Macroprudential Policy under Uncertainty*

Saleem Bahaj and Angus Foulis
Bank of England and Centre for Macroeconomics

We argue that uncertainty over the impact of macroprudential policy need not make a policymaker more cautious. Our starting point is the classic finding of Brainard that uncertainty over the impact of a policy instrument will make a policymaker less active. This result is challenged in a series of richer models designed to take into account the more complex reality faced by a macroprudential policymaker. We find that asymmetries in policy objectives, the presence of unquantifiable sources of risk, the ability to learn from policy, and private-sector uncertainty over policy objectives can all lead to more active policy.

JEL Codes: D81, E58.

1. Introduction

The macroprudential toolkit available to policymakers across several central banks is new and largely untested. For example, in the United Kingdom, the Bank of England's Financial Policy Committee (FPC) has, since the financial crisis, received powers to alter bank capital requirements and to place restrictions on the terms of household mortgages for macroprudential purposes. Neither of these policy tools has been used previously, so their impact and the Committee's reaction function remain unclear. Moreover, in contrast to

*We are grateful to the editor, Rafael Repullo, our discussant, Francois Gourio, and to participants at the 2016 Annual Conference of the *International Journal of Central Banking* for helpful comments and feedback. We would also like to further thank David Aikman, John Barrdear, Oliver Bush, Julia Giese, Rodrigo Guimaraes, Alexandre Kohlhas, Marc Hintersweiger, Sujit Kapadia, Frederic Malherbe, Roland Meeks, Tamarah Shakir, Rhiannon Sowerbutts, Misa Tanaka, and Martin Taylor. This paper contains the views of the authors and not necessarily those of the Bank of England, the MPC, the FPC, or the PRA board. The working paper version of this paper is Bahaj and Foulis (2016) and contains additional results that we reference in the text. Author contact: saleem.bahaj@bankofengland.co.uk; angus.foulis@bankofengland.co.uk.

monetary policy, where price stability can be judged against the rate of inflation, the objective of macroprudential policymakers, the stability of the financial system, is inherently unobservable. Thus macroprudential policymakers face a high degree of uncertainty over the impact and effectiveness of their tools and a target variable they cannot perfectly observe. In the face of this uncertainty, a prevalent view is that a cautious approach is warranted: if a policymaker is unsure what a tool does, she should use it gingerly. Indeed, this is a classic result from the literature on optimal policy under uncertainty as shown in Brainard (1967).

This paper takes the Brainard model as a starting point and asks: is the uncertainty faced by macroprudential policymakers sufficient to justify a cautious stance to macroprudential policy? The Brainard framework is stylized and static, and there are multiple reasons why a policymaker may want to overlook its conclusions. In this paper, we present the results from some simple extensions to this framework to illustrate how uncertainty could alter the behavior of policymakers. The analysis here is drawn from the existing literature, but our goal is to frame the issue of uncertainty in the macroprudential context.

As a starting point, and to fix ideas, we recast the Brainard model as a macroprudential policy problem where the policymaker attempts to stabilize the resilience of a financial system. In particular, we assume that the policymaker is trying to stabilize the level of financial stability denoted x about some target x^* through the use of a tool k (for example, a time-varying capital requirement such as the countercyclical capital buffer) that controls x imperfectly. The relationship between k and x is linear,

$$x = bk + u, \tag{1}$$

and is subject to two sorts of uncertainty. First, b, the parameter governing how k affects x, is uncertain with prior mean $\bar{b} > 0$ and variance σ_b^2. The failure of the policymaker to observe b perfectly could, for instance, reflect uncertainty over the impact of capital requirements on financial stability. Second, there is unobserved variation in the level of financial stability, u, which is independent of the policymaker's action and has prior mean 0 and variance σ_u^2. To simplify the exposition, we assume the two sources of uncertainty are uncorrelated in what follows. This is a standard assumption that model and shock uncertainty are not related; however, similar results do emerge in a more general setting.[1]

Policymakers should find an unstable financial system undesirable; however, an overly stable system may dampen economic activity and impose a burden on the financial intermediaries or consumers. To cite a cliche: policy should aim to avoid the stability of the graveyard. The value of x^* can, therefore, be thought of as the optimal level of financial stability, trading off a stable versus an active financial system. Similarly, adjusting k may also impose costs; for example, by forcing banks to pay the underwriting fees associated with equity issuance. Thus we assume the policymaker has the following objective:

$$W = -\frac{1}{2}\mathbb{E}((x - x^*)^2 + \lambda k^2). \tag{2}$$

The parameter $\lambda > 0$ captures the policymaker's view over the relative cost of stabilizing k versus x about their optimal levels.[2] Under these assumptions, it is straightforward to take the policymaker's first-order condition ($\mathbb{E}[b^2 k + bu - bx^* + \lambda k] = 0$) and show that the best choice of k under uncertainty (denoted k^u) is

$$k^u = \frac{\bar{b}x^*}{\bar{b}^2 + \sigma_b^2 + \lambda} < \frac{\bar{b}x^*}{\bar{b}^2 + \lambda} = k^c, \tag{3}$$

where k^c denotes the level of k the policymaker should choose if she faced no uncertainty. Policy is less active under uncertainty. Further, as uncertainty over the impact of policy, σ_b^2, increases, k^u falls. This means that the policymaker should use her tool "less" as uncertainty increases, and for any given level of uncertainty the policymaker should choose k at a lower level than if she was certain.

[1]The more general case of $Cov(b, u) \neq 0$ is discussed in the working paper version of this paper.

[2]Note that we are assuming that the policymaker has a symmetric objective. This is potentially unrealistic both for financial stability and the costs of changing k: low financial stability may be more worrisome than high financial stability, and cutting capital requirements may not impose much of a burden on the financial system relative to raising them. In section 2 we consider an asymmetric financial stability objective for the policymaker with low financial stability disproportionately costly.

The intuition for this result is simply that additional uncertainty over the tool is perceived to introduce additional volatility into the economy when it is used, which is undesirable from the policymaker's perspective. A policy instrument whose impact is uncertain should be used more sparingly. Further, when there is greater uncertainty, the instrument should be used less. In this static model we associate the degree of policy activism with the level of k chosen. A more natural interpretation of "activism" might be the responsiveness of policy to shocks. The model can support such an interpretation by considering the response of k to the desired level of financial stability x^*—greater uncertainty over the impact of policy will make the policy less responsive to the realized value of x^*. In section 5 we explicitly consider a model in which policy is set after observing shocks to financial stability.

An important feature of Brainard's analysis is that the form of uncertainty matters. Brainard's results are sometimes misleadingly cited as a general rule that a policymaker should do less in the face of uncertainty. However, note that σ_u^2 does not appear in k^u. Therefore, the second conclusion from this form of model is that being unsure over the state of the economy (for example, the inherent stability of the financial system) should not alter policymakers' behavior.

This paper argues that there are several types of uncertainty and multiple channels through which uncertainty can affect policymaking. The result that policy should be more cautious in the presence of uncertainty does not hold in general, particularly for specific examples that are relevant to macroprudential policy. As we shall see, if anything, the results speak to a more active policy stance in the face of uncertainty. This complements the need for policymakers to guard themselves against inaction bias. Financial stability risks are hard to measure (or unobservable), and actions to address them may have short-term costs making regulatory forbearance tempting. The lags associated with macroprudential policy instruments, both in terms of implementation and transmission, mean that there is the potential for policymakers to move too late to build resilience in the financial system ahead of crises.

To make these points, we consider several extensions to the Brainard model. Our first extension considers an asymmetric objective function for financial stability in which a crisis is disproportionately costly for the policymaker. If the policy tool is sufficiently effective on average, policy will become more active the greater this asymmetry. In a second extension, we recognize that macroprudential policy is concerned with rare events, the probability of which is difficult to quantify. In such a situation, the policymaker may wish to behave in a robust fashion, preparing for the worst-case scenario. This can also lead to more active policy. Our third extension considers a dynamic model which allows for learning. Using the tool today reduces uncertainty about its impact tomorrow, but may initially increase volatility. If the motivation to learn is sufficiently strong (i.e., the policymaker's discount rate is sufficiently low), optimal policymaking can become more active with greater uncertainty. In our final extension, we consider the interaction between private-sector uncertainty and the uncertainty facing the policymaker. In addition to directly affecting financial institutions, macroprudential actions may have a broader impact through signaling information about risks to financial stability. We show that this signaling channel will be less powerful when there is greater private-sector uncertainty about policy objectives. Consequently, when the private sector is more unsure about why the policymaker is acting, the policymaker will need to be more active in order to offset the diminished signaling power of the tool.

1.1 Related Literature

Brainard (1967) lays out the canonical case for higher uncertainty leading to diminished policy activism. Beyond this classic work, this paper has several links with the academic literature on policy under uncertainty. However, most of this previous work has focused on the monetary policy context; our contribution is to recast some of the findings in the context of macroprudential policy and discuss their relevance.

In terms of policy under fundamental (or Knightian) uncertainty, the approach of using robust control and min-max dates back to at least Wald (1950); a more detailed, modern treatment is available in Hansen and Sargent (2001, 2008). Barlevy (2009) offers additional exposition over the framework. Robust control and min-max based optimal policy problems generally deliver more aggressive policy action in response. However, Onatski and Stock (2002) show that this finding does not extend to all model perturbations.

On the topic of learning through policy actions to reduce uncertainty, several papers—for example, Orphanides and Williams (2007) and Wieland (2000, 2006)—show how the desire to learn can lead to policy being more aggressive in the face of uncertainty. An alternative strand of the literature (see, for example, Besley 2001, Landes 1998, and Mukand and Rodrik 2002) argues that cross-sectional variation in policymaking across countries or political parties should also be encouraged to foster innovation and allow for learning about the impact of alternative policies. In contrast, other authors have challenged the gains from actively introducing policy variation in order to learn. In a quantitative assessment, Svensson and Williams (2007) show that the benefits for experimentation can often appear very modest in a generic linear quadratic forward-looking setup allowing for model uncertainty. Cogley and Sargent (2005) also cast doubt on the benefits of setting policy with learning in mind and raise a more general point that the gains from more active policy for the purposes of learning could be limited if there are sufficient natural experiments.

The interaction of the uncertainty of the public and policymaker, alongside asymmetric information and the informational value of government policy, are discussed in the early contributions of King (1982) and Weiss (1982). More recently, Lorenzoni (2010) and Angeletos, Iovino, and La'O (2011, 2015) analyze how to optimally set policy in static models with dispersed information.

2. Asymmetric Objective Function

The model in the introduction has a *symmetric* financial stability loss function $(x - x^*)^2$, in which financial stability being above the target level x^* is equally as costly to the policymaker as financial stability being below target. Such objective functions have been frequently used for modeling inflation targeting, as high inflation and deflation can both be costly. However, it is not clear that this functional form is appropriate for financial stability: the losses associated with a financial crisis may be significantly greater than those imposed by having excessive stability. In this section we consider an extension to the Brainard model with an asymmetric loss function that captures this property.

Figure 1. Graphical Illustration of the Linex Function

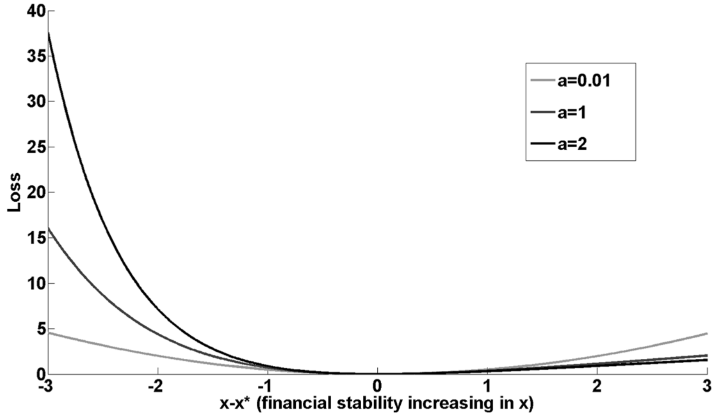

2.1 Modeling Framework

A tractable way to model an asymmetric loss function is to use the linex function (Varian 1975).[3] We consider the following objective function for the macroprudential policymaker:

$$W = -\frac{1}{a^2}\mathbb{E}\left[exp\left(a\left(x^* - x\right)\right) - a\left(x^* - x\right) - 1\right] - \frac{\lambda}{2}k^2, \qquad (4)$$

where $a > 0$, and the remaining variables are defined as in the introduction. An attractive property of equation (4) is that it nests the benchmark quadratic loss function used in the basic Brainard model, collapsing to it for $a \rightarrow 0$. When $a > 0$, equation (4) is asymmetric, with a greater loss when financial stability is low ($x < x^*$) than when it is too high ($x > x^*$). Further, the greater a is, the greater this differential, and the more costly low financial stability is relative to high financial stability. Figure 1 plots the financial stability loss function for a range of values of a (holding $k = 0$). As can be seen, for larger values of a, the costs of very low financial stability (to the left of 0) can be significantly higher than the costs of very high financial stability (to the right of 0).

[3]We would like to thank David Aikman for suggesting this functional form.

As in the basic Brainard model, the policy tool k influences financial stability x in a linear way,

$$x = bk + u, \tag{5}$$

with uncertainty over the impact of the tool, b, and shocks to financial stability, u. For tractability we assume that $b \sim N\left(\bar{b}, \sigma_b^2\right)$, $u \sim N\left(0, \sigma_u^2\right)$, and $Cov\,(b, u) = 0$. Using equation (5) and the properties of the log-normal distribution, the objective function can be written as[4]

$$W = -\frac{1}{a^2}\left[exp\left(a\left(x^* - \bar{b}k\right) + \frac{a^2\left(\sigma_b^2 k^2 + \sigma_u^2\right)}{2}\right) - a\left(x^* - \bar{b}k\right) - 1\right] - \frac{\lambda}{2}k^2. \tag{6}$$

This function is strictly concave and has a unique global maximum $\tilde{k} > 0$ (see the appendix for the proof).

In the appendix we show that if the tool is expected to be sufficiently powerful, that is to say that $\bar{b}$ is large enough that a negative realization of b is a remote possibility, then

$$\frac{d\tilde{k}}{da} > 0. \tag{7}$$

That is, in the presence of uncertainty, greater asymmetry in the policymaker's financial stability objective leads to more active policy. When the losses from financial instability are sufficiently greater than those from too much stability and the policy tool is reasonably effective, the policymaker is more active to make sure low-stability outcomes are avoided. Such activism brings large benefits in reducing the cost of a financial crisis with relatively small costs if financial stability ends up being too high.[5]

---

[4] Specifically, we make use of the result that for $Z \sim N\left(\mu, \sigma^2\right) : \mathbb{E}\left[exp\left(Z\right)\right] = exp\left(\mu + \frac{1}{2}\sigma^2\right)$.

[5] In the working paper version of this paper, we also show that if $\bar{b}$ and a are sufficiently large, then $\frac{d\tilde{k}}{d\sigma_b^2} > 0$ and $\frac{d\tilde{k}}{d\sigma_u^2} > 0$, as the asymmetry in the objective prompts the policymaker to insure against downside risks by being more active.

2.2 Discussion

There is good reason to think macroprudential policymakers should have an asymmetric preference when it comes to financial stability. The substantial skew with regards to negative economic outcomes when financial stability risks materialize means that the costs of missed downside risks may be much larger than benefits of erring towards looser policy. Theoretical models that deliver endogenous crises via occasionally binding constraints (see, for example, Bianchi and Mendoza 2010, Jeanne and Korinek 2013, Korinek 2011, and Korinek and Simsek 2014) illustrate how severe asymmetries can emerge in stylized macroeconomic models and show that this can provide a normative justification for macroprudential policies that are contractionary in normal circumstances. The intuition for the findings above are similar: as the policymaker is much more averse to low financial stability outcomes, she buys insurance against low financial stability by tightening policy. The cost of this is an overly tight policy choice in situations when the financial system is stable.

We see this reasoning in the mandates of macroprudential policymakers. In the United Kingdom, for instance, the Financial Policy Committee's primary objective is to identify, monitor, and take action to remove systemic risks, thus indicating a focus on the downside of risks to the financial system. In contrast, the mandate for monetary policy tends to be to avoid both high and low inflation, and in many cases the target is explicitly symmetric (see, for example, the mandate of the Sveriges Riksbank). However, there are reasons why asymmetries may also matter in the context of monetary policy. For example, if the zero lower bound exacerbates deflationary shocks, the policymaker will have a stronger preference for avoiding below-target inflation. On the other hand, inflation bias tends to be positive, hence central banks could be concerned that above-target inflation will be more damaging to their credibility. Patton and Timmermann (2007) provide some empirical evidence that monetary policymakers have asymmetric preferences. Specifically they show that the Federal Reserve's revealed preferences, based on forecast errors, are more averse to lower realizations of output growth.

3. Targeting the Worst-Case Outcome (Robustly)

The conclusions from Brainard (1967), and indeed most modeling frameworks where policymakers maximize the expectation of an objective, rely on policymakers being able to assign probabilities to potential future scenarios. The previous section considered what happened if more weight was placed on negative realizations but the policymaker was able to assess the likelihood of such realizations occurring. However, this could be an unrealistic way of describing the uncertainty faced by macroprudential policymakers. It is not always possible to say with confidence how likely a relevant outcome will be. There is an inescapable need for policymakers to make judgments about the state of the world that cannot be backed by statistical analysis (Svensson 2002). This sort of unquantifiable uncertainty is sometimes referred to as fundamental or Knightian uncertainty. And it is a form of uncertainty that may be particularly troublesome for macroprudential policymakers given the innovation and increasing complexity of the financial system, which can make risks difficult to quantify. Furthermore, the objective of macroprudential policymakers is defined in terms of the resilience of the financial system, suggesting a focus upon rare events whose likelihood is undetermined.

A popular method among economists for incorporating fundamental uncertainty and concerns over severe outcomes into policymaking decisions is to rely upon a robust control approach. This approach favors policies that avoid large losses across scenarios regardless of how likely any given scenario is. In practice this is implemented by what is called a min-max framework (Wald 1950): a policymaker sets policy to minimize her losses assuming that the parts of the problem she is fundamentally uncertain about are chosen to maximize her loss. This is equivalent to making the worst-case scenario as palatable as possible.

3.1 Modeling Framework

To introduce a robust control motive into the Brainard model, we modify the relationship between k and x to be

$$x = bk + u + v, \tag{8}$$

where with the exception of v, all variables are defined as in equation (1) in the introduction. The term v captures what we refer to as fundamental uncertainty over x, the level of financial stability. This is a source of uncertainty which is ambiguous, and the policymaker is unable to attach a probability distribution to it. A robust approach to dealing with the uncertainty captured by v is to choose policy on the assumption that the worst outcome has happened so as to avoid large losses. To implement this, the min-max framework has the policymaker choosing k to maximize its objective conditional on "nature" choosing v to minimize the objective

$$max_k min_v - \frac{1}{2}\mathbb{E}[(x - x^*)^2 + \lambda k^2] + \frac{\theta}{2}v^2,\ \theta > 1. \tag{9}$$

Both nature and the policymaker choose their action conditional on each other's choice, and a pure-strategy Nash equilibrium solves the model. In effect, the policymaker and nature play a phantom game with each other, and nature behaves as a fictitious evil agent that chooses v to ensure the worst-case outcome occurs. An important element to the robust control problem is the term θv^2. Using the definitions of Hansen and Sargent (2008), the parameter θ is called entropy and, at a fundamental level, both represents how wrong the policymaker thinks they can be about the true level of x and penalizes extreme realizations of v.[6] From the perspective of the optimization problem, θv^2 simply serves as a constraint on nature when choosing how bad the worst-case scenario can be. The constraint $\theta > 1$ ensures that nature's problem is convex and so a finite v is chosen.

It is straightforward to show that the solution for nature, taking the actions of the policymaker as given, is given by

$$v = \frac{\bar{b}k - x^*}{\theta - 1}. \tag{10}$$

Thus as θ goes to infinity, nature is constrained to choose $v = 0$ and the policymaker ignores the fundamental uncertainty; conversely, as

[6] An alternative interpretation, as in Hansen and Sargent (2001), is that v represents model misspecification: the policymaker does not know the true model driving financial stability and thus v reflects the perturbation of the true model from the model that is being relied upon.

$\theta \to 1$, the worst-case scenario becomes increasingly bad and the policymaker becomes more and more concerned about the ambiguity over v.

The policymaker's optimal choice of k, taking the actions of nature as given, is given by

$$k = \frac{\bar{b}x^* - \bar{b}v}{\bar{b}^2 + {\sigma_b}^2 + \lambda}. \tag{11}$$

Solving for the Nash equilibrium gives the robust control solution to k as

$$k = \frac{\frac{\theta}{(\theta-1)}\bar{b}x^*}{\frac{\theta}{(\theta-1)}\bar{b}^2 + {\sigma_b}^2 + \lambda}. \tag{12}$$

Inspecting this solution, it is clear that as $\theta \to \infty$ and fundamental uncertainty disappears, the robust control solution for k is the same as the Brainard solution given by equation (3) in the introduction. Note also that the equilibrium realization for v is

$$v = \frac{-x^*({\sigma_b}^2 + \lambda)}{\theta\bar{b}^2 + (\theta - 1)({\sigma_b}^2 + \lambda)} < 0. \tag{13}$$

The negative value of v means that the worst-case scenario, in this model, is one where there is a negative shock to financial stability. This is intuitive in the context of a macroprudential policymaker.

The main result of this section is that policy becomes more activist when there is greater fundamental uncertainty,[7]

$$\frac{dk}{d\theta} < 0. \tag{14}$$

Thus, in particular, as the solution coincides with Brainard when $\theta \to \infty$ and fundamental uncertainty disappears, the policymaker is more activist than in the case of Brainard when there is fundamental uncertainty (i.e., for $\theta \in (1, \infty)$).

The simple explanation for this finding is that increasing fundamental uncertainty makes the worst-case outcome, which the policymaker is preparing for, worse. As the worst-case outcome is one

[7]Recall that fundamental uncertainty increases as θ decreases.

Figure 2. Graphical Illustration of the Robust Control Model

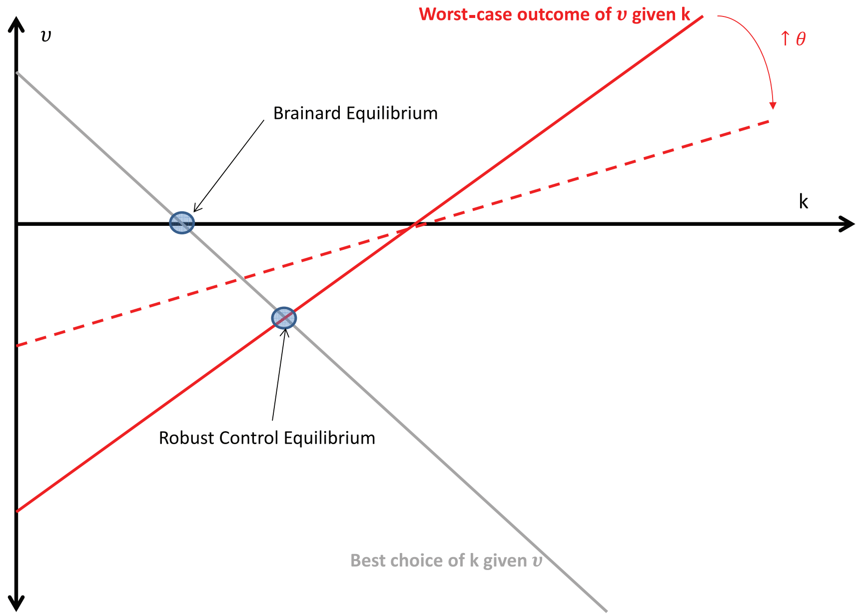

where there is a large negative shock to financial stability, this requires a tighter policy stance in response.

Figure 2 illustrates the intuition behind the robust control model graphically. The downward-sloping light grey line gives the best choice of k given v (given by equation (11)). The intuition for the downward slope is straightforward: a negative v implies a negative shock to financial stability and hence policy should be tightened to stabilize x. The upward-sloping line captures the "prepare for the worst" nature of the robust control problem: it is the value of v that gives the worst possible outcome to the policymaker given k and subject to the entropy constraint (given by equation (10)). This line is upward sloping because if loose policy is set, the worst-case outcome is a negative shock to financial stability. A robust policy choice is to set k where the two lines intersect: this means that the policymaker is choosing a policy which gives the best result if the very worst happens. In the standard Brainard problem in the introduction, $v = 0$, so the Brainard solution is equivalent to where the grey line crosses the k-axis. In the figure the intersection between the two lines is at a higher level of k than the Brainard point. Hence the robust control policy is more aggressive than Brainard would suggest. The intuition for this lies with the idea that more policy action is required if the most severe scenario—a negative financial stability shock—occurs, and the higher the fundamental uncertainty, the more severe is the worst case.[8] Increasing θ reduces fundamental uncertainty; this flattens the upward-sloping line[9] (shown by a dashed line in figure 2) and means that the worst-case outcome is not as bad, implying that policy does not need to be so activist.

3.2 Discussion

While a robust approach to policymaking usually suggests policy should be more active, it may say little about the direction in which policy should be more aggressive. To see this, consider the response of fiscal policy in the event of a severe recession. One extreme outcome may be high unemployment, negative growth, and deflation. Therefore, the robust response may be to embark upon aggressive fiscal expansion. However, an alternative scenario is that the recession puts sufficient strain on public finances to bring the government's solvency into question, leading to rising sovereign risk premia and additional stress on the financial system and the economy. So an alternative robust policy may be an aggressive tightening of the fiscal stance. Therefore, a policy designed to avoid either extreme outcome requires an aggressive shift in the policy stance, but the exact direction depends on which scenario policymakers are trying to avoid.

At face value, for the macroprudential context this lack of direction may seem less of an issue. Macroprudential policymakers are chiefly concerned by the extreme tail risks posed by financial crises.

---

[8]This is a modeling assumption. It is not the case that all robust control problems lead to a more aggressive policy, but it is a typical feature of this sort of model. See Onatski and Stock (2002) for a discussion in the context of monetary policy. Barlevy (2009) offers some counterexamples.

[9]Equation (10) can be rewritten as $k = \frac{v(\theta-1)+x^*}{b}$. Hence, nature's best response function intersects the k-axis at the point $k = \frac{x^*}{b}$ which is independent of θ. Therefore, any change in θ is a rotation of nature's best response line around the point $(\frac{x^*}{b}, 0)$, as illustrated in figure 2.

This asymmetry means that, at most points in time, a policymaker behaving in a robust way would be hawkish in setting her macroprudential tools. However, there is always a risk that under certain conditions severe negative outcomes can emerge if policy is set too tightly. In the spirit of Lucas (1987), one can imagine that the worst outcome of macroprudential policy is to lower the long-term trend growth rate. The question would then be whether an overly burdensome regulatory regime had a greater impact on long-term growth than financial crises.

Despite its popularity in the academic literature, a robust control approach can appear abstract when applied to practical policymaking. It has real-world applications nonetheless. For example, macroprudential policymakers often have a large number of potential indicators of risks to the financial system: a robust policymaking strategy would be to pay most attention to indicators that are signaling problems rather than focus on a measure of central tendency such as a simple average. Second, the principle of calibrating policy to prepare for severe outcomes is already embodied within many macroprudential frameworks via stress testing. Stress tests by definition provide a sense of economic outcomes if an extreme scenario emerges.

4. Learning about the Effects of Policy

A natural way to respond to uncertainty over a policy instrument is to attempt to learn about it. Research using evidence from other countries or from natural experiments, or by using calibrated theoretical models, can fill this gap. However, there is rarely a perfect substitute for using the instrument itself. Furthermore, in the macroprudential context, the framework for making and communicating changes to the tools is also new. Equivalent instruments have been used for microprudential purposes (such as bank capital requirements), but the signaling value (see section 5) and systemwide consequences of a macroprudential action may lead to a different impact, limiting what can be learned from previous uses of the tools.

In this section, we adapt the Brainard model to allow the policymaker to learn about the impact of her tool by observing what happens when she moves it. There is a trade-off: being active with a policy tool today leads to additional volatility, but by observing the tool's effect the policymaker will be less uncertain in future. This means that the policymaker has an incentive to be more active initially. While it may not be possible to directly measure or observe financial stability, so long as the policymaker can observe signals which convey some information about risks to the stability of the financial system (e.g., prices in financial markets), they will be able to learn about the effectiveness of their tool.

4.1 Modeling Framework

Here we lay out a two-period ($t = 1, 2$) version of the model presented in the introduction. Let

$$W_t := -\frac{1}{2}\mathbb{E}[(x_t - x^*)^2 + \lambda (k_t)^2]. \tag{15}$$

And suppose the policymaker's objective in period 1 is now to maximize

$$W_1 + \delta W_2 : \delta > 0. \tag{16}$$

As before, the policy instrument k_t has an uncertain impact on x_t described by

$$x_t = bk_t + u_t. \tag{17}$$

The parameter δ determines how much the policymaker values future periods. The other parameters share the interpretation they have in the introduction. For tractability we suppose that b and u_t are normally and independently distributed with the following initial priors: $b \sim N\left(\bar{b}, \sigma_b^2\right)$ and $u_t \sim N\left(0, \sigma_u^2\right)$.

The crucial assumption underpinning this two-period model is that by setting policy in period 1 the policymaker learns something about b in period 2. To embed this feature in the model, we have to specify the information available to the policymaker in each period. We assume that in period 1, when k_1 is chosen, the policymaker has the same information as in the Brainard model. At the end of period 1 the policymaker observes x_1, which, given k_1, provides an additional signal over b, reducing the uncertainty around the impact of the tool. However, since u_1 is not observed, the policymaker cannot perfectly distinguish between movements in x_1 that are due to k_1 or due to u_1. This results in a simple signal-extraction problem where the policymaker updates the expectation she has of b and uncertainty that surrounds that expectation (as shown in the appendix).

$$b|x_1 \sim N\left(\bar{b} + \frac{\sigma_b^2 k_1}{(k_1)^2\sigma_b^2 + \sigma_u^2}(x_1 - \bar{b}k_1),\ \frac{\sigma_b^2\sigma_u^2}{(k_1)^2\sigma_b^2 + \sigma_u^2}\right) \tag{18}$$

Equation (18) implies that the updated variance in period 2 is strictly less than in period 1 when $k_1 > 0$ ($Var(b|x_1) < \sigma_b^2$); thus uncertainty is reduced between the two periods. Moreover, increasing k_1 reduces $Var(b|x_1)$: the policymaker becomes more certain in the second period the more she acts in the first. The problem in period 2 is almost identical to the static Brainard model of the introduction, and the policymaker tries to maximize W_2 conditional upon beliefs about the mean and variance about b. The only difference is that now those beliefs are given by $\mathbb{E}(b|x_1)$ and $Var(b|x_1)$, which depend on the choice of k_1 in the previous period.

$$k_2 = \frac{\mathbb{E}(b|x_1)x^*}{\mathbb{E}(b|x_1)^2 + Var(b|x_1) + \lambda} \tag{19}$$

As a result, the policymaker faces a trade-off: actively using a tool in period 1 may not be immediately desirable due to the Brainard result from the introduction, but being active will make the policymaker better off in the future via lower uncertainty. To illustrate how this trade-off manifests and interacts with other parameters in the model, in figure 3 we present the optimal choice of k_1 across different levels of δ and ${\sigma_u}^2$.[10] These values of k_1 are presented in comparison to the optimal choice if the policymaker faced no uncertainty about b and if policy was set using the Brainard model.

Consider the left-hand panel of figure 3 first. We see the result that increasing δ monotonically increases k_1. The intuition for this is straightforward: as δ increases, the policymaker puts more weight on future outcomes. Hence, learning about how policy works in period 2 becomes a greater priority leading to a more activist policy. If $\delta = 0$

[10] It is clear from equations (18) and (19) that the policymaker's maximization problem can be set up in terms of k_1 only. This results in a highly non-linear function which is solved numerically using Monte Carlo methods.

Figure 3. Policy Choice with Learning

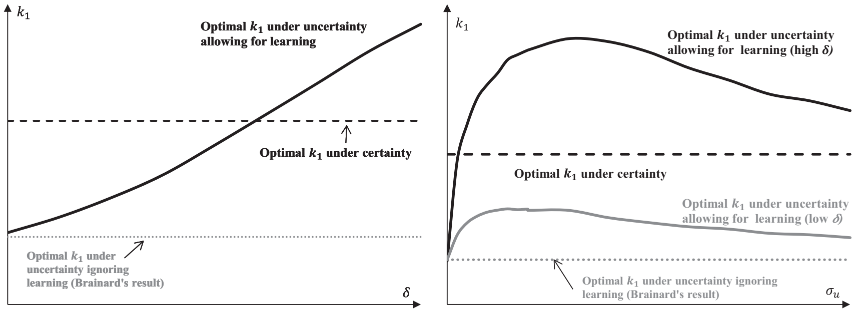

(i.e., there is no learning motive), the best choice for k_1 is simply in line with the Brainard model. For high values of δ (greater than 1) we can see that k_1 can be greater than it would be under certainty about b.[11] It is also possible to show, as we do in the appendix, that if the policymaker only cares about the second period (as they would when $\delta \to \infty$), then it is optimal for $k_1 \to \infty$ when there is no motive to stabilize k.

The right-hand panel of figure 3 considers how policy varies with ${\sigma_u}^2$. The irrelevance of uncertainty over the state of the economy for the policy choice does not extend to the learning case; instead we see a hump-shaped pattern with respect to ${\sigma_u}^2$. This is because how much can be learned about b by altering k_1 has a non-linear relationship with the uncertainty over u_1. If ${\sigma_u}^2$ is zero, there is nothing more that can be learned from altering k_1, as b is observed perfectly at the end of period 1 and thus the optimal k_1 is the same as Brainard. As ${\sigma_u}^2$ rises from zero, the uncertainty over b in period 2 increases, hence the policymaker chooses to be more active to learn more about the coefficient. However, as ${\sigma_u}^2$ becomes increasingly large, the ability of x_1 to be informative about b diminishes and the policymaker becomes less able to learn and is therefore less active.

[11] In reality policymakers will need to make decisions over many periods, thus the loss at period 2 could be thought of as a reduced-form way of capturing the continuation value of the problem. This would suggest that a large δ, potentially greater than 1, is appropriate.

At the limit ${\sigma_u}^2 \to \infty$, there is nothing that can be learned from x_1 and policy reverts to the Brainard solution.[12]

4.2 Discussion

The desire to learn is intimately related to how uncertain the policymaker thinks they will be in future. From today's perspective the macroprudential toolkit is largely untested. This is in contrast to the current state of monetary policy which can draw upon decades of academic research into its effects, substantial policy experience, and data sets that contain the necessary variation. In the macroprudential context, the extant empirical literature relies on difference-in-differences based empirical strategies, which can assess the impact of a small number of quasi-natural experiments but cannot account for differing macroeconomic climates or general equilibrium consequences. However, in the future, research into macroprudential policy will eventually come on stream and contribute to developing a better understanding of how the tools operate. In the meantime, policymakers today have to decide on the extent to which the current state of the literature can substitute for seeing the tool in action. Furthermore, using a policy generates both operational experience and the variation required to conduct detailed empirical studies.

The nature of the financial cycle means that the ability of policymakers to learn about the effectiveness of macroprudential policy differs substantially from the case of monetary policy. On the one hand, the relatively infrequent occurrence of financial crises means that, over time, there are relatively few events to learn from. On the other hand, the length of financial cycles means that, if policymakers believe that they can readily observe financial stability, there could be more opportunity to learn about the impact of policy tools ahead of crises.

The immediate practical issue with using a tool to learn is one of communication and political sensitivity. A macroprudential policymaker would probably struggle to articulate to banks that it was forcing them to raise more capital in order to determine the economic consequences. Deliberately experimenting with the economy with no other goal in mind is highly inadvisable and should run counter to

[12]These results for σ_u^2 are shown formally in the appendix.

macroprudential policy objectives. However, it is not necessary to implement policy that is harmful in order to learn. Though policy should not be set with wild abandon, the learning mechanism provides an argument to suggest that there are benefits from not being too cautious.

5. Policy Transparency and Private-Sector Uncertainty

It is not only policymakers that face uncertainty with regards to macroprudential policymaking. The public is also uncertain about risks to financial stability, the impact of macroprudential policy, and how the policymaker herself will behave. Furthermore, the policymaker may have informational advantages over the public through access to confidential regulatory data, supervisory intelligence, and the results of stress tests. Thus a macroprudential policymaker needs to consider how her policy actions interact with the information available to the public.

In the models used in the introduction and the previous three sections, the private sector has responded to policy actions in a mechanical fashion. In this section, we allow for strategic interaction between the policymaker and the private sector. The setup moves beyond a simple stabilization problem into a framework where (private) risk-taking is an endogenous choice (albeit modeled in a stylized fashion). The transmission mechanism of policy is similar, as is the eventual objective of the policymaker, to that discussed in the previous sections. However, this framework allows us to explore how private-sector uncertainty affects the transmission mechanism and, hence, how policymakers should behave in response.

We assume that the policymaker is better informed about the stability of the financial system than the private sector. As such, the policy action is an additional signal for the private sector. Crucially, this signal is imperfect: we also assume that the private sector is not fully aware of the objectives of the policymaker and hence cannot perfectly distinguish actions caused by risks to financial stability from those caused by an aversion to crises. An example of this would be if banks did not know if the policymaker set a high countercyclical capital buffer rate because it was aware of a specific financial stability risk or simply had a high preference for avoiding a crisis.

5.1 Modeling Framework

The model takes the form of a Stackelberg game. The policymaker first chooses the level of her policy instrument, k; then, having observed k, the private sector chooses their desired level of risk-taking, x. The policymaker internalizes the impact of her policy choice on the private sector's response when choosing her desired level of k.

We assume the private sector chooses x in order to maximize the following objective:

$$max_x \mathbb{E}\left[\left(-\frac{1}{2}(x - x^+ + u)^2 + \alpha x k\right)|k\right], \tag{20}$$

where α is a parameter whose interpretation we will return to shortly. There are two stochastic variables in this setup: u and c. The former enters the private sector's loss function; the latter is not directly relevant for the private sector but affects the preferences of the policymaker. Both u and c are uncorrelated random variables with positive variance, and the private sector has priors $u \sim N(0, \sigma_u^2)$ and $c \sim N(0, \sigma_c^2)$. This functional form has similarities with the stabilization problem described in the previous sections, with three main differences. First, x is now a choice variable of the private sector and there is a positive private target for risk-taking denoted $x^+ > 0$, capturing, for example, instability bias due to limited liability. Second, for given k, u interferes with the optimal level x that the private sector would prefer to target. A higher value of u implies that lower x is desired. This is consistent with the idea that u is an orthogonal shock that negatively affects financial stability and would make the private sector want to take less risk. The third difference is the cross-term $\alpha x k$. This term is missing from the models considered in the previous sections, and its presence allows for strategic interaction between the private sector and the policymaker. We assume $\alpha < 0$, which implies that, at the margin, tighter policy encourages less risk-taking by the private sector (this would be the case if, for instance, tighter capital requirements increase the private sector's skin in the game).[13]

---

[13] Our main result can also go through with α positive, which could occur for instance if tighter policy reduces the severity of the left tail of economic outcomes, encouraging risk-taking. In that case, equilibrium existence requires the assumption that $2\alpha\beta < 1$, where β is defined below.

The private sector takes k as given and observes its value before making its decision. The policymaker cares about u but has superior information about its realization; hence, k serves as a signal to the private sector about u. Formally, we can write the private sector's first-order condition as

$$x = \alpha k + x^+ - \mathbb{E}(u|k), \tag{21}$$

and so the private sector takes less risk as $\mathbb{E}(u|k)$ increases. The model is solved via the method of undetermined coefficients, and we postulate (and then verify) that the private sector's belief about u, having observed k, is linear in the choice of policy tool

$$\mathbb{E}(u|k) = \gamma_0 + \gamma_1 k. \tag{22}$$

The policymaker's problem is

$$max_k \left[-\frac{1}{2}(k - u - c)^2 + \beta x(k)k \right], \tag{23}$$

where from above, $x(k) = (\alpha - \gamma_1)k - \gamma_0 + x^+$. We also assume that the policymaker has full information (i.e., u and c are known). Again, we can think of a positive realization of u as a negative shock to financial stability which warrants tighter policy. In addition to u, the policymaker also has her own preferred target level for k, denoted c, which does not directly influence (and is unknown by) the private sector. Furthermore, we have a second cross-term: βxk. We assume that $\beta > 0$ to capture the idea that, at the margin, more risk-taking by the private sector makes tighter policy more desirable. Last, note that the policymaker internalizes that her decision over k will affect x, hence x enters the optimization problem as a function of policy.

These assumptions over preferences allow us to have a short exposition of the main mechanism we wish to highlight. The working paper version of this paper contains a larger model where preferences are more general (although still linear quadratic) and generates the same main result. The policymaker's problem has a unique maximum whenever

$$-\frac{1}{2} + \beta(\alpha - \gamma_1) < 0. \tag{24}$$

In the appendix we show that this condition is satisfied in equilibrium under the given parameter assumptions. Moreover, in equilibrium, $\gamma_1 > 0$. On this basis, taking the policymaker's first-order condition and solving for k gives

$$k = \frac{u + c + \beta(x^+ - \gamma_0)}{(1 - 2\beta(\alpha - \gamma_1))}. \tag{25}$$

From the perspective of the private sector, k is then normally distributed, and we can then solve for $\mathbb{E}(u|k)$ using the standard signal-extraction formula for normal distributions:

$$\mathbb{E}(u|k) = \mathbb{E}(u) + \frac{Cov(u,k)}{Var(k)}(k - \mathbb{E}(k)). \tag{26}$$

In the appendix we show that this is of the same form as equation (22), verifying the form of the postulated solution. We further show through equating coefficients that $\gamma_0 = \frac{-s\beta x^+}{1-\beta s}$ and

$$\gamma_1 = \frac{(1 - 2\beta\alpha)\, s}{1 - 2\beta s} > 0, \tag{27}$$

where $s = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}$ is a function of the signal-to-noise ratio, and equation (27) implies that we must also have $2\beta s < 1$. The private sector cares about u, but c, the policymaker's personal preference for tighter policy, does not influence its behavior. Hence, k is a useful signal to the extent that it provides information about u, but uncertainty about c adds noise to this signal. In particular, the greater the public-sector transparency about c, the lower σ_c^2 will be and the greater will be s. The greater the signal-to-noise ratio, the more informative k is as a signal about u. This is apparent upon inspecting equation (27), where we can see that $\frac{d\gamma_1}{ds} > 0$, consistent with the intuition that the sensitivity of $\mathbb{E}(u|k)$ to k is increasing in the signal-to-noise ratio.

In the appendix we also show that the policymaker's optimal choice is given by

$$k = \frac{1 - 2\beta s}{1 - 2\beta\alpha}\left(u + c + \frac{\beta x^+}{1 - \beta s}\right). \tag{28}$$

The main result of the model is then that the greater the signal-to-noise ratio, and hence the lower the private-sector uncertainty about policy preferences, the less active the policymaker is in responding to u. This is clear upon inspecting equation (28). The mechanism at work runs through the fact that the sensitivity of x to k is given by $\alpha - \gamma_1$. Recall that we assume that α is negative, meaning that absent any signaling channel, the private sector will take less risk as k is increased. The signaling channel amplifies the response of private-sector risk-taking to policy. When signaling is present, as s increases k becomes a more informative signal about the realization of u and consequently more weight is placed on the signal, resulting in γ_1 being larger. This in turn increases the responsiveness of x to k as the private sector view a larger part of the rise in k being due to a less stable financial system. Consequently, with the private sector taking less risk in response to a given policy choice, k does not need to be set as high, resulting in a lower sensitivity of k to u.

5.2 Discussion

This result highlights the importance of a policymaker ensuring that private agents understand why a policy action is taken, in order to maximize the signaling value of a policy decision. In practice, a macroprudential policymaker has a variety of ways of communicating to the private sector beyond her actions. The explanation of macroprudential policy decisions given alongside the action can help the public learn and extract an informative signal about the policymaker's views on financial stability risks. Alternatively, the macroprudential policymaker may signal their likely future actions in response to a buildup in risks by clarifying their "reaction function." The publication of core financial stability indicators may provide some clarity to the policymaker's reaction function. Communicating likely future actions may affect expectations about the evolution of the economy and hence influence behavior today.

Changes in private-sector uncertainty have affected the transmission mechanism of other policy tools. As noted by Beechey (2008), the news on May 6, 1997 that a Monetary Policy Committee would be created and the Bank of England would be granted operational independence with an explicit inflation target was followed by a substantial reduction in market-implied short- and long-term inflation expectations, arguably because this reduced uncertainty about the objectives of monetary policy. The forward rates of inflation compensation five to ten years ahead fell around 35 basis points on the day of the announcement, with a further decline over the following days. This in turn affected the transmission mechanism of monetary policy.

The obvious difference, however, between macroprudential and monetary policymakers is that in the context of the latter a precise, measurable target is available. Furthermore, the private sector has had more chance to learn about how monetary policy is conducted in practice and, correspondingly, policymakers' reaction functions.

6. Conclusion

As this paper has highlighted, there are several types of uncertainty and multiple channels through which uncertainty can affect policymaking. The well-known result of Brainard (1967) that policy should be more cautious in the presence of uncertainty does not hold in general, particularly for specific examples that are relevant to macroprudential policymakers. If anything, the need to learn about these relatively untested tools and the focus on avoiding tail risks speak to more active policymaking. Moreover, private-sector uncertainty over the financial stability objectives, preferences, and reaction function may diminish the potency of the signaling impact of macroprudential policy, requiring more active policy or communication. One limitation of the analysis here is that it takes place in context of static or two-period models. Hence it is a challenge to disentangle tighter steady-state macroprudential policy in response to increased risks from a more active response to economic shocks as they emerge. This would be an interesting area for future analysis.

Appendix 1. Asymmetric Objective Function

PROPOSITION 1. *The optimal choice of policy under asymmetry is given by the* $\tilde{k}$ *that satisfies*

$$\frac{-1}{a}\left(exp\left\{ax^* - a\bar{b}\tilde{k} + \tfrac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right\}\left(-\bar{b} + a\tilde{k}\sigma_b^2\right) + \bar{b}\right) - \lambda\tilde{k} = 0. \tag{29}$$

Further, $\tilde{k}\epsilon\left(0, \frac{\bar{b}}{a\sigma_b^2}\right)$.

Proof. The policy objective under asymmetry is W, as given by equation (6). The first derivative of W w.r.t. k is given by

$$-\left\{\frac{1}{a}\left[exp\left\{ax^* - a\bar{b}k + \tfrac{a^2}{2}\left(k^2\sigma_b^2 + \sigma_u^2\right)\right\}\left(-\bar{b} + ak\sigma_b^2\right) + \bar{b}\right] + \lambda k\right\}.$$

The second derivative of W w.r.t. k is given by

$$-\left\{\frac{1}{a^2}\left[e^{\left(ax^* - a\bar{b}k + \frac{a^2}{2}\left(k^2\sigma_b^2 + \sigma_u^2\right)\right)}\left\{-a\bar{b} + a^2k\sigma_b^2\right\}^2 \right.\right.$$
$$\left.\left. + e^{\left(ax^* - a\bar{b}k + \frac{a^2}{2}\left(k^2\sigma_b^2 + \sigma_u^2\right)\right)} a^2\sigma_b^2\right] + \lambda\right\}.$$

As this is always negative, W is a strictly concave function and hence equating the first-order condition (FOC) to 0 gives the unique global maximum.

Turning to the bounds on $\tilde{k}$, the first derivative, evaluated at $k = 0$, is

$$\frac{\bar{b}}{a}\left[exp\left(ax^* + \tfrac{a^2\sigma_u^2}{2}\right) - 1\right] > 0,$$

as all parameters are positive. Hence, as W is strictly concave, the derivative is positive iff $k < \tilde{k}$ and so we must have $0 < \tilde{k}$. Turning to the upper bound, if $k \geq \frac{\bar{b}}{a\sigma_b^2}$ then $-\bar{b} + ak\sigma_b^2 \geq 0$ and $\frac{dW}{dk} < 0$. However, as W is strictly concave, the derivative is negative iff $k > \tilde{k}$ and so we must have $\tilde{k} < \frac{\bar{b}}{a\sigma_b^2}$.

LEMMA 1. *Let $\tilde{k}$ be the policymaker's solution. Then $\frac{d\tilde{k}}{da} > 0$ iff the following expression is positive:*

$$e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)}{2}\right)} \left\{\frac{\bar{b}}{a} - \tilde{k}\sigma_b^2\right\} \left[\left(x^* - \bar{b}\tilde{k}\right) + a\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right]$$
$$+ \frac{\bar{b}}{a^2}\left[1 - e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)}{2}\right)}\right]. \tag{30}$$

Proof. The optimal policy choice $\tilde{k}$ is defined implicitly by equation (29). Let $g(\tilde{k}(a), a)$ be the left-hand side, then $g(\tilde{k}(a), a) \equiv 0$. By the implicit function theorem, $\frac{d\tilde{k}}{da} = \frac{\partial g(\tilde{k}(a),a)}{\partial a} / - \frac{\partial g(\tilde{k}(a),a)}{\partial \tilde{k}}$. As the objective function is strictly concave, $\frac{\partial g(\tilde{k}(a),a)}{\partial \tilde{k}} < 0$, as $g()$ is the first derivative of the objective function. Thus $\frac{d\tilde{k}}{da} > 0$ iff $\frac{\partial g(\tilde{k}(a),a)}{\partial a} > 0$. Taking the derivative of $g()$ w.r.t. a gives the result.

We first establish a useful lemma before turning to the main proposition.

LEMMA 2. *The following relationships hold as the expected effectiveness of the policy tool, $\bar{b}$, becomes arbitrarily large:*

$$\lim_{\bar{b}\to\infty} \tilde{k} = 0 \tag{31}$$

$$\lim_{\bar{b}\to\infty} \frac{\bar{b}}{a^2}\left[1 - exp\left(ax^* - a\bar{b}\tilde{k} + \tfrac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)\right] = 0 \tag{32}$$

$$\lim_{\bar{b}\to\infty} \bar{b}\tilde{k} = x^* + \frac{a\sigma_u^2}{2}. \tag{33}$$

Proof. To prove result (31), we establish that $\bar{b}\tilde{k}$ is bounded. We have established that $\bar{b}\tilde{k} > 0$, so it is sufficient to show that $\exists S \geq 0 : \bar{b}\tilde{k} \leq S \ \ \forall \bar{b} > 0$. If S exists, it follows that result (31) holds.[14]

[14]To see this, note that if $\exists S \in \mathbb{R}_+ : 0 < \bar{b}\tilde{k} \leq S \ \forall \bar{b}$, then $0 < \tilde{k} \leq \frac{S}{\bar{b}}$ and hence as $\lim_{\bar{b}\to\infty} \frac{S}{\bar{b}} = 0$, by the sandwich theorem we have $\lim_{\bar{b}\to\infty} \tilde{k} = 0$.

To show that S exists, we suppose for a contradiction that it does not. Then $\bar{b}\tilde{k} \to \infty$ as $\bar{b} \to \infty$. We rewrite FOC (29) as

$$\frac{-1}{a}\left(e^{\left(ax^*-a\bar{b}\tilde{k}+\frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2+\sigma_u^2\right)\right)}a\tilde{k}\sigma_b^2 \right.$$
$$\left. + \bar{b}\left\{1-e^{\left(ax^*-a\bar{b}\tilde{k}+\frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2+\sigma_u^2\right)\right)}\right\}\right) - \lambda\tilde{k} = 0. \tag{34}$$

To construct the contradiction, we first establish, under the supposition $\bar{b}\tilde{k} \to \infty$ as $\bar{b} \to \infty$, that

$$exp\left(ax^* - a\bar{b}\tilde{k} + \tfrac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right) \to 0 \text{ as } \bar{b} \to \infty. \tag{35}$$

To do this, we make use of the inequality $\tilde{k} \leq \frac{\bar{b}}{a\sigma_b^2}$, and thus the exponent in equation (35) must adhere to the following inequality:

$$ax^* - a\bar{b}\tilde{k} + \tfrac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right) \leq ax^* + \frac{a^2\sigma_u^2}{2} - a\bar{b}\tilde{k} + \frac{a^2\sigma_b^2}{2}\tilde{k}\left(\frac{\bar{b}}{a\sigma_b^2}\right)$$
$$= ax^* + \frac{a^2\sigma_u^2}{2} - \frac{a}{2}\bar{b}\tilde{k}. \tag{36}$$

As $\bar{b}\tilde{k}$ is unbounded by assumption, the right-hand side of equation (36) tends to $-\infty$ as $\bar{b} \to \infty$ and hence so too does the left-hand side. Thus, under our supposition, equation (35) holds.

Now, equation (35) implies $\bar{b}\left\{1-exp\left(ax^*-a\bar{b}\tilde{k}+\frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2+\sigma_u^2\right)\right)\right\} \to \infty$.

Further, $exp\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)a\tilde{k}\sigma_b^2 > 0$ and $\tilde{k} > 0$. Thus, the left-hand side of equation (34) tends to $-\infty$ as $\bar{b} \to \infty$, a contradiction, as the FOC no longer holds. Hence, we must have $\bar{b}\tilde{k}$ bounded, and thus result (31) holds.

Turning to the proof of result (32), rearranging equation (34), we have

$$\frac{\bar{b}}{a^2}\left[1 - e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)}\right]$$
$$= -\frac{\tilde{k}}{a}\left(\lambda + \sigma_b^2 e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)}\right).$$

The right-hand side $\to 0$ as $\bar{b} \to \infty$ because $\tilde{k} \to 0$ whilst, with $\tilde{k}$, $\bar{b}\tilde{k}$ bounded, the exponential term is also bounded.

Thus, we must have $\underset{\bar{b}\to\infty}{lim} \frac{\bar{b}}{a^2}\left[1 - e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)}\right] = 0$, showing result (32).

Finally, we must also then have $\underset{\bar{b}\to\infty}{lim}\, ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right) = 0$. Given $\underset{\bar{b}\to\infty}{lim}\, \tilde{k} = 0$, it follows that $\underset{\bar{b}\to\infty}{lim}\, \bar{b}\tilde{k} = x^* + \frac{a\sigma_u^2}{2}$. This proves result (33).

Using the lemma, we prove the main proposition of this section.

PROPOSITION 2. *When $\bar{b}$ is sufficiently large, policy becomes more active under greater asymmetry:* $\frac{d\tilde{k}}{da} > 0$.

Proof. To prove the proposition, we establish that equation (30) is positive as $\bar{b} \to \infty$. From the prior lemma,

$$\underset{\bar{b}\to\infty}{lim} \frac{\bar{b}}{a^2}\left[1 - e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)}\right] = 0$$
$$\underset{\bar{b}\to\infty}{lim}\, e^{\left(ax^* - a\bar{b}\tilde{k} + \frac{a^2}{2}\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right)\right)} = 1.$$

Moreover,

$$\lim_{\bar{b}\to\infty}\left(x^* - \bar{b}\tilde{k}\right) + a\left(\tilde{k}^2\sigma_b^2 + \sigma_u^2\right) = \left(x^* - \left(x^* + \frac{a\sigma_u^2}{2}\right)\right) + a\sigma_u^2$$
$$= \frac{a\sigma_u^2}{2} > 0.$$

As $\lim_{\bar{b}\to\infty} \tilde{k} = 0$, it follows that $\frac{\bar{b}}{a} - \tilde{k}\sigma_b^2 \to \infty$ and hence (30)$\to \infty$ as $\bar{b} \to \infty$. Thus, for sufficiently large $\bar{b}$, $\frac{d\tilde{k}}{da} > 0$.

Appendix 2. Learning about the Effects of Policy

LEMMA 3. *For a given choice of k_1, the posterior mean and variance of b (having observed x_1) are given by*

$$b|x_1 \sim N\left(\bar{b} + \frac{{\sigma_b}^2 k_1}{(k_1)^2{\sigma_b}^2 + {\sigma_u}^2}(x_1 - \bar{b}k_1),\ \frac{{\sigma_b}^2{\sigma_u}^2}{(k_1)^2{\sigma_b}^2 + {\sigma_u}^2}\right). \quad (37)$$

Proof. Given the distributional assumptions on b, u and given that k_1 is non-stochastic as chosen by the policymaker, we have that $x_1 \sim N\left(\bar{b}k_1, (k_1)^2\sigma_b^2 + \sigma_u^2\right)$. Then, as b, x_1 are both normally distributed, we have that

$$b|x_1 \sim N\left(\bar{b} + \frac{Cov\left(b, x_1\right)}{Var\left(x_1\right)}\left(x_1 - \mathbb{E}\left(x_1\right)\right),\ \sigma_b^2 - \frac{\left(Cov\left(b, x_1\right)\right)^2}{Var\left(x_1\right)}\right). \quad (38)$$

Now, $Cov\left(b, x_1\right) = Cov\left(b, bk_1 + u\right) = k_1\sigma_b^2$; applying this to equation (38) and rearranging completes the proof.

LEMMA 4. *Suppose $\sigma_u^2 = 0$ or $\sigma_u^2 \to \infty$; then k_2 is independent of k_1 and the solution for k_1 is as in the static case.*

Proof. From equation (37), it is clear that when $\sigma_u^2 \to \infty$, $b|x_1 \sim N\left(\bar{b}, {\sigma_b}^2\right)$ and so k_2 is independent of k_1 and hence the solution for k_1 is as in the static Brainard case.

Now suppose $\sigma_u^2 = 0$. Then from equation (37) $b|x_1$ has zero variance, and so in the second period the value of b is known perfectly for any choice of $k_1 \neq 0$. Thus, both $b|x_1$ and $b^2|x_1$ are known perfectly in the second period. Hence k_2 is independent of k_1 and thus the optimal solution for k_1 will be as in the static Brainard case.

PROPOSITION 3. *When $\lambda = 0$, when the policymaker only cares about the second period's welfare, it is optimal to set k_1 arbitrarily large.*

Proof. The relevant objective of the policymaker is now $\mathbb{E}(W_2)$, where $\mathbb{E}(.)$ denotes the prior expectation of the policymaker before x_1 is realized. $\mathbb{E}(W_2)$ has no closed-form solution; however, using the law of iterated expectations we can write $\mathbb{E}(W_2) = \mathbb{E}(\mathbb{E}(W_2 \mid x_1))$, where $\mathbb{E}(. \mid x_1)$ is the posterior expectation of the policymaker once x_1 is realized. As shown below, $E(W_2 \mid x_1)$ can be written in closed form. Therefore, the logic of the proof will be that, for $\lambda = 0$, $\mathbb{E}(W_2 \mid x_1)$ is maximized when $k_1 \rightarrow \infty$ for *any* realization x_1. In turn, this implies that $\mathbb{E}(W_2)$ is maximized when $k_1 \rightarrow \infty$.

Once x_1 is realized, we know that the optimal choice of k_2 is given by

$$k_1 = \frac{\mathbb{E}(b \mid x_1)x^*}{\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda}.$$

Substituting this into $\mathbb{E}(W_2 \mid x_1)$ yields

$$\mathbb{E}(W_2 \mid x_1) = -\frac{1}{2}E\left[\left(\frac{b\mathbb{E}(b \mid x_1)x^*}{\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda} + u_2 - x^*\right)^2 \right.$$
$$\left. + \frac{\lambda\mathbb{E}(b \mid x_1)^2(x^*)^2}{(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda)^2} \mid x_1\right].$$

Expanding the brackets gives $\mathbb{E}(W_2 \mid x_1) =$

$$-\frac{\mathbb{E}(b^2 \mid x_1)\mathbb{E}(b \mid x_1)^2(x^*)^2}{2(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda)^2} + \frac{\mathbb{E}(b \mid x_1)^2(x^*)^2}{\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda}$$
$$-\frac{\lambda\mathbb{E}(b \mid x_1)^2(x^*)^2}{2(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda)^2} + t.i.p.,$$

where $t.i.p.$ denotes terms independent of policy. We can then factorize to get

$$\mathbb{E}(W_2 \mid x_1) = -\frac{\mathbb{E}(b \mid x_1)^2 (x^*)^2}{2(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda)^2} \big(\mathbb{E}(b^2 \mid x_1) - 2(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda) + \lambda\big) + t.i.p.$$

Using the fact that $\mathbb{E}(b^2 \mid x_1) = \mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1)$, the expression simplifies to

$$\mathbb{E}(W_2 \mid x_1) = \frac{\mathbb{E}(b \mid x_1)^2 (x^*)^2}{2(\mathbb{E}(b \mid x_1)^2 + Var(b \mid x_1) + \lambda)} + t.i.p.$$
$$= \frac{(x^*)^2}{2\left(1 + \frac{Var(b|x_1)+\lambda}{\mathbb{E}(b|x_1)^2}\right)} + t.i.p.$$

Now when $\lambda = 0$, the maximum value that $\mathbb{E}(W_2 \mid x_1)$ takes is $\frac{(x^*)^2}{2}$, as all terms are positive. This occurs when $Var(b \mid x_1) = 0$ for any given realization of x_1 (except the knife-edge case where $\mathbb{E}(b \mid x_1)^2 = 0$ and $\mathbb{E}(W_2 \mid x_1)$ is independent of the policy choice). As described in the main text, $Var(b \mid x_1) = \sigma_b^2 \sigma_u^2 / ((k_1)^2 \sigma_b^2 + \sigma_u^2)$; therefore $Var(b \mid x_1) \to 0$ as $k_1 \to \infty$. Hence, $\mathbb{E}(W_2 \mid x_1)$ tends to its maximum when $k_1 \to \infty$. By the logic above, this implies that $\mathbb{E}(W_2)$ also attains its maximum. This completes the proof.

Appendix 3. Policy Transparency and Private-Sector Uncertainty

PROPOSITION 4. *Given our assumptions over parameters, a solution exists where private-sector beliefs are of the form $E(u|k) = \gamma_0 + \gamma_1 k$. Moreover, $\gamma_1 > 0$ and $\frac{dk}{du}$ is decreasing in $s = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}$. Further, the optimal policy choice under private-sector uncertainty is given by*

$$k = \frac{1 - 2\beta s}{1 - 2\beta\alpha} \left(u + c + \frac{\beta x^+}{1 - \beta s} \right).$$

Proof. Suppose that the private sector has a posterior expectation of the form $\mathbb{E}(u|k) = \gamma_0 + \gamma_1 k$. As shown in the text, so long as

$-\frac{1}{2} + \beta(\alpha - \gamma_1) < 0$, the optimality condition of the policymaker is then given by

$$k = \frac{u + c - \beta(\gamma_0 - x^+)}{(1 - 2\beta(\alpha - \gamma_1))}.$$

As u, c are normally distributed, so too is k and so, using the formula for conditional normal distributions,

$$\mathbb{E}(u|k) = \mathbb{E}(u) + \frac{cov(u,k)}{Var(k)}(k - \mathbb{E}(k)).$$

Now, with u, c uncorrelated,

$$\frac{cov(u,k)}{Var(k)} = \frac{\frac{\sigma_u^2}{(1-2\beta(\alpha-\gamma_1))}}{\frac{\sigma_u^2+\sigma_c^2}{(1-2\beta(\alpha-\gamma_1))^2}} = (1 - 2\beta(\alpha - \gamma_1))\frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}$$

$$\mathbb{E}(k) = \frac{-\beta(\gamma_0 - x^+)}{(1 - 2\beta(\alpha - \gamma_1))}.$$

Thus

$$\begin{aligned}\mathbb{E}(u|k) &= (1 - 2\beta(\alpha - \gamma_1))\frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}\left(k + \frac{\beta(\gamma_0 - x^+)}{(1 - 2\beta(\alpha - \gamma_1))}\right) \\ &= \frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}\left((1 - 2\beta(\alpha - \gamma_1))k + \beta(\gamma_0 - x^+)\right).\end{aligned}$$

This confirms that the postulated linear form for the private sector's posterior belief, $\mathbb{E}(u|k) = \gamma_0 + \gamma_1 k$, is correct. Equating the coefficients on k gives

$$\gamma_1 = (1 - 2\beta(\alpha - \gamma_1))\frac{\sigma_u^2}{\sigma_u^2 + \sigma_c^2}.$$

Define $s = \frac{\sigma_u^2}{\sigma_u^2+\sigma_c^2}$, where s is a function of the signal-to-noise ratio, and $s \in (0, 1)$. Rearranging this equation gives

$$\gamma_1 (1 - 2\beta s) = (1 - 2\beta\alpha) s.$$

As $2\beta\alpha < 0 < 1$, the right-hand side is non-zero, and thus we must have $\gamma_1 \neq 0$, $(1 - 2\beta s) \neq 0$, giving

$$\gamma_1 = \frac{(1 - 2\beta\alpha)\, s}{1 - 2\beta s}.$$

Now,

$$\begin{aligned} 1 - 2\beta(\alpha - \gamma_1) &= \frac{(1 - 2\beta s) + 2\beta\,(1 - 2\beta\alpha)\, s - 2\beta\alpha\,(1 - 2\beta s)}{1 - 2\beta s} \\ &= \frac{(1 - 2\beta\alpha)\,[(1 - 2\beta s) + 2\beta s]}{1 - 2\beta s}. \end{aligned}$$

Tidying gives

$$1 - 2\beta(\alpha - \gamma_1) = \frac{(1 - 2\beta\alpha)}{1 - 2\beta s}.$$

From equation (24), we have a valid equilibrium iff $1-2\beta(\alpha-\gamma_1) > 0$, and hence iff $\frac{(1-2\beta\alpha)}{1-2\beta s} > 0$. Thus, in any equilibrium in which private-sector conditional beliefs are of the linear form, we must have $\gamma_1 > 0$. Moreover, given this, with $2\alpha\beta < 0 < 1$ we must also have $2\beta s < 1$.

Now, the equilibrium policy choice will be

$$k = \frac{(u + c + \beta(x^+ - \gamma_0))\,(1 - 2\beta s)}{(1 - 2\beta\alpha)},$$

and so

$$\frac{dk}{du} = \frac{(1 - 2\beta s)}{(1 - 2\beta\alpha)} > 0,$$

which is decreasing in s. Finally, consider the coefficient γ_0. Equating the constant coefficient for $\mathbb{E}(u|k)$ gives $\gamma_0 = \beta s(\gamma_0 - x^+)$. Given that $2\beta s < 1$, $\beta s \neq 1$, and we must have $\gamma_0 = \frac{-s\beta x^+}{1-\beta s}$, and $x^+ - \gamma_0 = \frac{x^+}{1-s\beta}$.

References

Angeletos, G.-M., L. Iovino, and J. La'O. 2011. "Optimal Monetary Policy with Informational Frictions." NBER Working Paper No. 17525.

———. 2015. "Efficiency and Policy with Endogenous Learning." Mimeo.

Bahaj, S., and A. Foulis. 2016. "Macroprudential Policy under Uncertainty." Working Paper No. 584, Bank of England.

Barlevy, G. 2009. "Policymaking under Uncertainty: Gradualism and Robustness." *Economic Perspectives* 33 (Q II): 38–55.

Beechey, M. J. 2008. "Lowering the Anchor: How the Bank of England's Inflation-Targeting Policies Have Shaped Inflation Expectations and Perceptions of Inflation Risk." Finance and Economics Discussion Series No. 2008-44, Board of Governors of the Federal Reserve System.

Besley, T. 2001. "Political Institutions and Policy Competition." Technical Report, World Bank.

Bianchi, J., and E. G. Mendoza. 2010. "Overborrowing, Financial Crises and 'Macroprudential' Taxes." NBER Working Paper No. 16091.

Brainard, W. C. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review* 57 (2): 411–25.

Cogley, T., and T. J. Sargent. 2005. "The Conquest of US Inflation: Learning and Robustness to Model Uncertainty." *Review of Economic Dynamics* 8 (2): 528–63.

Hansen, L. P., and T. J. Sargent. 2001. "Acknowledging Misspecification in Macroeconomic Theory." *Review of Economic Dynamics* 4 (3): 519–35.

———. 2008. *Robustness.* Princeton University Press.

Jeanne, O., and A. Korinek. 2013. "Macroprudential Regulation versus Mopping Up after the Crash." NBER Working Paper No. 18675.

King, R. G. 1982. "Monetary Policy and the Information Content of Prices." *Journal of Political Economy* 90 (2): 247–79.

Korinek, A. 2011. "Systemic Risk-taking: Amplification Effects, Externalities, and Regulatory Responses." ECB Working Paper No. 1345.

Korinek, A., and A. Simsek. 2014. "Liquidity Trap and Excessive Leverage." NBER Working Paper No. 19970.

Landes, D. S. 1998. *The Wealth and Poverty of Nations: Why Some Are So Rich and Some So Poor.* London: Little, Brown and Company.

Lorenzoni, G. 2010. "Optimal Monetary Policy with Uncertain Fundamentals and Dispersed Information." *Review of Economic Studies* 77 (1): 305–38.

Lucas, R. E. 1987. *Models of Business Cycles.* Oxford: Blackwell.

Mukand, S., and D. Rodrik. 2002. "In Search of the Holy Grail: Policy Convergence, Experimentation, and Economic Performance." NBER Working Paper No. 9134.

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 (01): 85–110.

Orphanides, A., and J. C. Williams. 2007. "Inflation Targeting under Imperfect Knowledge." *Economic Review* (Federal Reserve Bank of San Francisco) 1: 1–23.

Patton, A. J., and A. Timmermann. 2007. "Testing Forecast Optimality Under Unknown Loss." *Journal of the American Statistical Association* 102 (480): 1172–84.

Svensson, L. E. 2002. "Monetary Policy and Real Stabilization." In *Rethinking Policy*, 261–312. Proceedings of an Economic Policy Symposium sponsored by the Federal Reserve Bank of Kansas City, held in Jackson Hole, Wyoming, August 29–31.

Svensson, L. E., and N. M. Williams. 2007. "Bayesian and Adaptive Optimal Policy under Model Uncertainty." NBER Working Paper No. 13414.

Varian, H. R. 1975. "A Bayesian Approach to Real Estate Assessment." In *Studies in Bayesian Econometrics and Statistics: In Honor of Leonard J. Savage*, ed. L. J. Savage, S. E. Fienberg, and A. Zellner, 195–208. Amsterdam: North-Holland.

Wald, A. 1950. *Statistical Decision Functions.* New York: John Wiley & Sons.

Weiss, L. 1982. "Information Aggregation and Policy." *Review of Economic Studies* 49 (1): 31–42.

Wieland, V. 2000. "Monetary Policy, Parameter Uncertainty and Optimal Learning." *Journal of Monetary Economics* 46 (1): 199–228.

———. 2006. "Monetary Policy and Uncertainty about the Natural Unemployment Rate: Brainard-Style Conservatism versus Experimental Activism." *B.E. Journal of Macroeconomics* 6 (1, March): 1–34.