Monetary Policy, Asset Prices, and Misspecification

Robert Tetlow

Introduction

The period from 1995Q1 to 2000Q2 was an unusual one for the US economy. Productivity growth, which had averaged 1 1/4 per cent per year over the previous 20 years, climbed by more than a percentage point. Over the same boom period, the federal funds rate was remarkably stable—perhaps in response to core inflation rates that mostly fell. From 1952 to 1994, stock market capitalization fluctuated between 30 and 100 per cent of nominal GDP. From there, it rocketed to a peak of 188 per cent of GDP in 2000Q1. Over the next two years, however, the stock market retraced nearly all of its post-1994 gains, and the economy slid into recession. Two questions immediately arise. The first concerns the role of the apparent stock market bubble in bringing about the recession. The second, following from the first, is whether there was more that the Fed could have done to tame the bubble and avoid the recession.

That there is some likelihood the stock market bust played a role in the recession is demonstrated in Figure 1. It shows the ratio of stock market wealth, as well as business expenditures on high-tech equipment and software (E&S), both as a share of nominal GDP.[1] The shaded bars are the

1. Stock market wealth comes from the Flow of Funds Accounts. It is approximately equal to the market capitalization of the Wilshire 5000 stock index. Very much the same impression could be drawn from a graph of the raw data (that is, without scaling by nominal GDP) or by redefining the numerator to include broader categories of business fixed-investment expenditures.

Figure 1
Equipment and software expenditures and stock market capitalization
(share of nominal GDP)

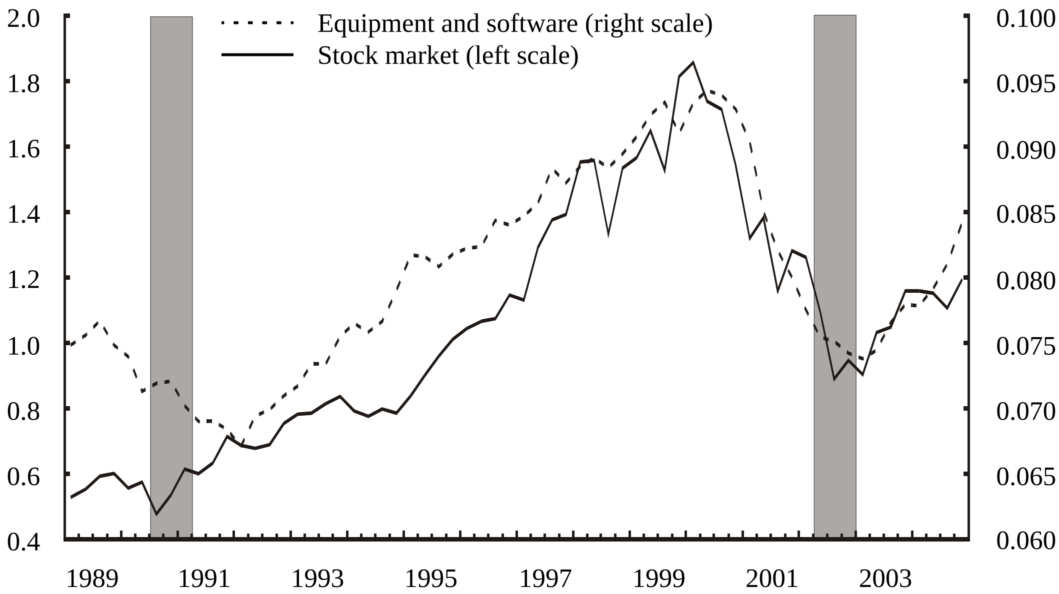

National Bureau of Economic Research recession periods. Three salient facts can be drawn from this figure.

First, clearly, both investment expenditures and stock market wealth increased dramatically through the latter half of the 1990s, before falling back sharply.[2] Second, the decline in the stock market preceded the decline in investment. And third, unlike in the 1991 recession (and indeed unlike most recessions), investment led the business cycle instead of trailing it.

Bernanke and Gertler (BG, 1999) argue that the quiescence of monetary policy was the correct response; that monetary policy should respond only to the *projected* effects of stock market movements on inflation and perhaps output, but not to perceived stock market bubbles per se. To central bankers, this advice seems sound: Asset prices appear to be too untrustworthy to be

2. Just to provide a longer-term perspective on the late 1990s, the ratio of nominal E&S expenditures to nominal GDP broke its historical record of 8.54 per cent (set in 1979) in 1996, and continued to climb from there. Our series begin in 1960. Stock market wealth broke its record share of nominal GDP—of 1.00, originally set in 1968Q4—at the end of 1995, and peaked at 1.85 in 2000Q1. The ratio is available dating back to 1947.

responded to directly; they give too many false signals, and too little is known about their determinants in real time.[3]

And yet the logic from control theory is also straightforward and points in the opposite direction: If asset prices (or financial wealth) are state variables in a macroeconomic system, they should be responded to like any other state variable. The fact that they are measured with error only means that care should be taken to filter the information correctly. The uncertainty inherent in the measurement of asset prices and in the origins of shocks to asset prices may result in attenuation of the response to asset-price movements, but it will not be generally optimal to fail to respond to such movements altogether. Cecchetti, Genberg, Lipsky, and Wadhwani (CGLW, 2000) and Mussa (2003) formulate arguments in favour of leaning against asset-price fluctuations on largely these grounds.

BG (1999) and CGLW couch their arguments in the language of inflation targeting, a sensible approach given the rising popularity of inflation targeting among central banks. At the same time, however, the recent experience in the United States should give advocates of inflation targeting some pause. If it is true that the bursting of the stock market bubble in 2000 was the proximate cause of the recession of 2001, and if the Fed can be described as having followed a policy of inflation targeting—keeping inflation on track but not directly responding to escalating equity prices—then either the recession was the best of all possible worlds, or inflation targeting alone is not a sufficient policy.

It seems clear that asset markets are prone to non-fundamental outcomes—that is, to bubbles or fads. Such phenomena are non-linear in nature in that they sometimes build up in a continuous fashion but revert to fundamentals in a discrete manner. From a technical standpoint, this raises problems, because efficient tools for computing optimal Taylor-type-rule coefficients rely on linearity of the model with Gaussian shocks. It is not a straightforward task to forecast their implications for future output and inflation and devise the appropriate response. Moreover, from a behavioural standpoint, one might argue that it is unreasonable to expect agents to form rational expectations of the effects of phenomena that are observed only once every 20 years.

3. According to Bullard and Schaling (2002), reacting to asset prices—or more specifically in this case, equity prices—can also increase the range of model instability. That is, there are more combinations of structural (non-policy) coefficients and policy-rule parameters for which the model is unstable when the authority reacts to equity prices than when it does not.

In this paper, we use a variant of the Bernanke-Gertler-Gilchrist (BGG, 1999) model to reassess the case for responding to bubbles. We embellish the version of the model used by BG (1999), adding structure to enhance dynamic propagation and make the model more consistent with the data. With this model, we contribute two things. First, we compute the (approximately) optimal weight on a stock market term of the outcome-based and inflation-forecast-based policy rules, in the presence of a full set of stochastic shocks. The use of forward-looking rules is important here, because current fluctuations in stock market valuation have effects on output and inflation over extended periods of time. The reliance on such rules is very much in line with the policy advice of BG (1999). This gives an upper bound, for this model and calibration, of the good the Fed can do in responding to asset-price developments. Second, we drop the full-information assumption and instead assume that the Fed has doubts about its model of the economy. We model the authority as believing that it is controlling a different economy than in fact is the case. Note that this is related to, but different from, exercises where the authority is assumed to be unsure of the drivers of asset prices. Both of these contributions are novel to this paper.

Our application is for the US economy in the presence of stock market bubbles. That said, consistent with the arguments of Batini and Nelson (2000), among others, we believe that the same logic can be applied to exchange rate, commodity price, and, with some modification, land price bubbles.

The rest of the paper proceeds as follows. In section 1, we introduce our model: a variant of the same BGG model used by both BG (1999, 2001) and CGLW.[4] The model differs from its predecessors in the allowance of richer dynamics and a more complete set of stochastic shocks. In the same section, we describe our bubble process and the calibration of the models. Section 2 computes the optimal outcome-based (or, equivalently, Taylor-type) rules, with and without a term for equity prices, and with and without knowledge of the model. The final section offers concluding remarks.

1 The Building Blocks

1.1 The model

In most respects, the basic BGG model is a straightforward New Keynesian dynamic general-equilibrium model but adds a financial accelerator to the model's propagation mechanism. Firms finance capital spending with a

4. Carlstrom and Fuerst (1997) is another creditable example of a financial accelerator model. We use BGG to maximize comparability with the earlier literature in this subject area.

mixture of external and borrowed funds. Financial market frictions imply a wedge between the cost of internal and external finance. The cost of external finance is a decreasing function of the net worth of the firm, because higher net wealth implies a higher collateralized value of the firm. This means that shocks—including "non-fundamental" ones—that raise the value of the firm relax a constraint on capital accumulation and induce investment. This is a useful feature of the model, since it arguably captures many of the stories that go along with speculative booms and busts. In the late 1990s in the United States, for example, the financial press was replete with stories characterizing the unusual ease with which firms could raise funds.[5] Firms are owned by entrepreneurs who plan over finite horizons to purchase physical capital, rent labour, and produce output. Households choose work, consumption, and savings over an infinite horizon. The government operates monetary policy through the calibration and application of a Taylor-type interest rate feedback rule.[6]

The basic BGG model is embellished in several ways. Like BG (1999, 2001), we use a "hybrid" Phillips curve specification that allows for a lagged term in inflation in addition to the forward-looking term that is familiar from the canonical New Keynesian model.[7] However, we also add or adjust three features of the BG (1999, 2001) implementation. First, we allow adjustment costs to investment in the form of Casares and McCallum (2000). They specify adjustment costs of the form $C(i_t) = \psi i_t^{\eta}$ with $\psi > 0, \eta > 1$. A value of $\eta = 2$ would be garden-variety quadratic adjustment costs; we adopt their mid-range case from their Table 5, p. 26, of $\eta = 2.5$ along with $\psi = 2000$. Second, we allow for external habit persistence in consumption, as in Abel (1990). Whereas the canonical New Keynesian consumption function models consumption as purely forward looking, habit persistence allows a lag of consumption to enter the consumer's decision rule.[8] Each of these first two alternations is intended to impart persistence into the model and thereby create more realistic model dynamics. The greater persistence, on the other hand, should make the welfare consequences of policy mistakes larger than would otherwise be the case. Our third change concerns the channels through which a stock market bubble may operate. Whereas BG

5. See, e.g., Kaplan (2003), who presents interesting numbers on initial public offerings.
6. The use of feedback rules in place of monetary targeting is quickly becoming standard. Nonetheless, one could recast the policy decisions in this paper in terms of money, provided one were to assume a stable money-demand function. However, the comparability of this work with previous research would be impaired by such a step.
7. See, e.g., Woodford (2003), chapter 3. Amato and Laubach (2002) show how a portion of rule-of-thumb price-setters can provide a microfoundation for the hybrid Phillips curve.
8. It also allows a second lead, date $t + 2$. In any case, for plausible calibrations of habits, the degree of persistence in consumption imparted by this formulation is not all that large.

(1999, 2001) allowed only consumption to be affected by stock market bubbles, we also allow investment to be misallocated because of bubbles. We do this by allowing investment decisions to respond to observed stock market values rather than the "fundamental *Q*." This takes on board the observation of Dupor (2001), who argues that inefficient shocks to firms' investment schedules may render a case for activist monetary policy responses to bubbles. The evidence, shown in Figure 1, would also suggest that there is a case on empirical grounds for this extension.

1.2 The bubble process

In the rest of this section, we explain the addition of exogenous stock market bubbles. Our formulation is almost identical to BG (1999, 2001) and CGLW (2000).

Assume that the market price of capital, S, varies from the "fundamental" price, Q, owing to bubbles or fads, so that the existence of the bubble can be summarized by the difference between the two: $U_t = S_t - Q_t$. If a bubble exists, it persists with probability, p, and conditional on its persistence, it grows at rate a/p:

$$U_{t+1} = \frac{a}{p} U_t R^q_{t+1} \qquad \text{if the bubble persists}$$

$$U_t = 0 \qquad \text{otherwise},$$

where R^q_{t+1} is the relative stochastic discount rate at which dividends are discounted, and a is the expected growth rate of the bubble with $p < a < 1$. With this restriction, the unconditional expectation of the bubble in period $t+1$ is $a < 1$, while the expectation conditional on the bubble not bursting is $a/p > 1$. In other words, if the bubble doesn't burst, it grows. In calibrating the bubble process, in most instances we assume $a = 0.99$ and $p = 0.5$, the same assumptions as those of Bernanke and Gertler.[9] This means that once a bubble is initiated, it will (almost) double if it does not burst. To ensure that a single outsized event does not dominate results, we further assume that a bubble never lasts more than five periods.[10]

The bubble process has two noteworthy features. First, it is a (virtually) rational bubble in that the expected rate of return on holding capital,

9. Were we to assume that $a = 1$, we would be assuming a rational bubble. In most of what follows, however, we assume $a = 0.99$ to ensure that the model is stationary, while staying arbitrarily close to a rational bubble.

10. The odds of a bubble lasting longer than five periods is only one in thirty-two, in any case.

conditional on a bubble, is the same as the opportunity cost of funds. Thus, the persistence of the bubble does *not* depend on "irrational exuberance." Second, the bubble is exogenous. As in most of the literature on this subject, we do not attempt to explain how bubbles originate. Similarly, we allow no channel for monetary policy to affect the bubble directly. There are advantages and disadvantages to the bubble process we use. The disadvantages are that no theory is adopted to explain why bubbles arise; and a potentially important, if obscure, channel whereby monetary policy can work—a channel from policy actions to private agents' beliefs—is omitted. The advantages of the process are that it is simple, transparent, and it does not depend on arbitrary assumptions regarding investor beliefs, other than the assumption that bubbles can exist in the first place. It has been used before, in BG (2001). Finally, there is reason to hope that by eschewing the modelling of a possibly controversial channel for policy to act on beliefs, the results derived here will be more broadly applicable than otherwise.

To fix ideas on how the bubble process works, Figure 2 shows a particular realization of a bubble. The solid line shows the bubble itself, which arrives in period t and happens to have a magnitude of unity. The dashed line shows the surprise to private agents owing to the bubble's initiation and continuation. As shown, the bubble lasts three periods before bursting in period $t+3$. At period $t+1$, agents aware of the existence of a bubble expect that with probability $1-p = 0.5$, it will burst and that with probability $p = 0.5$, it will continue. If it continues, it doubles in size. Thus, the expected rate of return on holding stock market assets in period $t+1$ is $R^u_{t+1} = E_t(R^s_{t+1} - R^k_{t+1}) = p \cdot 1 + (1-p) \cdot -1 =$, meaning that expected excess returns are zero: the bubble is a rational bubble. In period $t+1$, in fact, the bubble does not burst, a realization that engenders a surprise of 1. Now the agents face the same decision as in the previous period, except that the stakes have doubled. When the bubble continues in period $t+2$, the surprise is 2; when it finally bursts, the surprise is –4.

As mentioned above, the calibration we use has the initiation of a bubble governed by a Poisson distribution with arrival probability 0.02, or about once every 13 years. Since there was a stock market crash in October 1987 and another spread over 2001–02, this would appear to be about the right frequency. In our experiments, we simulate 5,000 periods, so that a bubble occurs, on average, 100 times in a run. Given the initiation of a bubble, the size of the bubble is determined by a standard mean-zero Gaussian distribution.

The bubble process, as just discussed, completes the model up to the policy rule. That said, the relative complexity of the BGG model, combined with space constraints, does not allow for a detailed discussion of the model. BG

Figure 2
Three-period bubble realization
($p = 0.5, U_t = 1$)

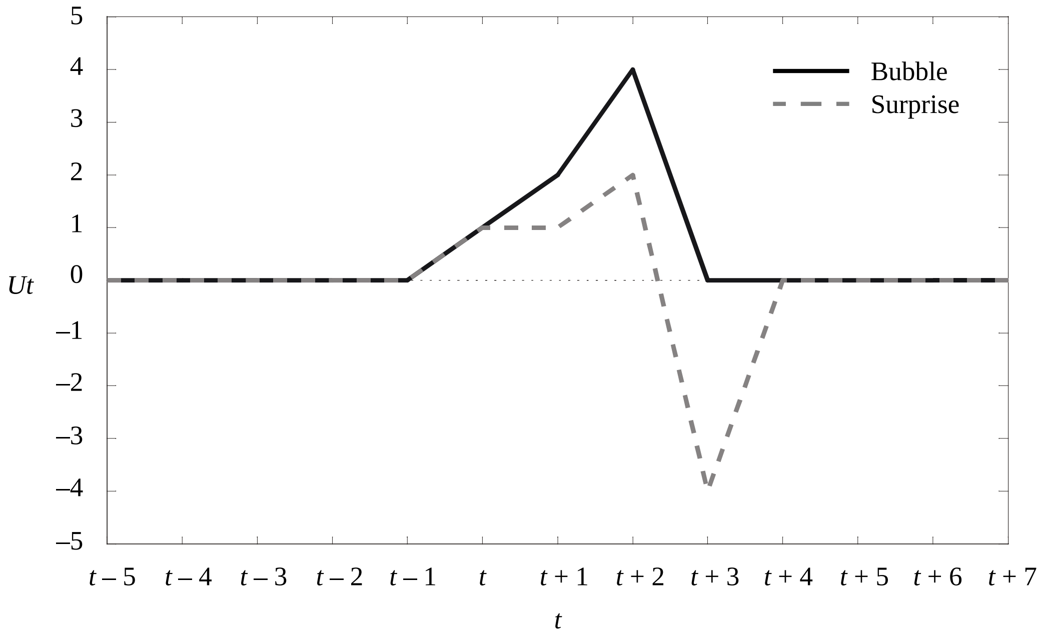

(1999) provide some discussion and BGG lay out the model in considerable detail. For those who are interested, the complete model is shown in Appendix 1. However, to give a bit of an idea of how the model works, Figure 3 shows the model's response to a one-time shock to trend total-factor productivity (the "*z* shock" in the left-hand column of charts) and to the initiation of a bubble (the "*u* shock"). The responses shown are conditional on two policies to be discussed later, one that responds solely on the forecast of inflation one quarter ahead—the dashed line—and another that feeds back on the change in the value of the stock market as well as inflation—the solid line. The first row of the figure shows the output-gap responses ("*ygap*"), the second row shows the inflation rate ("*inf*"), and the third shows the nominal funds rate ("*rn*").

We would argue that the model's responses look sensible; both shocks produce hump-shaped economic dynamics, as would, for example, a vector autoregression (VAR). That basic point aside, two interesting observations can be taken from this figure. First, the policy responses to both shocks differ substantially depending on whether or not policy responds directly to stock prices. This suggests that the policy design decision under study here is consequential. Second, the policy response to a productivity shock differs markedly from what is appropriate to a bubble shock. This ably sketches the dilemma faced by the Fed in the late 1990s: To the extent that the boom of

Figure 3
Model impulse responses to selected shocks

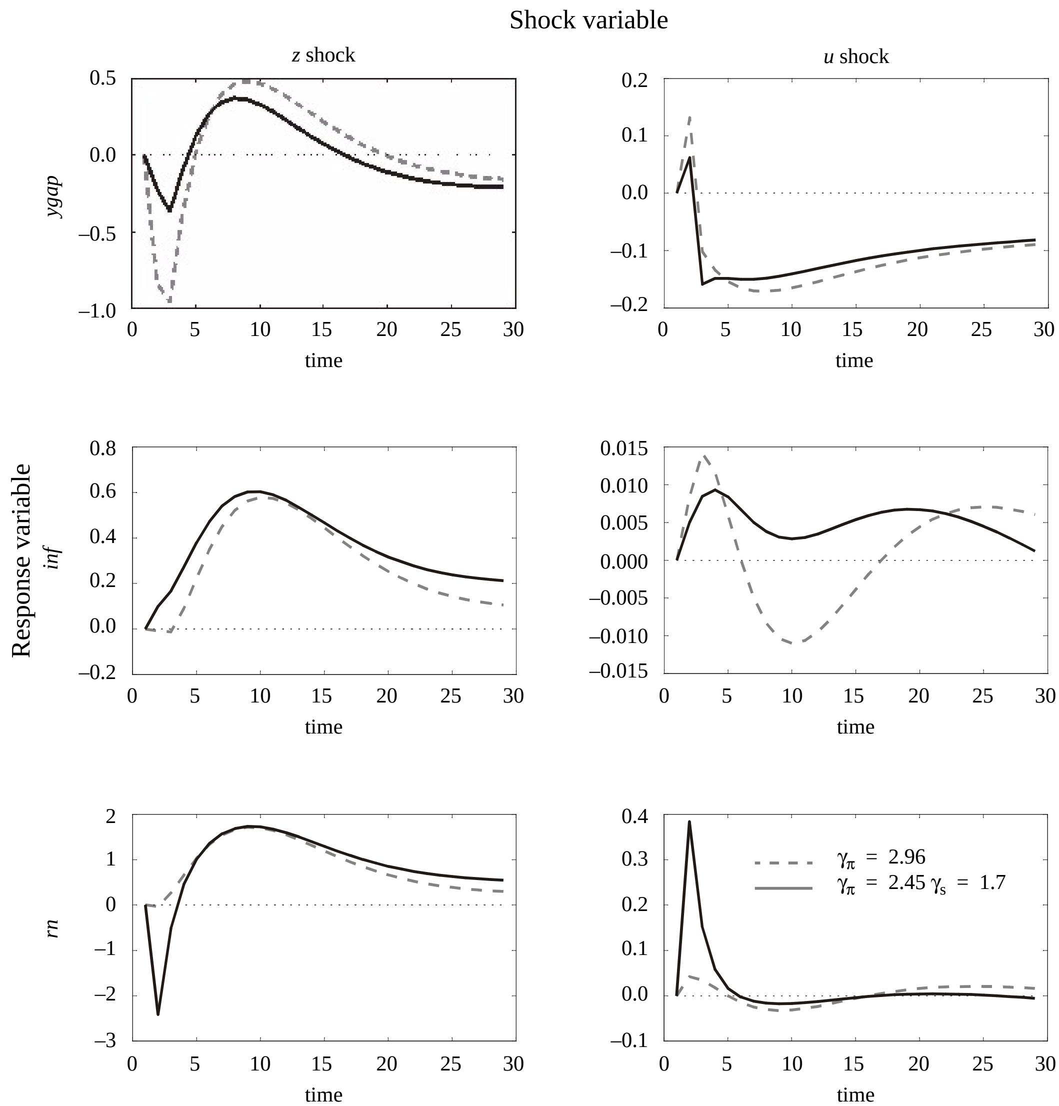

that period was "fundamental" in that it was being driven by productivity, the appropriate response of policy is accommodative. If, however, the boom is mostly a bubble phenomenon, the appropriate response is restrictive.

1.3 Certainty equivalent policy

In the certainty equivalent policy experiments we consider, the government is assumed to minimize a quadratic loss function as follows:

$$\underset{\langle\gamma_i\rangle}{MIN} E_0 \sum_{i=0}^{T} \lambda \tilde{y}^2_{t+i} + (1-\lambda)\pi^2_{t+i}, \qquad (1)$$

where $\tilde{y}_t$ is the output gap, $\pi_t = 400(P_t/P_{t-1} - 1)$ is the inflation rate, and γ_i is a vector of policy-rule parameters that will be explained presently.[11] Notice that no term appears for instrument smoothing; nor is there a term in some measure of the stock market itself. This means that in what follows, the efficacy of reacting directly to stock market developments is modelled as the means to an end and not as a goal of policy in itself. This formulation is in keeping with what is now the standard approach in the literature.

The target rate of inflation is taken to be a positive constant large enough to avoid the zero-bound problem on nominal interest rates and is normalized out of the equation. The minimization is subject to the rest of the model, the variance-covariance matrix of stochastic disturbances, and the form of the policy rule. As noted, policy is assumed to be governed by a Taylor-type rule:

$$R_t = r^* + \gamma_\pi E_t \pi_{t+i} + \gamma_y \tilde{y}_t + \gamma_s (s_t - s_{t-1}) \qquad i = 0 \text{ or } 1 . \qquad (2)$$

Several aspects of equation (2) are worth noting. First, the inflation term can appear as either a contemporaneous term or with a one-period lead. The former is a traditional outcome-based Taylor-type rule, while the latter has been dubbed by some as an inflation-forecast-based (IFB) rule. IFB rules are touted for their ability to encompass a great deal of information in a single object: the inflation forecast. The idea is that the entire model within which the rule is embedded is used to solve for the inflation forecast so that feeding back on the lead of inflation implicitly feeds back on all of the states that are relevant for inflation determination—including the output gap.[12] IFB rules have their detractors, however, mostly because of the presumed lack of robustness of such rules to model misspecification.[13] Second, equation (2) shows the stock price entering in first differences. Both advocates and

11. Variables in upper case should be understood to mean levels, while lower case designates 100 times the logarithm.

12. The earliest use of IFB rules is in the Bank of Canada's QPM model beginning in 1991; see, inter alia, Coletti et al. (1996) for a discussion of the model and the IFB rule therein. Since then, its popularity has grown. Svensson (2002) argues that IFB rules—like all "simple rules"—are less than completely efficient, since the way state variables are used in formulating the forecast in equation (2) is not the same as they would be used in the targeting regime he promotes. Finan and Tetlow (1999) show that simple outcome-based rules perform very close to the optimal rule in small models but somewhat less well in large-scale rational-expectations models.

13. Levin, Wieland, and Williams (2003) study the robustness of IFB rules, finding that they are (surprisingly) robust provided that the lead horizon on inflation is short, as it is here. Critics argue that the models studied by Levin, Wieland, and Williams are too similar to do justice to the issue of model uncertainty.

detractors of direct feedback on asset prices argue that central banks should not attempt to "prick" bubbles; rather, the most they should do is "lean against the wind" of asset-price changes. Formulating the stock price in log differences, as opposed to deviations from fundamental *levels*, is consistent with this interpretation of the role of policy in that it does not require the policy-maker to know the equilibrium level of stock prices. Third, equation (2) includes both a stock-price variable and an output-gap term in addition to the usual inflation variable. In fact, the primary cases we are interested in are the ones studied by BG, which involve the restrictions $i = 1$ and $\gamma_{\tilde{y}} = 0$ with comparisons of $\gamma_s = 0$ and $\gamma_s > 0$, that is, an inflation-forecast-targeting rule with or without feedback on the change in the stock price, but no output-gap variable.[14] This focus has the advantage of allowing a close comparison to the earlier results of BG as well as reducing the already significant computational cost of searching over optimal coefficients. That said, Cecchetti, Genberg, and Wadhwani (2003) speculate that the absence of feedback on the output gap in BG (2001) might be one reason why the BG conclusions differ from those of CGLW, and so we shall devote some space to this issue.

The generic experiments, as we call them, differ from BG in only small ways. Of course, the model differs in some ways. We also differ in that we consider a broader range of stochastic shocks to the model, adding shocks to tastes (consumption) and to government expenditures in addition to the productivity shocks and bubble shocks studied by BG.[15] Finally, we differ in the range of rules we permit in that we consider outcome-based rules. Outside of the generic experiments, however, we consider the issue of model uncertainty, doing so through the lens of robust control.

1.4 Robust-control policies

There are at least three different approaches to robust control. What they share is a focus on the distinction between *uncertainty* in the sense of

14. Two differences in our formulation, relative to BG, are that (i) we assume that the government reacts to the change in the stock price rather than to the gap between stock prices and steady-state stock prices; and (ii) we assume feedback on the contemporaneous change in stock prices, not lagged stock prices. The former assumption stems from our belief that stock market fundamentals are difficult to measure at any time. The latter assumption stems from our belief that actual stock prices are easy to measure in real time.

15. Specifically, in our base-case experiments, we assume a variance-covariance matrix of forcing shocks that is $\text{diag}\begin{bmatrix}1 & 1 & 1 & 0.1 & 4\end{bmatrix}$, where the first three shocks are to the Phillips curve, consumption, and government expenditures, respectively; the fourth shock is to trend total-factor productivity and the fifth shock is the bubble shock. Note that the variance of 4 for the bubble shock is only applicable when a bubble shock arrives. In some instances, we will allow the productivity and bubble shocks to covary.

Knight, which is non-parametric in nature, and the concept of *risk*, which can be taken as parametric. Risk is a statistical concept for which there exist straightforward techniques for managing. Uncertainty is more profound and arguably more plausible for the issue of asset-price bubbles, since the infrequency and unfamiliarity of bubbles militate against a reliable quantification by econometric methods, which typically require large samples to be efficacious. The approach to robust control we adopt is *structured Knightian uncertainty*, where the uncertainty is assumed to be located in one or more specific parameters of the model, but where the true values of these parameters are known only to be bounded between minimum and maximum conceivable values. Among the expositors of this approach to model uncertainty are von zur Muehlen (1982), Giannoni (2001, 2002), and Tetlow and von zur Muehlen (2004).[16] This particular variant of robust control is arguably the most intuitive and practical of the choices. To illustrate how structured model uncertainty is characterized, let us summarize our model in general state-space form by the following expression:

$$x_t = Bx_{t-1} + CR_t + \varepsilon_t, \tag{3}$$

where x is a vector of endogenous (state) variables, including $\tilde{y}$ and π, and R is the control variable, the same short-term interest rate in the policy-maker's reaction function. Structured model uncertainty posits that the policy-maker takes equation (3) to be the *reference model*, thinking that it is approximately correct, but harbouring uncertainty about some subset of the model's structural parameters, either B or C. Moreover, the policy-maker is assumed not to have a parametric estimate of this uncertainty—a standard error, or some such thing—but rather is more generally wary of errors of indeterminate origin or magnitude. This may arise either because of suspected misspecification—something that, unlike sampling error, does not lend itself to parametric estimates—or because the phenomenon of interest occurs too infrequently to expect parametric estimators to extract from the data. Either or both of these phenomena may be at work in current circumstances. Let us consider the misspecification of a single parameter within the matrix B, and let us call it b_{jk}. In the absence of a reliable statistical estimate of the error in b_{jk}, Gilboa and Schmeidler (1989) show that the policy-maker's problem naturally leads to a min-max solution wherein the policy-maker acts as though choosing a loss-minimizing policy

16. Among the other two notions of structured model uncertainty in the sense of Knight is unstructured model uncertainty, where the uncertainty is non-parametric and its location is unclear. See Hansen and Sargent (2005) and references therein. The third method differs from the other two in that the authority is assumed to choose a policy rule that maximizes the set of models for which the economy is stable. See, e.g., Onatski and Stock (2002) and Tetlow and von zur Muehlen (2001).

conditional on the reference model and subject to the loss-maximizing choice of b_{jk}, where:

$$b_{jk}^* = \underset{b}{\text{argmax}} \, [\underline{b}_{jk}, \bar{b}_{jk}] \qquad \text{s.t. (1) to (3)}, \tag{4}$$

where $\underline{b}_{jk}$ is the lower bound on possible values of b_{jk} as conceived by the policy-maker, and $\bar{b}_{jk}$ is the corresponding upper bound. The common metaphor is that in the absence of information on what value b_{jk} could take, the optimal strategy for the policy-maker is to protect against the worst-case outcome for the parameter, that is, to act as though there was an "evil agent" that chooses the worst possible value for b_{jk} within bounds. The policy-maker then acts as the leader in a Stackelberg game, choosing the best policy-rule parameters, $\gamma_i^*, i = \pi, \tilde{y}$, conditional on b_{jk}^*, the vigilent's choice at b_{jk}.

The main parameter of interest for our min-max problem will be a, the expected growth rate of bubbles, although we shall also investigate p, the survival probability. The likely size and growth of bubbles seem to have been in play in the late 1990s stock market bubble in the United States and consequently appear to be the obvious candidates for analysis.

Formally, the problem to be solved is:

$$\underset{\gamma_i}{\text{MIN}} \quad \underset{b_{jk} \in [\underline{b}_{jk} \; \bar{b}_{jk}]}{\text{MAX}} \quad E_{t-1} \sum_{m=0}^{T} \lambda \tilde{y}_{t+m}^2 + (1-\lambda)\pi_{t+m}^2$$

$$\text{s.t. (2) to (3)} \tag{5}$$

and subject to any coefficient restrictions on $\gamma_i, i = \tilde{y}, \pi, s$ as applicable for the problem at hand. The next section presents some results.

2 Results

In this section, we will present results from stochastic simulations of the model with optimization of policy-rule coefficients. The first subsection considers straightforward experiments involving the base-case calibration along with some sensitivity analysis. The second subsection considers the implications of possible model misspecification and the policy response to misspecification.

In all instances, simulations were conducted over 5,000 periods with a Poisson arrival rate of 0.02 for bubble shocks. Such an arrival rate is consistent with a bubble arising every 13 years on average, or about 100 times in each sample. By some arguments, a (negative) bubble in stock prices occurred in the United States in the mid-1970s, leading to a stock market crash in 1987, with another bubble and subsequent crash arising in 2000, so the chosen arrival rate seems reasonable. The economy we are studying contains non-linear and non-Gaussian elements. As such, it does not lend itself to solution using algebraic methods. Accordingly, a grid-search procedure was utilized to find the optimal parameterization of equation (2).

2.1 Basic results

In this subsection, we discuss results for experiments in which the standard quadratic loss function in equation (2) is minimized subject to the model, the variance-covariance matrix of forcing shocks, and the specification of either an outcome-based Taylor-type rule or an IFB rule. We begin with results from IFB rules, summarized in Table 1. The first column in the body of the table shows the weight on the (squared) output gap in the loss function. Three different sets of preferences are highlighted. The rest of the table shows the optimal coefficients for one-, two-, and three-parameter rules, and the corresponding losses. The losses have been normalized such that the loss under the rule feeding back on the inflation forecast and the change in asset prices is equal to unity. Except for the fact that the asset-price term appears in (log) changes, this is the form of rule upon which BG focused. We refer to this scenario as the base case and to the performance of the economy under these circumstances as the base-case loss.

Thus, line 1 shows that the base-case rule bears a feedback coefficient on future inflation of 2.45, a coefficient on the change in the stock market of 1.70, and a loss of unity. The second row shows that the optimal one-parameter rule—that is, the optimal rule subject to the restriction of no (direct) feedback on the stock market—carries a coefficient on future inflation of 2.96, a little higher than the coefficient in line 1, but not substantially so. More important, the loss column shows that the incremental loss from restricting oneself to responding directly to inflation alone is about 11 per cent of the base-case loss. While this is not trivial, it would be difficult to argue that a loss of this measure is a major concern.

Note that the impulse responses shown previously in Figure 3 are conditioned on the policy rules described on lines 1 and 2. Some perspective on these rules is provided by Figure 4, which shows the stability mapping for the model. The horizontal axis of the figure plots the feedback coefficient on the lead of inflation, γ_{π}, while the vertical axis maps the coefficient on the change in the stock market, γ_s. The dark regions to the northeast and southwest, marked “stable,” represent policy-rule parameterizations that

Table 1
Optimal coefficients and performance of inflation-forecast-based rules
(base-case calibration)

row	loss function	rule coefficients			loss
	λ	γ_π	γ_s	$\gamma_{\tilde{y}}$	L
(1)	0.5	2.45	1.70	–	1
(2)		2.96	–	–	1.11
(3)		2.60	1.70	1.80	0.93
(4)		11.44	–	35.36	0.95
(5)	0.9	2.10	1.43	–	1
(6)		2.54	–	–	1.14
(7)	0.1	8.49	5.79		1
(8)		6.74	–	–	1.14

ensure saddle-point stability of the model. The lighter region running from the northwest to the southeast represents the parameterizations of the policy rule that permit indeterminacy in solutions. Finally, the region in the southeast corner is the unstable region. The vertical line at unity for γ_π is the naïve borderline for stability for models, marking the satisfaction of the so-called Taylor principle. The figure shows that while $\gamma_\pi > 1$ is a useful feature in that there is a large region of stability that satisfies this condition, it is neither necessary nor sufficient for stability in this model.

Also shown in the figure are the positions of the two rules from lines 1 and 2 of Table 1. The figure shows that these policies are fairly close to regions of indeterminacy. What this means is that in principle small misperceptions in the structure of the true model that could result in perturbations of the optimal coefficients of the rules considered here could put the (actual) economy in the indeterminate region. The resulting dynamics of the economy would be subject to drifting inflation governed by random beliefs—sunspots—that by definition are difficult to describe a priori. As Lubik and Schorfheide (2004) have emphasized, the observational implications of sunspot equilibria in monetary models include greater persistence and larger (or more) shocks than would otherwise be the case. Tetlow and von zur Muehlen (2005) describe tools for avoiding such an outcome.

Returning to Table 1, line 3 shows the optimal coefficients for the three-parameter rule. Since it allows feedback on a broader set of variables, this rule must outperform the base-case rule. In this instance, however, the improvement is remarkably small. Just as important, the feedback on the stock market does not differ from that of line 1. Evidently, output response and stock market response are not strong substitutes in this model.

Figure 4
Model stability mapping

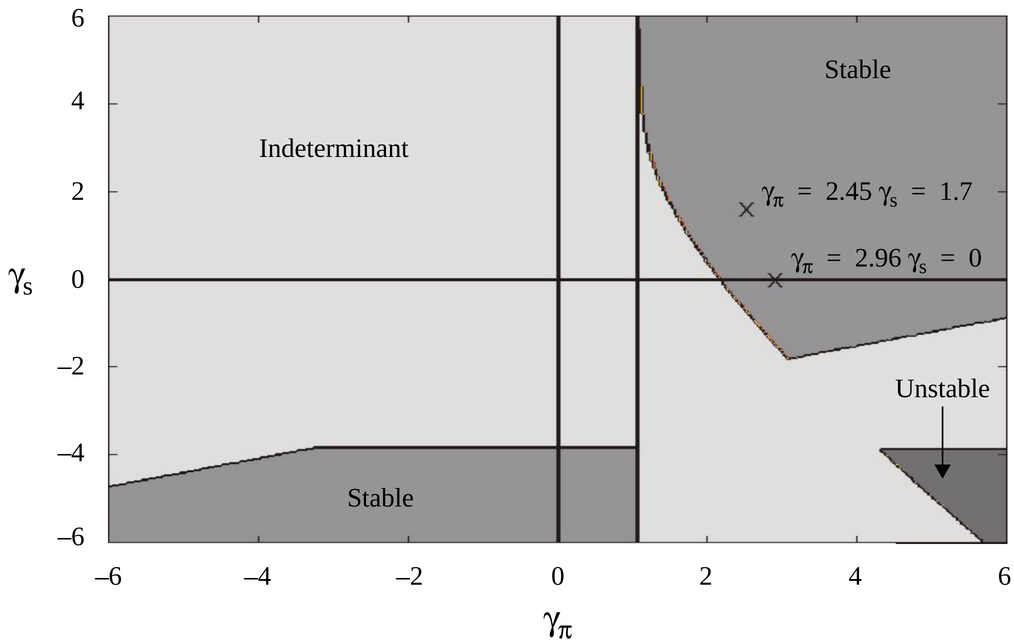

Line 4 shows the optimal coefficients for a rule that feeds back on (future) inflation and the output gap, but not on stock prices. Note that in this case the feedback coefficients on both inflation and the output gap are quite large. In fact, the contours of the loss surface are such that while these coefficients are optimal for the problem at hand, rules with smaller feedback coefficients exist, not unlike those on line 1 that perform close to as well as the rule shown.[17] When we compare the last two columns of these two rows, it is evident that feedback on stock prices is not crucial to macroeconomic performance once feedback on the output gap is permitted. Were the central bank to eschew feedback on the output gap, say, on the grounds that the gap cannot be measured accurately, the comparison of the last column of lines 1 and 2 would be germane: there, it is shown that the incremental gain, while positive, is small.[18]

17. The topology of the (negative of the) loss surface looks like a ridge line in $\gamma_\pi - \gamma_{\tilde{y}}$ space, which means that the optimal coefficients rise together with the height of the ridge, changing very little. This feature appears to be robust to changes in model parameterization.
18. This finding, which is also true for the version of the BGG model used by BG (2001), answers the speculative claim of Cecchetti, Genberg, and Wadhwani (2003, 436) to the effect that the lack of importance that BG attribute to the stock market may be attributable to choices regarding the inclusion or exclusion of the output gap.

Lines 5 through 8 of the table repeat the information in lines 1 and 2 but for very different sets of policy preferences. Lines 5 and 6 are for a monetary authority that places a large weight on output stabilization (and a correspondingly low weight on inflation stabilization) in its decision making. Lines 7 and 8 cover the case of preferences skewed in the opposite direction. The basic conclusions under these two sets of preferences are the same as for the base-case preferences, namely, that while allowing feedback on the stock market is helpful, it is not overwhelmingly so.

Let us turn now to the outcome-based rules shown in Table 2. Note that the losses in this table are *still* normalized to the same policy two-parameter IFB rule shown in row 1 of Table 1. Two conclusions can be drawn from this table. First, we see that outcome-based rules are markedly inferior to the IFB counterparts, at least when policy-makers care substantially about inflation stabilization. It follows that forecasting matters and thus that the quality of the forecast also matters. Second, just as in the case of the IFB rules, the addition of feedback on stock prices does relatively little for economic performance. Given the similarity of the outcome-based results to those for IFB rules, and the similarity of results for alternative preferences, we henceforth restrict our attention to IFB rules for base-case preferences.[19]

The computations shown in Tables 1 and 2 were carried out under the assumption that bubble shocks and productivity shocks are uncorrelated. Given the historical experience of the 1990s, where the tech boom seemingly begat the stock market boom, it is arguably more reasonable to assume that the two shocks are correlated. Table 3 adopts this assumption, allowing a 0.9 covariance between the productivity shocks and the emergence of a bubble shock, while holding all else constant.

As before, the losses are normalized in the case where the Fed feeds back on expected future inflation and the contemporaneous change in stock prices (line 1 of the table).[20] Line 2 shows that the decision to omit feedback on the

---

19. The results in CGLW favouring feedback on the stock market depend, at least in part, on assessing welfare as changes in output rather than the output gap, with the argument that many of the fluctuations in potential output are non-fundamental and should therefore not be regarded “desirable.” See their footnote 12, p. 22. We thank Steve Cecchetti for pointing this out. In this paper, potential output depends on the actual capital stock and the fundamental (not the observed, stock market) value of that capital. The CGLW argument is a subtle one in that while one might argue that the capital accumulation induced by a bubble shock should not be there, once it is there, it is obviously part of the productive capacity of the economy, and society should use it as efficiently as possible. And whatever argument there is for excluding the influence of bubble shocks on capital accumulation and hence on potential, it is not true for productivity shocks.

20. For the record, the loss in the case with correlated shocks is about 50 per cent higher than in the first line of Table 1.

Table 2
Optimal coefficients and performance of outcome-based rules
(base-case calibration)

row	loss function	rule coefficients			loss
	λ	γ_π	γ_s	$\gamma_{\tilde{y}}$	L
(1)	0.5	2.05	1.65	–	1.24
(2)		2.37	–	–	1.42
(3)		4.03	3.99	10.35	1.03
(4)		4.54	–	13.30	1.18
(5)	0.9	1.74	1.28	–	1
(6)		1.98	–	–	1.13
(7)	0.1	4.51	3.31	–	1
(8)		4.41	–	–	1.24

Table 3
Optimal coefficients and performance of inflation-forecast-based rules
(correlated productivity and bubble shocks)

line	loss function	rule coefficients			loss
	p	γ_π	γ_s	$\gamma_{\tilde{y}}$	L
(1)	0.5	3.05	3.57	–	1
(2)		4.66	–	–	1.24
(3)		4.43	4.49	2.78	0.97
(4)		8.00	–	11.11	1.07

stock market is now a moderately costly one: the loss is about 24 per cent higher, significantly more costly than in Table 1. Line 3 shows that, unlike in Table 1, feedback on stock prices and on the output gap is largely complementary in that adding feedback to the output gap elicits a larger γ_s (and γ_π) than otherwise. Line 4 shows that replacing the stock price with the output gap in the rule can bring policy performance close to that of the base-case rule, but only if one is prepared to accept fairly extreme responses to inflation and output. The large values for γ_π and $\gamma_{\tilde{y}}$ suggest a possible lack of robustness to specification errors in the model.[21]

Since we regard the positive covariance of bubble and productivity shocks as a plausible feature of the model, we maintain this assumption for all subsequent experiments.

21. A looser convergence criterion in the algorithm, together with low-number starting values, renders an "optimal" policy of $\gamma_\pi = 3.97$ and $\gamma_{\tilde{y}} = 4.79$ but generates a normalized loss of 1.26. This reflects the model's topology, as noted previously.

The computations in Table 2 were carried out under the assumption that once initiated, a bubble persisted with probability $p = 0.5$. Table 4 explores the significance of this assumption for our results by recomputing optimal policies for a small perturbation of the continuation probabilities ranging between 0.45 and 0.55. Two salient facts can be gleaned from the table. First, the aggressiveness of optimal policies varies inversely with probability of continuation. This result obtains both, because a lower probability of continuation means that surprises from bubbles are larger, and because of the covariance of bubbles and productivity shocks. The second salient fact from the table is that the incremental loss from responding solely to forecasts of inflation is sharply decreasing in the continuation probability. In particular, low values of propagation lead to large deteriorations in policy performance under pure IFB targeting, owing to the large and persistent errors that arise when bubbles do propagate. Thus, while expected returns are independent of p, the variance of returns—which is what effectively enters the loss function—is not. In short, the case for responding directly to bubbles strengthens when bubbles are dramatic events.

Table 5 examines the implications of differences in the variance of the forcing bubble shock, holding constant the arrival rate and the continuation probability. Note that feedback on the stock market is at least marginally useful even when there are no bubble shocks. This is partly because the value of the capital stock is a state variable in the model, and also because the parsimony of the policy rule allows the stock market to proxy for other state variables that would appear in an optimal control rule.[22] The response of optimal coefficients to higher variances of the bubble shock is to raise somewhat the feedback coefficient on inflation in the pure IFB cases. What is not independent of these shocks is the loss associated with restricting direct feedback on the stock market. The differences among two-parameter rules is comparatively small. What is not small is the difference in economic performance between pure (one-parameter) IFB rules and two-parameter rules when the bubble shocks are large. When bubble shocks have a variance as large as 9, the incremental cost of directly ignoring the bubble is fairly significant. It is difficult to measure “fundamentals” of stock prices, even long after bubbles have burst, so there can be little precision in calibrating the magnitude of bubble shocks. Nonetheless, the bubbles produced by sequences of shocks with a variance of 9 are very large. This is why our base-case calibration uses a variance of 4, a conservative choice that in fairly large samples adds only marginally to the variance of output relative to the

22. Which begs the question: Why not use the fully optimal rule? The argument is that optimal rules are too fragile to be used in worlds where models are only approximations of reality, because their specification depends on the fine points of interaction between states that might not be modelled correctly.

Table 4
Optimal coefficients and performance of IFB Taylor-type rules
(alternative bubble continuation probabilities; covarying shocks)

line	continuation probability λ	rule coefficients γ_π	γ_s	loss L
(1)	0.45	3.57	3.91	1
(2)		5.96	–	1.45
(3)	0.50	3.05	3.57	1
(4)		4.66	–	1.24
(5)	0.55	2.53	3.03	1
(6)		3.35	–	1.14

Policy rule: $R_t = r^* + \gamma_\pi E_t \pi_{t+1} + \gamma_y \tilde{y}_t + \gamma_s (s_t - s_{t-1})$; variance-covariance matrix of forcing shocks: $diag\left[\sigma_\pi^2\ \sigma_c^2\ \sigma_g^2\ \sigma_z^2\ \sigma_u^2\right] = diag\left[1\ 1\ 1\ 0.1\ 4\right]$ with Poisson arrivals of bubble shocks at rate 0.02 and $cov_{u,z} = 0.9$. See Appendix 1 for details of the model.

Table 5
Optimal coefficients and performance of inflation-forecast-based rules
(alternative magnitudes of bubble shocks)

line	variance of bubble shock σ_u^2	rule coefficients γ_π	γ_s	target variable variances $\sigma_{\tilde{y}}^2$	σ_π^2	loss L
(1)	0	2.42	2.09	2.95	3.33	1
(2)		2.78	–	2.90	3.69	1.05
(3)	4	3.08	3.56	4.83	3.25	1
(4)		4.70	–	7.50	3.26	1.24
(5)	9	10.37	31.88	19.86	4.42	1
(6)		3.47	–	31.57	3.11	1.45

Policy rule: $R_t = r^* + \gamma_\pi E_t \pi_{t+1} + \gamma_y \tilde{y}_t + \gamma_s (s_t - s_{t-1})$; variance-covariance matrix of forcing shocks: $diag\left[\sigma_\pi^2\ \sigma_c^2\ \sigma_g^2\ \sigma_z^2\ \sigma_u^2\right] = diag\left[1\ 1\ 1\ 0.1\ column2\right]$ with Poisson arrivals of bubble shocks at rate 0.02. See Appendix 1 for details of the model.

case where there are no bubble shocks (shown in line 1). Nonetheless, the results on lines 5 and 6 of Table 5 do stand as a warning against complacency on bubbles.

2.2 Robust results

The results in Tables 3 and 4 suggest that there are possible worlds in which feeding back on (the change in the) stock market would be a welfare-improving policy, relative to a one-parameter pure IFB rule. However, the conditions under which this is so are fairly restrictive. One must have either large bubble shocks or large bubble surprises to make the case for directly responding to stock market developments. The results so far, however, have been for a relatively well-informed central bank and a symmetrically informed private sector. Under these circumstances, the forecast of inflation appearing in the policy rule can be assured of doing a good job of summarizing the states of the model economy. If the policy-maker's model were misspecified, however, there would be two potentially important implications for performance. First, the optimal coefficients in the rule would be incorrect, based on the wrong model. Second, the inflation forecast itself would be misspecified. In this section, we consider the implications of misspecification of this sort using the structured robust-control policies discussed in section 1.4.

We examine two sources of misspecification. The first source of misspecification is beliefs on the rate of growth of bubbles, conditional on the bubble not bursting. As already noted, the rational-bubble case sets the growth rate at unity so that investing in the stock market is a fair bet. The literature notes that the conditions under which a bubble can exist in a rational-expectations environment are very restrictive; see, e.g., Blanchard and Fischer (1989, chapter 5). Yet, as several contributors to the volume from the recent Federal Reserve Bank of Chicago/World Bank conference on asset price bubbles make clear, the real world seems to be replete with bubbles (see Hunter, Kaufman, and Pomerleano 2003). One way to describe the role of monetary policy—and in particular, the role of monetary policy in a world of uncertainty—is to keep the economy out of trouble. To the extent that this is so, the object of concern should not be rational bubbles as such, since investors taking fair bets under symmetric and nearly complete information present little risk to the economy. A more problematic scenario, if it exists, is "irrational bubbles," that is, bubbles that do not obey the rules governing linear rational-expectations models. The second source of uncertainty, given less time here, is the continuation probability of bubbles.

In our base-case model, the (linearized) bubble process follows:

$$u_{t+1} = a \cdot u_t + \varepsilon_{u,t+1} \qquad a = 0.99 . \tag{6}$$

For our first experiment, we consider a range of possible (true) values for a, with the lower bound set at $\underline{a} = 0.90$ and the upper bound set about as close to unity as is feasible: $\bar{a} = 0.9999$. Relative to the base case, this range of uncertainty is not symmetric, of course, but it reflects the balance of risks inherent in holding to the prior belief that bubbles are rational. Nonetheless, to explore the implications of this asymmetry, we study the borderline cases where the reference model either has $a = 0.90$ or has $a = 0.9999$, but where the policy-maker wishes to consider hedging against any value for a between those two values. As already noted, the policy-maker's objective is then to choose a vector of feedback coefficients to minimize the loss function, equation (1), subject to the perceived, or reference model, the variance covariance matrix of forcing shocks, the form of the policy rule, equation (2), and the loss-maximizing choice of $a \in \left[\underline{a}\ \bar{a}\right]$. In this instance, the solution to this min-max problem arrives at a corner solution for a; that is, the loss-maximizing choice for a will be either $\underline{a}$ or $\bar{a}$.

Table 6 shows the results for the robustness with respect to bubble persistence, under preferences that assign equal penalties of one-half on squared output and inflation gaps. The upper panel of the table shows the results when the reference model is $a = 0.90$. The first two rows cover the case where the local worst case is the reference model itself, meaning there is no robustness in play. These two lines serve as benchmark cases. Lines 3 and 4 show the effects of protecting against a misspecification of a in the reference model: $a \in \left[0.9\ 0.9999\right]$. The bottom panel—lines 5 to 8—shows the same experiment except for a reference model with $a = 0.9999$.

The way to interpret the table is to read off of the last column on the right the cost of protecting against misspecifications using the rules (and inflation forecasts) of selected models. So the first two lines of each panel show the cases where the perceived model and the worst-case models are the same; that is, these are the cases where the authority chooses not to protect against misspecification. These two lines should be compared with the next two, which show the cost of protection against a world of $a = 0.9999$ in a $a = 0.90$ reality. Similarly, the reverse case, where reality is $a = 0.9999$ and $a = 0.90$ is being protected against is compared in lines 5 and 6 versus 7 and 8. The answer, in a nutshell, is that there is little difference among the policies in terms of their performance within the range of values for a against which the policy-maker attempts to protect—at least for the modest shocks we use here. In other words, the misspecification studied here would not change whatever conclusion one might draw from Table 3.

The middle columns showing the rule coefficients give a hint as to why these tepid results obtain. The optimal coefficients for these models do not vary a great deal. That by itself is not fully informative since the inflation forecasts upon which the rules are feeding back differ. What Table 6 shows, however, is that the inflation forecasts also differ little. This is a manifestation of the stabilizing power of rational expectations. Our experiments have two key features. First, the private sector knows what the monetary authority is doing, even if the authority is unclear about the model. That is, the private sector has better information than does the policy-maker. This seems a

Table 6
Robust policies and performance of inflation-forecast-based rules
(alternative conditional growth rates of bubbles;
equal weights in loss function)

row	**model** a	**boundaries** $[\underline{a}\ \bar{a}]$	**rule coefficients** γ_π	γ_s	**target variable variances** $\sigma^2_{\tilde{y}}$	σ^2_π	**loss** L
(1)	0.90	0.90	3.14	3.38	3.11	3.29	1
(2)		0.90	4.15	–	2.01	3.26	1.29
(3)		0.9999	3.07	3.58	3.12	3.29	1.02
(4)		0.9999	4.66	–	2.04	3.28	1.33
(5)	0.9999	0.9999	3.07	3.58	8.15	3.26	1
(6)		0.9999	4.66	–	10.99	3.28	1.25
(7)		0.90	3.14	3.38	8.16	3.23	1.03
(8)		0.90	4.15	–	11.12	3.31	1.30

Policy rule: $R_t = r^* + \gamma_\pi E_t \pi_{t+1} + \gamma_y \tilde{y}_t + \gamma_s(s_t - s_{t-1})$; variance-covariance matrix of forcing shocks: $diag\left[\sigma^2_\pi\ \sigma^2_c\ \sigma^2_g\ \sigma^2_z\ \sigma^2_u\right] = diag\left[1\ 1\ 1\ 0.1\ 4\right]$ with $cov(u, z) = 0.9$ and Poisson arrivals of bubble shocks at rate 0.02. See Appendix 1 for details of the model. The loss function assigns equal weights to squared output and inflation gaps.

reasonable assumption, albeit perhaps a strong one. Second, the policies chosen by our ill-informed policy-maker always stay in the stable region of the model, that is, the northeast and southwest regions shown in Figure 4. Together, these two features establish a strongly stabilizing force in the economy. Had the chosen policies ended up in the indeterminate or unstable regions—a possibility, given the misspecifications considered—the answers would have been much different.

In addition to the experiment on robustness over bubble persistence, we also experimented with uncertain bubble duration, as well as with a few structural model parameters. For reasonable ranges of uncertainty, the answers were broadly the same as those just described.

Conclusion

This paper has examined the role of monetary policy in responding to stock market bubbles. The analysis centred around extensions of the Bernanke-Gertler-Gilchrist (1999) model, a New Keynesian model with a financial accelerator mechanism. Our efforts were concentrated in three directions. First, we embellished the model, adding more persistence and stronger behavioural links between investment and the stock market to test the breadth of applicability of the argument of Bernanke and Gertler (1999, 2001) that monetary policy should react to asset prices only insofar as they affect the forecast of future inflation. Second, we broadened the list of experiments and preferences to which the model was subjected. Third, we examined the implications of model uncertainty in the sense of Knight for policy design and performance. We interpret our results as mostly supportive of the hands-off view advanced by Bernanke and Gertler, with some reservations. Under the base-case calibration of the model, we found no more than modest gains from responding directly to stock prices. Similarly, we found little reason to engage in robust responses to model uncertainty in the key area of the bubble process, at least for balanced preferences and modest bubble shocks. Put simply, so long as policy is seen to be strongly stabilizing, a policy of pure inflation-forecast targeting does a reasonable and robust job.

A potential fly in the ointment is that there are alternative calibrations of the bubble process for which responding to bubbles is more efficacious. In particular, when the probability of bubble persistence is small, the resulting surprises from bubbles that propagate are large. Similarly, when the magnitude of bubble shocks is large, so are the surprises and the costs. Both instances considerably strengthen the case for responding directly to stock market developments. This finding is a bit problematic given that the measurement of stock market fundamentals is difficult and thus that the measurement of bubbles is difficult as well. There is little guidance in the data regarding what a sensible process might be. The case for responding directly to (perceived) bubbles becomes stronger when bubbles become large, and private agents are surprised by their growth when they are common and when they are correlated with productivity shocks.

Looking ahead, uncertainty about the measurement of fundamentals adds to the complexity of the issue. Both Bernanke and Gertler (2001) and Cecchetti, Genberg, and Wadhwani (2003) point to the detection of bubbles as a key issue. The results shown here suggest that failing to react systematically to large developments in stock markets can be costly, while ignoring small bubbles is less worrisome. This suggests that a non-linear feedback rule that responds to bubbles only when they become large enough that they become important macroeconomic phenomena—and when their size leaves little doubt that fundamentals cannot be the sole driving factor—may be a welfare-improving strategy. This line of research seems a fruitful direction in which to head. In a related vein, modelling the measurement of fundamentals in quasi-real time would also be advantageous. That said, neither course of action can be taken on at low cost; the computational challenges are impressive.

Appendix 1
A Version of the Bernanke-Gertler-Gilchrist Model

$$y_t = cy\ c_t + cey\ ce_t + iy\ i_t + gy\ g_t \quad \text{(A1)}$$

$$c_t = -\sigma rr_t + \phi_{c1} E_t c_{t+1} + \phi_{c2} E_t c_{t+2} + \phi_{c3} c_{t-1} + \varepsilon_{c,t} \quad \text{(A2)}$$

$$rr_t = E_t f_{t+1} + \psi(nw_t - q_t - u_t - k_t) \quad \text{(A3)}$$

$$f_t = (1-\nu)(mc_t + y_t - k_t) + \nu q_t - q_{t-1} + \nu\varepsilon u_t - \varepsilon u_{t-1} \quad \text{(A4)}$$

$$s_t = \phi(E_{t-1} f_t - \nu E_{t-1} k_t + (1-\nu) i_{t+1}) \quad \text{(A5)}$$

$$nw_t = \chi[f_t - (1-nk)(rr_{t-1} + \psi(k_{t-1} + q_{t-1}))]$$

$$+ \nu u_t - \varepsilon[1 + (1-nk)(\psi - (1-bx))]u_{t-}$$

$$+ [(1-nk)\psi + nk]nw_{t-1} + (1-\kappa\ rk)(nk/\kappa)y_t \quad \text{(A6)}$$

$$ce_t = (kn/\psi)[nw_t - (1-\kappa\ rk)(nk/\kappa)y_t] \quad \text{(A7)}$$

$$k_{t+1} = \delta i_t + (1-\delta)k_t \quad \text{(A8)}$$

$$y_t = z_t + \alpha k_t + (1-\alpha)h_t \quad \text{(A9)}$$

$$mc_t = (1/\sigma)c_t + \gamma_h h_t - y_t \quad \text{(A10)}$$

$$E_{t-1}\pi_t = \lambda mc_t + \theta_b \pi_{t-1} + \theta_f E_{t-1}\pi_{t+1} + \mu_t^{\pi} \quad \text{(A11)}$$

$$rr_t = rn_t - E_t \pi_{t+1} \quad \text{(A12)}$$

$$rn_t = \gamma_\pi E_t \pi_{t+1} + \gamma_u(\Delta q_t + \Delta u_t + \Delta k) \quad \text{(A13)}$$

$$g_t = \rho_g g_{t-1} + \mu_t^g \quad \text{(A14)}$$

$$z_t = \rho_z z_{t-1} + \mu_t^z \quad \text{(A15)}$$

$$u_t = [bx\ rk/(1\text{-}\delta)]u_{t-1} + \mu_t^u \quad \text{(A16)}$$

Table A1
Key model parameters

parameter	description	value
σ	inverse elasticity of intertemporal substitution	5
ϕ_{c3}	coefficient on lagged consumption in consumption Euler equation	0.33086
ψ	elasticity of financial leverage premium	0.05
ε	extent to which entrepreneurs participate in consumption	0.75
ϕ	elasticity of investment with respect to Tobin's Q	0.5641
χ	wealth accumulation constant (from linearization)	1.9794
bx	bubble propagation parameter $a((1-\delta)/(rk))$	0.9604
δ	quarterly rate of capital depreciation	0.025
rk	steady-state rate of return on capital $1/\beta + 0.02/4$	1.0151
β	subjective rate of time preference	0.99
α	capital's share of income	0.33
ν	linearization constant $(1-\delta)/(\alpha/(\mu \cdot ky) + 1 - \delta)$	0.9605
θ_f	weight on forward expectations in price equation	0.59579
θ_b	weight on lagged inflation in price equation	0.4012
λ	elasticity of inflation with respect to marginal cost	0.025827
ρ_g	propagation of government expenditure shocks	0.95
ρ_z	propagation of total-factor productivity shocks	0.99
a	conditional rate of propagation of bubbles $bk(rk/(1-\delta))$	0.99

References

Abel, A.B. 1990. "Asset Prices under Habit Formation and Catching Up with the Joneses." *American Economic Review* 80 (2): 38–42.

Amato, J.D. and T. Laubach. 2002. "Rule-of-Thumb Behaviour and Monetary Policy." Finance and Economics Discussion Series Paper No. 2002–05. Board of Governors of the Federal Reserve System.

Batini, N. and E. Nelson. 2000. "When the Bubble Bursts: Monetary Policy Rules and Foreign Exchange Market Behavior." Bank of England. Unpublished.

Bernanke, B.S. and M. Gertler (BG). 1999. "Monetary Policy and Asset Price Volatility." In *New Challenges for Monetary Policy,* 77–128. A symposium sponsored by the Federal Reserve Bank of Kansas City. Kansas City: Federal Reserve Bank of Kansas City.

———. 2001. "Should Central Banks Respond to Movements in Asset Prices?" *American Economic Review Papers and Proceedings* 91 (2): 253–57.

Bernanke, B.S., M. Gertler, and S. Gilchrist (BGG). 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." In *Handbook of Macroeconomics,* Volume 1C, edited by J.B. Taylor and M. Woodford, 1341–93. Amsterdam: North-Holland.

Blanchard, O.J. and S. Fischer. 1989. *Lectures on Macroeconomics.* Cambridge, MA: MIT Press.

Bullard, J.B. and E. Schaling. 2002. "Why the Fed Should Ignore the Stock Market." *Federal Reserve Bank of St. Louis Review* 84 (2): 35–41.

Carlstrom, C.T. and T.S. Fuerst. 1997. "Agency Costs, Net Worth, and Business Fluctuations: A Computable General Equilibrium Analysis." *American Economic Review* 87 (5): 893–910.

Casares, M. and B.T. McCallum. 2000. "An Optimizing IS-LM Framework with Endogenous Investment." National Bureau of Economic Research Working Paper No. 7908.

Cecchetti, S.G., H. Genberg, J. Lipsky, and S. Wadhwani (CGLW). 2000. *Asset Prices and Central Bank Policy.* International Center for Monetary and Banking Studies and Centre for Economic Policy Research.

Cecchetti, S.G., H. Genberg, and S. Wadhwani. 2003. "Asset Prices in a Flexible Inflation Targeting Framework." In *The Implications for Monetary, Regulatory, and International Policies,* edited by W.C. Hunter, G.G. Kaufman, and M. Pomerleano, 427–44. Cambridge, MA: MIT Press.

Coletti, D., B. Hunt, D. Rose, and R. Tetlow. 1996. *The Bank of Canada's New Quarterly Projection Model, Part 3, The Dynamic Model: QPM.* Technical Report No. 75. Ottawa: The Bank of Canada.

Dupor, B. 2001. “Nominal Price versus Asset Price Stabilization.” The Wharton School. Unpublished.

Finan, F.S. and R. Tetlow. 1999. “Optimal Control of Large, Forward-Looking Models: Efficient Solutions and Two Examples.” Finance and Economics Discussion Series Paper No. 1999–51. Federal Reserve Board.

Giannoni, M.P. 2001. “Robust Optimal Monetary Policy in a Forward-Looking Model with Parameter and Shock Uncertainty.” Federal Reserve Bank of New York. Unpublished

———. 2002. “Does Model Uncertainty Justify Caution? Robust Optimal Policy in a Forward-Looking Model.” *Macroeconomic Dynamics* 6 (1): 111–44.

Gilboa, I. and D. Schmeidler. 1989. “Maxmin Expected Utility with a Non-unique Prior.” *Journal of Mathematical Economics* 18 (2): 141–53.

Goodhart, C. and B. Hofmann. 2000. “Asset Prices and the Conduct of Monetary Policy.” London School of Economics. Unpublished.

Hansen, L.P. and T.J. Sargent. 2005. *Robustness*. Princeton University Press.

Hunter, W.C., G.G. Kaufman, and M. Pomerleano (ed.). 2003. *Asset Price Bubbles: The Implications for Monetary, Regulatory, and International Policies*. Cambridge, MA: MIT Press.

Kaplan, S.N. 2003. “Valuation and New Economy Firms.” In *Asset Price Bubbles: The Implications for Monetary, Regulatory, and International Policies*, edited by W.C. Hunter, G.G. Kaufman, and M. Pomerleano, 391–401.

Kent, C. and P. Lowe. 1997. “Asset-Price Bubbles and Monetary Policy.” Reserve Bank of Australia Research Discussion Paper No. 9709.

Levin, A., V. Wieland, and J.C. Williams. 2003. “The Performance of Forecast-Based Monetary Policy Rules Under Model Uncertainty.” *American Economic Review* 93 (3): 622–45.

Lubik, T.A. and F. Schorfheide. 2004. “Testing for Indeterminacy: An Application to U.S. Monetary Policy.” *Journal of American Economic Review* 94 (1): 190–217.

Mussa, M. 2003. “Asset Prices and Monetary Policy.” In *Asset Price Bubbles: The Implications for Monetary, Regulatory, and International Policies*, edited by W.C. Hunter, G.G. Kaufman, and M. Pomerleano, 41–50. Cambridge, MA: MIT Press.

Onatski, A. and J.H. Stock. 2002. “Robust Monetary Policy under Model Uncertainty in a Small Model of the U.S. Economy.” *Macroeconomic Dynamics* 6 (1): 85–110.

Svensson, L.E.O. 2002. "Inflation Targeting: Should It Be Modeled as an Instrument Rule or a Targeting Rule?" *European Economic Review* 46 (4–5): 771–80.

Tetlow, R.J. and P. von zur Muehlen. 2001. "Robust Monetary Policy with Misspecified Models: Does Model Uncertainty Always Call for Attenuated Policy?" *Journal of Economic Dynamics and Control* 25 (6–7): 911–49.

———. 2004. "Avoiding Nash Inflation: Bayesian and Robust Responses to Model Uncertainty." *Review of Economic Dynamics* 7 (4): 869–99.

———. 2005. "Robustifying Learnability." <http://www.roberttetlow.com>. Unpublished.

von zur Muehlen, P. 1982. "Activist vs. Non-Activist Monetary Policy: Optimal Rules Under Extreme Uncertainty." Finance and Economics Discussion Series Paper No. 2001–2. Federal Reserve Board.

Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton: Princeton University Press.