single variable to summarize the relevant change in financial structure simplifies the calculations below, and allows a fully optimal commitment to be described using fairly simple equations. I do not mean, of course, to suggest that in practice an adequate model will take account of only one dimension of financial structure, still less that precisely the same single variable will be the most useful one for all countries. Local institutional details are likely to matter even more for this aspect of the model than for structural relations such as (2.1) and (2.2); the present analysis is intended as an illustration of a general approach

to the problem, rather than as a presentation of a formula that can be directly applied once one correctly associates local data series with the various symbols.

It remains to connect leverage (or financial risk-taking more generally) with the other endogenous variables of the model. I shall slightly simplify the dynamics of the Cúrdia-Woodford model by postulating a simple law of motion of the form

$$L_t = \rho L_{t-1} + \xi y_t + v_t, \tag{2.3}$$

where v_t is an exogenous disturbance term and the coefficients satisfy $0 < \rho < 1, \xi > 0$. This reduced-form relation combines two structural equations, one relating the growth of aggregate bank assets (and hence the aggregate leverage L_t) to the rate at which new loans are originated, and one relating the rate of new borrowing to the level of aggregate activity (and hence to the output gap y_t).

The rate of new lending (the equilibrium volume of intermediation) is an increasing function of the level of activity if an increase in incomes increases the demand for intermediation, as in the derivation of the “XD curve” in Woodford (2010). In the model of Curdia and Woodford (2009), the relation is increasing because expenditure by the borrowers is assumed to be more sensitive to variations in the interest rate at which they can borrow than expenditure by savers is sensitive to variations in the interest rate earned on their savings. A shift in monetary policy that increases aggregate expenditure also necessarily increases the share of expenditure by borrowers, and so increases spending by borrowers more than it increases their incomes, requiring an increase in new borrowing. The disturbance v_t may reflect any of a variety of factors: a shift in the relationship between bank assets and leverage, or more generally, in the degree of risk that banks must take in order to finance assets of a given volume, perhaps because of events that reduce bank capital or shift the penalties associated with risk-taking; a shift in the relationship between the level of economic activity and the output gap, due for example to a productivity shock; or a shift in the degree to which an expansion of demand requires credit expansion, due for example to shocks with different effects on borrowers and savers. To the extent that the changes in these relationships can be treated as exogenous (and unaffected by monetary policy), they can be lumped together in a single composite disturbance term.

Our simple structural model then consists of the three equations (2.1) – (2.3) together with the regime-switching model of the evolution of financial conditions, to determine the four endogenous variables $\pi_t, y_t, L_t, \Omega_t$ each period, given the central bank’s adjustment of the path of i_t and the paths of the exogenous disturbances (including the timing of the regime transitions).

2.2 A Criterion for Optimal Policy

I shall further assume that the goal of policy is to minimize a loss function of the form

$$\frac{1}{2} E_0 \sum_{t=0}^{\infty} \beta^t [\pi_t^2 + \lambda_y y_t^2 + \lambda_\Omega \Omega_t^2], \quad (2.4)$$

for some weights $\lambda_y, \lambda_\Omega > 0$, and a discount factor $0 < \beta < 1$ that is the same as in (2.2). This adds to the usual inflation and output-gap stabilization goals, standard in accounts of flexible inflation targeting (e.g., Svensson, 2011), an additional goal of reducing the incidence of financial crises (reflected by the assumption that $\lambda_\Omega > 0$). Cúrdia and Woodford (2009) show that an objective of this form can be justified as a quadratic approximation to the goal of maximizing the average expected utility of households in their DSGE model with credit frictions, in the case that all disturbances are small enough for a second-order Taylor series expansion to give a correct ranking of welfare under alternative policies.⁹ The additional term represents the welfare consequences of the distortion of the composition of expenditure as between credit-constrained and unconstrained households (or between borrowers and savers); because this distortion is minimized when $\Omega_t = 0$ (i.e., the marginal utilities of income of the two types are equal), the welfare effects of this distortion can be approximated by a term of the form $\lambda_\Omega \Omega_t^2$.¹⁰

⁹Of course, in the present analysis, we do not necessarily wish to assume that the difference between $\bar{\Omega}$ and $\underline{\Omega}$ is small, so that the Taylor series expansion invoked by Cúrdia and Woodford may not be valid. However, the conclusion that welfare is reduced by spending a greater fraction of time in the crisis state — which is the key implication of (2.4) for the analysis below — is also true in the Cúrdia-Woodford model, even in the case that financial distortions in the crisis state are large.

¹⁰In the derivation in Cúrdia and Woodford (2009), an additional term appears in the loss function relating to resources consumed by the intermediary sector; I here assume those to be negligible, in order to simplify the analysis.

We may now consider how policy should be conducted in order to minimize (2.4). If we abstract from the possibility that the zero lower bound on nominal interest rates can constrain the central bank’s ability to move i_t as far as it would otherwise wish, then the problem can be reduced to the choice of state-contingent paths for the variables $\{\pi_t, y_t, L_t, \Omega_t\}$ consistent with constraints (2.2) – (2.3) and the transition equation for the regimes so as to minimize (2.4). Equation (2.1) can then simply be solved to determine the required path for the short-term nominal interest rate $\{i_t\}$ in order to implement the desired evolution of the other variables.

It is perhaps useful first to consider the special case in which the transition probability γ_t is independent of L_t . In this case, the evolution of the regimes (and hence the evolution of the financial distortion factor $\{\Omega_t\}$) is purely exogenous, and the third term in the objective is independent of policy. It then suffices to consider how policy affects the expected discounted value of the other two terms in the objective, which is to say, the traditional inflation and output-gap stabilization objectives. Furthermore, the evolution of $\{L_t\}$ has no welfare consequences in this case, so we can ignore constraint (2.3). The problem then reduces to the choice of state-contingent paths for inflation and the output gap to minimize the traditional loss function in analyses of “flexible inflation targeting” subject to the aggregate-supply constraint (2.2). Both the objective and constraint are thus as in standard treatments (such as Clarida *et al.*, 1999) that abstract from financial frictions, with the exception that the exogenous fluctuations in Ω_t result in an additional additive (exogenous) disturbance term in (2.2). Essentially, variations in financial conditions represent another source of “cost-push” disturbances, in addition to the u_t shock already allowed for in Clarida *et al.* (1999).

In this case, as is well known, the first-order conditions for an optimal policy commitment¹¹ imply that the evolution of inflation and the output gap must satisfy

$$\lambda_y y_t - \kappa_y \varphi_t = 0 \tag{2.5}$$

$$\pi_{t+1} + \varphi_{t+1} - \varphi_t = 0 \tag{2.6}$$

in each possible state of the world for all $t \geq 0$, where φ_t is a Lagrange multiplier associated with the constraint (2.2) in any given state of the world in period t . (There is also an

¹¹See either Clarida *et al.* (1999) or Woodford (2011a, sec. 1) for derivation of these conditions.

additional, generally different, constraint linking the initial period inflation rate π_0 and the initial multiplier φ_0 ;¹² but we need not be concerned with the form of this additional condition here, as it affects the form of the optimal target criterion for the initial period only.)

Elimination of the Lagrange multipliers then implies that the evolution of inflation and output must satisfy

$$\pi_t + \phi_y(y_t - y_{t-1}) = 0 \tag{2.7}$$

in every period $t \geq 1$, where $\phi_y \equiv \lambda_y/\kappa_y > 0$. Moreover, this condition (plus a modified version of the condition that applies in the initial period only) suffices to uniquely determine the state-contingent evolution of inflation and output.¹³ Hence satisfaction of the *target criterion* (2.7) is necessary and sufficient for the evolution of inflation and output to correspond to an optimal policy commitment, and the optimal policy can be implemented through a *forecast targeting* procedure, under which each time that policy is reconsidered, the central bank verifies that its intended forward path for policy continues to imply forward paths for inflation and the output gap that are expected (conditional on the economy's state at the time of the exercise) to satisfy the target criterion at all future horizons.¹⁴

Alternatively, the first-order conditions (2.5) – (2.6) imply that under optimal policy, there must exist a constant price level target p^* such that

$$p_t + \phi_y y_t = p^* \tag{2.8}$$

for all $t \geq 0$, where p_t is the log of the general price level, so that $\pi_t \equiv p_t - p_{t-1}$. Note that under this formulation, it is possible to write a target criterion with the same form in all periods $t \geq 0$, rather than needing a different criterion for the initial period. A criterion of the form (2.8) implies that (2.7) holds for all $t \geq 1$ regardless of the value of p^* ; it is then necessarily possible to choose the value of p^* so that (2.8) also holds at $t = 0$ under

¹²The form of this condition differs depending whether we consider unconstrained “Ramsey policy” or optimal policy “from a timeless perspective,” as discussed in Woodford (2011a, sec. 1).

¹³That is, there is a unique set of bounded state-contingent paths for inflation and the output gap that satisfy both (2.2) and (2.7) each period, and this is the optimal state-contingent evolution of these variables.

¹⁴The advantages of forecast targeting as a practical approach to implementation of an optimal policy commitment are discussed in Svensson (2005, 2011), Svensson and Woodford (2005) and Woodford (2007).

the optimal policy. This alternative formulation of the optimal target criterion — as what might be called a *flexible price-level target* — is particularly convenient when we introduce the additional complications associated with endogenous financial crises.

In the special case that γ_t is unaffected by variations in leverage, then, optimal policy continues to be described by a target criterion that involves only the projected paths of inflation and of the output gap, and that has *exactly the same* quantitative form as in a model without any credit frictions. Monitoring financial conditions is nonetheless necessary for the conduct of monetary policy for two reasons: (i) the state-contingent paths for inflation and output that will be feasible and at the same time consistent with the target criterion will generally depend on the (exogenous) evolution of financial conditions, owing to the “cost-push” effects of financial crises when they occur; and (ii) the interest-rate path required to bring about a given evolution of inflation and output will also depend on financial conditions, owing to the “financial headwinds” terms in (2.1). This simple result is essentially the reason why Cúrdia and Woodford (2009) find that the target criterion (2.7) continues to provide a good approximation to optimal policy even in the presence of (largely exogenous) financial disruptions.

However, matters are more complex if we allow γ_t to be an increasing function of L_t , as assumed earlier. In this case (as shown in the Appendix), the first-order condition (2.5) takes the more general form

$$z_t = \beta \rho E_t z_{t+1} - \beta \xi X_t, \quad (2.9)$$

where z_t is the expression on the left-hand side of (2.5), and X_t is a variable defined as

$$X_t \equiv \gamma'_t(L_t) \cdot \Delta V_{t+1|t} \quad (2.10)$$

if $\Omega_t = \underline{\Omega}$ (that is, when the economy is in the normal state), while $X_t \equiv 0$ if $\Omega_t = \bar{\Omega}$. Finally, in expression (2.10), the variable $\Delta V_{t+1|t}$ indicates the difference that occurrence of the crisis state as opposed to the normal state will make to the expected continuation value in period $t + 1$, conditional on the state in period t , *i.e.*,

$$\Delta V_{t+1|t} \equiv E[V_{t+1}|h_t, \Omega_{t+1} = \bar{\Omega}] - E[V_{t+1}|h_t, \Omega_{t+1} = \underline{\Omega}], \quad (2.11)$$

where V_{t+1} is the policymaker’s continuation loss looking forward from period $t + 1$ (defined more precisely in the Appendix), and h_t is the state of the world (history of all exogenous

disturbances, including the history of regime transitions) in period t . In addition, first-order condition (2.6) continues to apply.

As before, (2.6) implies that under an optimal policy commitment, there must exist a constant log price level p^* such that $\varphi_t = p^* - p_t$ for all $t \geq 0$, which in turn implies that

$$(p_t - p^*) + \phi_y y_t = \kappa_y^{-1} z_t \quad (2.12)$$

for all $t \geq 0$. When γ_t is independent of L_t , (2.10) implies that $X_t \equiv 0$ at all times, so that (2.9) can be solved forward to yield the (unique bounded) solution $z_t = 0$ at all times, from which it follows that optimal policy requires that (2.7) be satisfied at all times. But more generally, solving (2.9) forward yields

$$z_t = -\beta\xi \sum_{T=t}^{\infty} (\beta\rho)^{T-t} E_t X_T, \quad (2.13)$$

substitution of which into (2.12) yields the more general target criterion

$$p_t + \phi_y y_t + \phi_X E_t \sum_{T=t}^{\infty} (\beta\rho)^{T-t} E_t X_T = p^*, \quad (2.14)$$

where $\phi_X \equiv \beta\xi/\kappa_y > 0$. As in the special case, conformity to the criterion (2.14) in all periods is a necessary and sufficient condition for bounded state-contingent processes $\{\pi_t, y_t, L_t\}$ to constitute an optimal policy commitment.

Thus even in the more general case, in which the probability of occurrence of a crisis is endogenous, optimal policy is characterized by conformity to a certain type of flexible price-level target. The difference now is that the target criterion does not depend solely on the paths of the general price level and of the aggregate output gap; it also depends on the projected evolution of another variable, X_t , which for shorthand I shall refer to as the *marginal crisis risk*. As indicated by the definition (2.10), this variable measures the rate at which the expected loss from the occurrence of a financial crisis increases per unit increase in leverage. Under the assumptions that I have made above, the value of this quantity should always be non-negative: equal to zero when the economy is in a crisis state, but always (at least slightly) positive under normal conditions. Even conditional on being in the normal regime ($\Omega_t = \underline{\Omega}$), though, the size of this positive quantity is likely to vary over time. It may be larger than normal either because the amount of damage that a crisis would do (through

the distortion of the allocation of resources that will result from a disruption of financial intermediation) is currently greater than usual (a larger value of $\Delta V_{t+1|t}$), or because the amount by which a marginal increase in leverage would increase the probability of a crisis is currently greater than usual (a larger value of $\gamma'_t(L_t)$). If $\gamma_t(L)$ is a strictly convex function, as assumed above, then one reason for the marginal crisis risk to be higher than normal could simply be that leverage L_t is currently higher than its normal level.

3 How Much of a Change in the Policy Framework is Necessary?

The optimal target criterion (2.14) implies that the central bank should be willing, at least to some extent, to trade off a greater degree of stability of conventional stabilization objectives — namely, price stability and output-gap stability — for the sake of greater stabilization of the marginal crisis risk. Under certain circumstances — specifically, when the current or anticipated near-term marginal crisis risk is unusually elevated — the target criterion implies that an ideal policy would tighten monetary conditions to the point that the price level and/or the output gap undershoot the levels that would otherwise have been considered desirable for these variables. (The target criterion specifies only that there should be undershooting of a certain linear combination of these two variables; but because of the Phillips-curve relation linking short-term variations in these two variables, it is likely that *both* variables would need to undershoot their target paths.) In this sense, the model implies that it is appropriate to use monetary policy to “lean against” a credit boom, even if this requires both inflation and the output gap to be below their medium-run target values for a time.

This statement requires some immediate qualifications. First, the analysis above is purely qualitative; considerable research is still needed to provide a sound empirical basis for a quantitative specification of crucial relationships such as the endogenous transition probability $\gamma_t(L)$ or the law of motion (2.3). Even the correct definition and measurement of the variable L_t in these relationships is far from obvious. I have referred to it as the degree of “leverage” for shorthand, but surely the risk of occurrence of a crisis depends on other aspects of the

balance sheets of financial institutions as well, such as maturity mismatch and the degree to which multiple institutions are exposed to the same (or highly correlated) risks. In all likelihood, a model that could be used for practical policymaking would have to replace the scalar variable L_t assumed in the simple exposition above by a *vector* of financial risk factors; the law of motion (2.3) would have to be replaced by a corresponding equation *system*. Research should probably focus first on identifying the risk factors that are most important in explaining variations in the endogenous transition probability γ_t ; the approach of Davig and Hakkio (2010) is an example of how this can be done, though they do not consider the predictive value of financial-sector balance-sheet variables. Once the key financial risk factors have been identified, attention can turn to the development of a structural econometric model of the evolution of those variables, with particular attention to the connection between the risk factors and other endogenous variables that are (directly or indirectly) influenced by the central bank's interest-rate policy. Only with a quantitative empirical model of this kind in hand will it be possible to say anything very specific about the way or degree to which it is appropriate to "lean against" a credit boom.

Indeed, the merely qualitative discussion above does not yet establish how large the variations in the final term on the left-hand side of (2.14) are likely to be, and hence it does not really allow one to say whether a conventional approach to policy (one simply focused on conformity with a criterion like (2.7)) might not still represent a fairly good approximation to optimal policy. It is certainly possible that, at least over certain periods of time, variations in the forecast of the marginal crisis risk would be fairly minimal, so that in practice a policy based on (2.14) would essentially amount to stabilization of a linear combination of the price level and the output gap. If we suppose that γ_t remains negligible for all values of L_t below some threshold, then $\gamma'_t(L_t)$ will also remain negligible for L_t below that threshold, so that the marginal crisis risk X_t should be minimal as well. If L_t is far enough below the critical threshold to have little probability of crossing it over the horizon that is relevant given the exponentially declining weights in the final term on the left-hand side of (2.14), then the third term in the target criterion might be small and stable for several meetings of the policy committee in a row. Under such circumstances, simple and conventional rules of thumb for monetary policy decisions would likely be adequate. However, even under circumstances of

this kind, the more general target criterion (2.14) does not give an *incorrect* signal about whether policy is on track; and reference to the more general target criterion would be necessary in order to answer whether a simpler rule of thumb continues to be adequate or not. It is plausible to suppose that under at least some circumstances, the additional correction for marginal crisis risk *would* be of quantitative significance. At any rate, further efforts to quantify the relations involved in such a calculation would seem appropriate, before dismissing out of hand the possibility that non-trivial adjustment of the acceptable paths for the price level and the output gap might be appropriate at some times.

Nor does the analysis offered here imply in any way that the conventional monetary policy should be assigned the primary responsibility for containing risks to financial stability, so that other regulatory and supervisory safeguards are unnecessary. To the contrary, because the analysis identifies reasons for a tension to exist between the conventional stabilization goals (represented by the first two terms in the loss function (2.4)) and the goal of reducing the distortions resulting from financial crises (over and above their consequences for the stability of inflation and the output gap, as represented by the third term in the loss function), it implies that the existence of additional policy instruments — that could ensure that significant variations in marginal crisis risk never occur, even when conventional interest-rate policy is used purely to minimize the variability of inflation and the output gap — should allow better outcomes on *both* dimensions. Hence the development of such tools, possibly including new instruments of “macroprudential policy” as discussed in Woodford (2011b), is highly desirable to extent that it proves to be practical. Nonetheless, it cannot be claimed that such tools already exist and have proven their effectiveness, so that there is no ground at present to dismiss the relevance of financial stability considerations for monetary policy deliberations.

In the case that additional macroprudential policy tools (such as variable capital requirements) exist, and can be flexibly adjusted in response to changing conditions, it would be necessary to extend the framework sketched above to allow the law of motion (2.3) and/or the function $\gamma_t(L)$ to be shifted by these instruments of policy. In that case, an optimal policy regime would involve optimal adjustment of both instruments in response to economic disturbances, and the conditions required for such a joint policy to be optimal could be

described by a pair of target criteria, one for each instrument, as in Woodford (2011b). But the target criterion for optimal monetary policy would continue to be of the form (2.14); the intended adjustment of the macroprudential instrument would simply be an additional factor to take into account in forming projections of the future paths of the price level, the output gap and the marginal crisis risk under alternative forward paths for monetary policy. Only under quite optimistic assumptions should one expect allowance for macroprudential policy to completely eliminate variations in projected marginal crisis risk, so that monetary policy decisions can be made without even considering this variable.

Yet while the analysis here suggests that financial stability conditions should be taken into account in monetary policy deliberations — and research on the kind of quantitative models needed to analyze this issue should probably be a large part of the agenda for central-bank research staffs in the near term — it hardly follows that the traditional goals of monetary stabilization policy should no longer be important, or even that inflation-forecast targeting should not still prove useful as an organizing framework for monetary policy deliberations. In the model sketched in the previous section, optimal policy can still be characterized by the fulfillment of a certain linear relationship between the projected paths of three variables, and so a forecast-targeting procedure provides a useful practical approach to the implementation of such a policy. The three variables that must be projected in such a procedure are a price index, a measure of the output gap, and the “marginal crisis risk”; the first two of these are the same two variables as are emphasized in traditional accounts of “flexible inflation targeting” (and in the procedures of forecast-targeting central banks such as the Bank of England, the Riksbank, and the Norges Bank).

It is quite possible that, much of the time, the monetary policy deliberations required by the proposed criterion would be quite similar to those that would be undertaken by a “flexible inflation-targeting” central bank that neglected financial frictions altogether. Under the assumptions made above, γ'_t should be an increasing function of the degree of leverage in the financial sector; but the marginal crisis risk X_t need not be equally sensitive to variations in the degree of leverage at all times. If, as suggested above, the crisis risk is negligible for degrees of leverage below some critical level, then γ'_t (and hence the value of X_t) will be small and relatively insensitive to variations in L_t , as long as L_t remains well below that critical

level. At such times, verification that the economy's projected evolution conforms to the target criterion would reduce to a simple comparison of the projected paths of inflation and of the output gap. Under this conjecture, it would only be at times of particular financial imbalance that the additional term in the target criterion would become a quantitatively significant factor in policy deliberations. At such times, less weight would be put on the traditional stabilization goals, in order to reduce the risk of financial crisis; but this would be done in a way that remained completely continuous with the approach followed in more normal times.

Moreover, even during times when financial stability concerns cannot be set aside, the proposed target criterion would continue to imply a clear medium-run target for the inflation rate; for a commitment to ensure that the economy's projected evolution is consistent with the target criterion at all horizons would necessarily imply that departures of the inflation rate from its optimal long-run level (zero, in the simple model presented here) were purely transitory. In fact, the proposed target criterion implies something stronger: it implies that the forecast of the long-run price level should remain constant over time, so that there is a commitment to *eventually reverse* any changes in the price level that result from temporary departures of the inflation rate from its medium-run target. Thus the inflation target is not merely one that must be satisfied *prospectively*, in the absence of shocks that cannot be foreseen well in advance — it is one that the central bank should commit to fulfill *ex post*, and regardless of the disturbances that may occur, as long as the realized inflation rate is averaged over a sufficiently long horizon. Thus while the target criterion does not involve only the projected inflation rate, it remains true that the policy commitment involves a target for inflation in a sense which is not true for any other variable: the policy commitment about the long run is *only* a commitment about the (cumulative) rate of inflation over the long run, and it is only this variable for which there is a commitment to a definite numerical magnitude which is independent of subsequently realized disturbances.

It might be thought that the credibility of a central bank's commitment to its supposed inflation target would inevitably be weakened by a more complex target criterion of the kind proposed, which makes the justifiable short-run departures from the inflation target depend on additional variables beyond the output gap — and moreover, variables that will likely

present even more controversial measurement issues than those connected with the output gap, and that may depart from their normal levels for longer periods of time than has been typical for the output gap (at least, during the “Great Moderation” period in which inflation targeting became popular), so that one might fear that these additional terms in the target criterion could be used to justify departures from the supposed inflation target for years on end. This could easily be a problem, under a certain conception of inflation targeting, under which the target criterion is purely forward-looking and relates only to the rate of inflation that is projected for some fairly short horizon (two to three years in the future), with “bygones allowed to be bygones.”

The target criterion (2.14) proposed above is actually quite different, because of the *error correction* implied by a commitment to a price-level target rather than merely to a forward-looking inflation target. Under the criterion proposed above, any departure of the price level from its long-run target path that is justified by an assessment of variations in the projected marginal crisis risk will subsequently have to be reversed. Moreover, a given degree of elevation of the assessed marginal crisis risk will *not* justify ongoing inflation on one side of the long-run target rate, even if it persists for years: it would only justify a given size of one-time increase or decrease in the *price level*, and after this adjustment of the price level has occurred, the persistence of the abnormal crisis risk would require the price level to continue to grow at a rate equal to the long-run inflation target. Then, when the abnormal conditions eventually subside, even the deviation of the price level from its long-run target path would have to be reversed.

A credible commitment to a (modified) price-level target of the kind proposed here should also help to mitigate a common fear about proposals to use monetary policy to “lean against” credit booms. This the fear that, if the credit boom occurs during a period when interest rates are being kept low because inflation is undershooting its medium-run target level, tightening policy to restrain the growth of leverage in the financial sector would run the risk of tipping the economy into a deflationary spiral — a risk that central banks treat with exceeding caution, owing to the fear that policy easing will cease to be possible once deflationary expectations set in, owing to the zero lower bound on nominal interest rates. (This fear was surely a critical factor in the Fed’s decision during the mid-2000s not to

raise interest rates more quickly, despite warnings of a housing “bubble.”) Such a disaster scenario has some plausibility in the case of a central bank with a forward-looking inflation target, for once deflationary expectations result in the zero bound being reached, and as a consequence in the generation of deflation greater than the central bank would wish, the nature of the target that the central bank would reasonably be expected to pursue later — refusing to allow any excess inflation, even to offset unwanted past deflation — would itself tend to generate a “vicious circle” in which deflation is simultaneously a consequence of and a justification for deflationary expectations (Krugman, 1998; Eggertsson and Woodford, 2003). But there is much less ground for such fears under a policy commitment of the kind proposed here. With a commitment to an invariant long-run price level target, any period in which the price level falls below its target path should immediately create an anticipation of a future period of higher-than-average inflation to return to the target path. Allowing inflation below the target rate in order to lean against a credit bubble should not create expectations of continuing deflation, because such episodes would predictably be followed periods of corrective inflation. And even if the zero bound were to bind temporarily, the amount of unwanted deflation that would result should be modest, because the farther the price level falls below the target path, the greater the amount of “catch-up” inflation that should be expected in the future.

I therefore conclude that inflation targeting frameworks can and should be adapted to take account of the possibility of intermittent disruptions of financial intermediation of the kind experienced in 2007-2009. Taking this challenge seriously will require a new research program, in order to put quantitative flesh on the stylized model sketched above. But it should not, in my view, require substantial modification of the fundamental structure of forecast targeting as a framework for making decisions about interest-rate policy, nor abandonment of central banks’ commitment to a quantitative definition of medium-run price stability. These important innovations of the past twenty years are likely to remain highly useful despite the additional challenges that must now be faced. To the extent that fundamental reconsideration is needed, it should be of the interpretation of inflation targeting as an approach that focuses purely on the prospective rate of inflation a few years in the future (rather than on the cumulative realized increase in prices), and on the length of the horizon

over which inflation should be expected to return to the target rate, rather than on the criteria that should determine how large a transitory departure from price stability is justified. These are dimensions on which the theoretical literature on ideal inflation targeting regimes had already argued that actual inflation-targeting procedures could be improved, even before the global financial crisis and even abstracting from concerns for financial stability (e.g., Woodford, 2007). Both the possibility of sometimes hitting the zero lower bound and the possibility of sometimes needing to use interest-rate policy to restrain the growth of risks to financial stability make these reforms of inflation-targeting practice all the more urgent. But the reforms that are needed are a natural extension of the logic of inflation-forecast targeting, rather than a repudiation of its central aims.

References

- [1] Bernanke, Ben S., Thomas Laubach, Frederic S. Mishkin, and Adam S. Posen, *Inflation Targeting: Lessons from the International Experience*, Princeton: Princeton University Press, 1999.
- [2] Borio, Claudio, and Mathias Drehmann, “Assessing the Risk of Banking Crises – Revisited,” *BIS Quarterly Review*, March 2009, pp. 29-46.
- [3] Clarida, Richard, Jordi Gali, and Mark Gertler, “The Science of Monetary Policy: A New Keynesian Perspective,” *Journal of Economic Literature* 37: 1661-1707 (1999).
- [4] Cúrdia, Vasco, and Michael Woodford, “Credit Frictions and Optimal Monetary Policy,” working paper, Federal Reserve Bank of New York, August 2009.
- [5] Cúrdia, Vasco, and Michael Woodford, “The Central-Bank Balance Sheet as an Instrument of Monetary Policy,” *Journal of Monetary Economics* 58: 54-79 (2011).
- [6] Davig, Troy, and Craig Hakkio, “What Is the Effect of Financial Stress on Economic Activity?” Federal Reserve Bank of Kansas City *Economic Review*, Second Quarter 2010, pp. 35-62.
- [7] DeGrauwe, Paul, “There is More to Central Banking than Inflation Targeting,” www.voxEU.com, posted November 14, 2007.
- [8] Del Negro, Marco, Andrea Ferrero, Gauti Eggertsson, and Nobuhiro Kiyotaki, “The Great Escape? A Quantitative Evaluation of the Fed’s Non-Standard Policies,” working paper, Federal Reserve Bank of New York, March 2010.
- [9] Eggertsson, Gauti B., and Michael Woodford, “The Zero Bound on Interest Rates and Optimal Monetary Policy,” *Brookings Papers on Economic Activity* 2003:1, pp. 139-211.
- [10] Gertler, Mark, and Peter Karadi, “A Model of Unconventional Monetary Policy,” *Journal of Monetary Economics* 58: 17-34 (2011).

- [11] Giavazzi, Francesco, and Alberto Giovannini, “The Low-Interest-Rate Trap,” www.voxEU.com, posted July 19, 2010.
- [12] Hakkio, Craig, and William Keeton, “Financial Stress: What Is It, How Can It Be Measured, and Why Does It Matter?” Federal Reserve Bank of Kansas City *Economic Review*, Second Quarter 2009, pp. 5-50.
- [13] King, Mervyn A., “Changes in U.K. Monetary Policy: Rules and Discretion in Practice,” *Journal of Monetary Economics* 39: 81-97 (1997).
- [14] Krugman, Paul, “It’s Baaack! Japan’s Slump and the Return of the Liquidity Trap,” *Brookings Papers on Economic Activity* 1998:2, pp. 137-187.
- [15] Leijonhufvud, Axel, “Central Banking Doctrine in Light of the Crisis,” www.voxEU.org, posted May 13, 2008.
- [16] Svensson, Lars E.O., “Monetary Policy with Judgment: Forecast Targeting,” *International Journal of Central Banking* 1: 1-54 (2005).
- [17] Svensson, Lars E.O., “Inflation Targeting,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, vol. 3B, Amsterdam: Elsevier Press, 2011.
- [18] 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.
- [19] Woodford, Michael, “Forecast Targeting as a Monetary Policy Strategy: Policy Rules in Practice,” NBER Working Paper no. 13716, December 2007.
- [20] Woodford, Michael, “Financial Intermediation and Macroeconomic Analysis,” *Journal of Economic Perspectives* 24(4): 21-44 (2010).
- [21] Woodford, Michael, “Optimal Monetary Stabilization Policy,” in B.M. Friedman and M. Woodford, eds., *Handbook of Monetary Economics*, vol. 3B, Amsterdam: Elsevier Press, 2011a.

[22] Woodford, Michael, “Monetary Policy and Financial Stability,” a presentation at the NBER Summer Institute, Cambridge, Massachusetts, July 15, 2011b.

A APPENDIX

Here I explain the derivation of the first-order conditions characterizing the optimal policy commitment. Let us recall that the planner’s problem is to choose state-contingent paths for the variables $\{\pi_t, y_t, L_t, \Omega_t\}$ consistent with constraints (2.2) – (2.3) and the transition equation for the regimes so as to minimize (2.4).

This problem can be expressed in a recursive form, if we suppose that in each period t , the planner observes the current values of the exogenous disturbances (including whether a regime transition has occurred), and conditional on these, chooses values of y_t and L_t and a state-contingent commitment specifying a target for π_{t+1} in each of the possible states that may be reached in the following period. These choices must be consistent with equations (2.2) – (2.3) and with the state-contingent target for π_t chosen in the previous period. Under this specification of the planner’s strategy set each period, the constraints are all “backward looking,” and it is possible to solve for an optimal policy commitment (equivalent to optimal choice of an infinite-horizon state-contingent plan at some initial date) using dynamic-programming methods.¹⁵

Let $V_t(\pi_t, L_{t-1}; \Omega_t)$ be the minimum attainable value for the *continuation loss*

$$\frac{1}{2} E_t \sum_{\tau=t}^{\infty} \beta^{\tau-t} [\pi_{\tau}^2 + \lambda_y y_{\tau}^2 + \lambda_{\Omega} \Omega_{\tau}^2]$$

looking forward from period t , conditional upon the current realization of the exogenous disturbances, the inflation target π_t previously committed to for this state, the lagged level of leverage L_{t-1} , and the current regime Ω_t . (The arguments indicating the current realization of the exogenous disturbances have been suppressed, and are instead represented by the time subscript on the function $V_t(\cdot)$.) In such a state, values y_t, L_t and the state-contingent commitment $\pi_{t+1}(\cdot)$ are chosen to minimize

$$\frac{1}{2} [\pi_t^2 + \lambda_y y_t^2 + \lambda_{\Omega} \Omega_t^2] + \beta E_t [V_{t+1}(\pi_{t+1}, L_t; \Omega_{t+1})] \quad (\text{A.1})$$

subject to the inflation pre-commitment π_t , constraints (2.2) – (2.3), and the transition law for the regime Ω_{t+1} . The minimized value of (A.1) is then the value of $V_t(\pi_t, L_{t-1}; \Omega_t)$.

¹⁵See Woodford (2011a) for further discussion and illustration of methods that can be used to characterize optimal policy commitments in models of this kind.

In this recursive version of the problem, the values of π_t and Ω_t are given, so that one can equivalently state the problem as one of minimizing

$$\frac{1}{2} \lambda_y y_t^2 + \beta \mathbb{E}[V_{t+1}(\pi_{t+1}, L_t; \Omega_{t+1}) | h_t]$$

subject to the above constraints, where h_t is the complete history of realization of shocks through period t (so that the previous notation $E_t[\cdot]$ can here equivalently be written as $E[\cdot | h_t]$). One can also eliminate y_t as a choice variable by using (2.3) to substitute for y_t as a function of the path of leverage and of exogenous disturbances. The problem can then alternatively be stated as the choice of L_t and the state-contingent commitment $\pi_{t+1}(\cdot)$ to minimize

$$\begin{aligned} \frac{\lambda_y}{2\xi^2} (L_t - \rho L_{t-1} - v_t)^2 &+ \beta \Gamma_t(L_t; \Omega_t) \mathbb{E}[V_{t+1}(\pi_{t+1}, L_t; \bar{\Omega}) | h_t, \Omega_{t+1} = \bar{\Omega}] \\ &+ \beta (1 - \Gamma_t(L_t; \Omega_t)) \mathbb{E}[V_{t+1}(\pi_{t+1}, L_t; \underline{\Omega}) | h_t, \Omega_{t+1} = \underline{\Omega}] \end{aligned}$$

subject to the inflation pre-commitment π_t and the constraint (2.2), where $\Gamma_t(L_t; \Omega_t)$ is the conditional probability of the crisis state occurring in period $t + 1$.

Differentiating the Lagrangian for this problem with respect to L_t , one obtains the first-order necessary condition

$$z_t + \beta \xi \mathbb{E}_t[V_{L,t+1}(\pi_{t+1}, L_t; \Omega_{t+1})] + \beta \xi \Gamma_{L,t}(L_t; \Omega_t) \Delta V_{t+1|t} = 0, \quad (\text{A.2})$$

where

$$z_t \equiv \lambda_y x_t - \kappa_y \varphi_t, \quad (\text{A.3})$$

φ_t is a Lagrange multiplier associated with the constraint (2.2), and

$$\Delta V_{t+1|t} \equiv \mathbb{E}[V_{t+1} | h_t, \Omega_{t+1} = \bar{\Omega}] - \mathbb{E}[V_{t+1} | h_t, \Omega_{t+1} = \underline{\Omega}].$$

Similarly, differentiating the Lagrangian with respect to the value of π_{t+1} in any of the possible states of the world in period $t + 1$, one obtains a first-order necessary condition

$$V_{\pi,t+1}(\pi_{t+1}, L_t; \Omega_{t+1}) - \varphi_t = 0 \quad (\text{A.4})$$

for each possible state in period $t + 1$. Note that φ_t depends only on the history h_t (as there is only one constraint (2.2) corresponding to each possible history up until period t), as does the variable $V_{t+1|t}$.

Conditions (A.2)–(A.4) are also first-order conditions for a solution to the problem of minimizing (A.1). Applying the envelope theorem to that problem (which defines the value function $V_t(\cdot)$), we can evaluate the partial derivatives of the value function as

$$V_{\pi,t}(\pi_t, L_{t-1}; \Omega_t) = \pi_t + \varphi_t, \quad (\text{A.5})$$

$$V_{L,t}(\pi_t, L_{t-1}; \Omega_t) = -\rho\xi^{-1} z_t, \quad (\text{A.6})$$

where z_t is again defined as in (A.3).

Using (A.5) to substitute for V_π in (A.4), we can alternatively write this first-order condition in the form (2.6) given in the text. Similarly, using (A.6) to substitute for V_L in (A.2), we can alternatively write this first-order condition in the form (2.9) given in the text, where

$$X_t \equiv \Gamma_{L,t}(L_t; \Omega_t) \Delta V_{t+1|t}. \quad (\text{A.7})$$

Thus we obtain first-order conditions (2.6) and (2.9) as stated in the text. Finally, the assumption made in the text about the form of the transition probabilities between regimes implies that

$$\begin{aligned} \Gamma_{L,t}(L_t; \underline{\Omega}) &= \gamma'_t(L_t), \\ \Gamma_{L,t}(L_t; \bar{\Omega}) &= 0. \end{aligned}$$

Substituting this for $\Gamma_{L,t}$ in (A.7), we obtain the definition of the “marginal crisis risk” variable X_t given in the text.