# Policymaking under uncertainty: Gradualism and robustness

**Gadi Barlevy**

## Introduction and summary

Policymakers are often required to make decisions in the face of uncertainty. For example, they may lack the timely data needed to choose the most appropriate course of action at a given point in time. Alternatively, they may be unable to gauge whether the models they rely on to guide their decisions can account for all of the issues that are relevant to their decisions. These concerns obviously arise in the formulation of monetary policy, where the real-time data relevant for deciding on policy are limited and the macroeconomic models used to guide policy are at best crude simplifications. Not surprisingly, a long-standing practical question for monetary authorities concerns how to adjust their actions given the uncertainty they face.

The way economists typically model decision-making under uncertainty assumes that policymakers can assign probabilities to the various scenarios they might face. Given these probabilities, they can compute an *expected loss* for each policy—that is, the expected social cost of the outcomes implied by each policy. The presumption is that policymakers would prefer the policy associated with the smallest expected loss. One of the most influential works on monetary policy under uncertainty based on this approach is Brainard (1967). That paper considered a monetary authority trying to meet some target—for example, an inflation target or an output target. Brainard showed that under certain conditions, policymakers who face uncertainty about their economic environment should react less to news that they are likely to miss their target than policymakers who are fully informed about their environment. This prescription is often referred to as "gradual" policy. Over the years, gradualism has come to be viewed as synonymous with caution. After all, it seems intuitive that if policymakers are unsure about their environment, they should avoid reacting too much to whatever information they do receive,

given that they have only limited knowledge about the rest of the environment.

Although minimizing expected loss is a widely used criterion for choosing policy, in some situations it may be difficult for policymakers to assign expected losses to competing policy choices. This is because it is hard to assign probabilities to rare events that offer little historical precedent by which to judge their exact likelihood. For this reason, some economists have considered an alternative approach to policymaking in the face of uncertainty that does not require knowing the probability associated with all possible scenarios. This approach is largely inspired by work on robust control of systems in engineering. Like policymakers, engineers must deal with significant uncertainty— specifically, about the systems they design; thus they are equally concerned with how to account for such uncertainty in their models. Economic applications based on this approach are discussed in a recent book by Hansen and Sargent (2008). The policy recommendations that emerge from this alternative approach are referred to as *robust* policies, reflecting the fact that this approach favors policies that avoid large losses in all relevant scenarios, regardless of how likely they are. Interestingly, early applications of robust control to monetary policy seemed to contradict the gradualist prescription articulated by Brainard (1967), suggesting that policymakers facing uncertainty should respond more aggressively to news that they are likely to miss their target than policymakers facing no uncertainty. Examples of such findings include Sargent (1999), Giannoni (2002), and Onatski and Stock (2002);

*Gadi Barlevy is a senior economist and economic advisor in the Economic Research Department at the Federal Reserve Bank of Chicago. He thanks Marco Bassetto, Lars Hansen, Charles Manski, and Anna Paulson for their comments on earlier drafts.*

their results contradict the conventional wisdom based on Brainard (1967), which may help to explain the tepid response to robust control in some monetary policy circles.

In this article, I argue that aggressiveness is not a general feature of robust control and that the results from early work on robust monetary policy stem from particular features of the economic environments those papers studied.[1] Similarly, gradualism is not a generic feature of the traditional approach to dealing with uncertainty based on minimizing expected losses—a point Brainard (1967) himself was careful to make. I explain that the way policymakers should adjust their response to news that they are likely to miss their target depends on asymmetries in the uncertain environment in which they operate. As we shall see, both the traditional and robust control approaches dictate gradualism in the environment Brainard (1967) considered, while both dictate being more aggressive in other environments that capture elements in the more recent work on robust control.

My article is organized as follows. First, I review Brainard's (1967) original result. Then, I describe the robust control approach, including a discussion of some of its critiques. Next, I apply the robust control approach to a variant of Brainard's model and show that it implies a gradualist prescription for that environment—just as Brainard found when he derived optimal policy, assuming policymakers seek to minimize expected losses. Finally, using simple models that contain some of the features from the early work on robust monetary policy, I show why robustness can recommend aggressive policymaking under some conditions.

## Brainard's model and gradualism

I begin my analysis by reviewing Brainard's (1967) model. Brainard considered the problem of a policymaker who wants to target some variable that he can influence so that the variable will equal some prespecified level. For example, suppose the policymaker wants to maintain inflation at some target rate or steer output growth toward its natural rate. Various economic models suggest that monetary policy can affect these variables, at least over short horizons, but that other factors beyond the control of monetary authorities can also influence these variables. Meeting the desired target will thus require the monetary authority to intervene in a way that offsets changes in these factors. Brainard focused on the question of how this intervention should be conducted when the monetary authority is uncertain about the economic environment it faces but can assign probabilities to all possible scenarios

it could encounter and acts to minimize expected losses computed using these probabilities.

Formally, let us refer to the variable the monetary authority wants to target as $y$. Without loss of generality, we can assume the monetary authority wants to target this variable to equal zero. The variable $y$ is affected by a policy variable set by the policymaker, which I denote by $r$. In addition, $y$ is affected by some variable $x$ that the policymaker can observe prior to setting $r$. For simplicity, suppose $y$ depends on these two variables linearly; that is,

1)   $y = x - kr,$

where $k$ measures the effect of changes in $r$ on $y$ and is assumed to be positive. For example, $y$ could reflect inflation, and $r$ could reflect the short-term nominal interest rate set by the monetary authority. Equation 1 then implies that raising the nominal interest rate would lower inflation, but that inflation is also determined by other variables, as summarized by $x$. These variables could include shocks to productivity or the velocity of money. If $x$ rises, the monetary authority would simply have to set $r$ to equal $x/k$ to restore $y$ to its target level of 0.

To incorporate uncertainty into the policymaker's problem, suppose $y$ is also affected by random variables whose values the policymaker does not know, but whose distributions are known to him in advance. Thus, let us replace equation 1 with

2)   $y = x - (k + \varepsilon_k)\, r + \varepsilon_u,$

where $\varepsilon_k$ and $\varepsilon_u$ are independent random variables with means 0 and variances $\sigma_k^2$ and $\sigma_u^2$, respectively. This formulation assumes the policymaker is uncertain both about the effect of his policy, as captured by the $\varepsilon_k$ term that multiplies his choice of $r$, and about factors that directly affect $y$, as captured by the additive term $\varepsilon_u$. The optimal policy depends on how much loss the policymaker incurs from missing his target. Suppose the loss is quadratic in the deviation between the actual value of $y$ and its target—that is, the loss is equal to $y^2$. The policymaker will then choose $r$ so as to minimize his expected loss, that is, to solve

3)   $\min_r E\left[y^2\right] = \min_r E\left[\left(x - \left(k + \varepsilon_k\right)r + \varepsilon_u\right)^2\right].$

Brainard (1967) showed that the solution to this problem is given by

4)   $r = \dfrac{x}{k + \sigma_k^2/k}.$

Equation 4 is derived in appendix 1. Uncertainty about the effect of policy will lead the policymaker to attenuate his response to $x$ relative to the case where he knows the effect of $r$ on $y$ with certainty. In particular, when $\sigma_k^2 = 0$, the policymaker will set $r$ to undo the effect of $x$ by setting $r = x/k$. But when $\sigma_k^2 > 0$, the policy will not fully offset $x$. This is what is commonly referred to as gradualism: A policymaker who is unsure about the effect of his policy will react less to news about missing the target than he would if he were fully informed. By contrast, the degree of uncertainty about $\varepsilon_u$, as captured by $\sigma_u^2$, has no effect on policy, as evident from the fact that the optimal rule for $r$ in equation 4 is identical regardless of $\sigma_u^2$.

To understand this result, note that the expected loss in equation 3 is essentially the variance of $y$. Hence, a policy that leads $y$ to be more volatile will be considered undesirable given the objective of solving equation 3. From equation 2, the variance of $y$ is equal to $r^2\sigma_k^2 + \sigma_u^2$, which is increasing in the absolute value of $r$. An activist (aggressive) policy that uses $r$ to offset nonzero values of $x$ thus implies a more volatile outcome for $y$, while a passive (gradual) policy that sets $r = 0$ implies a less volatile outcome for $y$. This asymmetry introduces a bias toward less activist policies. Even though a less aggressive response to $x$ would cause the policymaker to miss the target on average, he is willing to do so in order to make $y$ less volatile. Absent this asymmetry, there would be no reason to attenuate policy. This explains why uncertainty in $\varepsilon_u$ has no effect on policy: It does not involve any asymmetry between being aggressive and being gradual, since neither affects volatility. Although Brainard (1967) was careful to point out that gradualism is an optimal reaction to certain types of uncertainty, his result is sometimes misleadingly cited as a general rule for coping with uncertainty regardless of its nature.

## The robust control approach

An important assumption underlying Brainard's (1967) analysis is that the policymaker knows the probability distribution of the variables that he is uncertain about. More recent work on policy under uncertainty is instead motivated by the notion that policymakers may not know what probability to attach to scenarios they are uncertain about. For example, there may not be enough historical data to infer the likelihood of various situations, especially those that have yet to be observed but remain theoretically possible. Without knowing these probabilities, it will be impossible to compute an expected loss for different policy choices as in equation 3. This necessitates an alternative criterion to judge what constitutes a good

policy. The robust control approach argues for picking the policy that minimizes the damage that the policy could possibly inflict—that is, the policy under which the largest possible loss across all potential outcomes is smaller than the largest possible loss under any alternative policy. A policy chosen under this criterion is known as a robust policy (or a robust strategy). Such a policy ensures the policymaker will not incur a bigger loss than the unavoidable bare minimum. This rule is often associated with Wald (1950, p. 18), who argued that this approach, known as the minimax (or minmax) rule, is "a reasonable solution of the decision problem when an a priori distribution ... does not exist or is unknown." For a discussion of economic applications of robust control as well as related references, refer to Hansen and Sargent (2008).

Before I consider the consequences of adopting the robust control approach for choosing a target as in Brainard (1967), I first consider an example of an application of robust control both to help illustrate what it means for a policy to be robust in this manner and to discuss some of the critiques of this approach. The example is known as the "lost in a forest" problem, which was first posed by Bellman (1956) and which spawned a subsequent literature that is surveyed in Finch and Wetzel (2004).[2] Although this example differs in several respects from typical applications of robust control in economics, it remains an instructive introduction to robust control. I will point out some of these differences throughout my discussion when relevant.

The lost in a forest problem can be described as follows. A hiker treks into a dense forest. He starts his trip from the main road that cuts through the forest, and he travels in a straight line for one mile into the forest. He then lies down to take a nap, but when he wakes up he realizes he forgot which direction he came from. He wishes to return to the road—not necessarily to the point where he started, but anywhere on the road where he can flag down a car and head back to town. He would like to do so using the shortest possible route, which if he knew the location of his starting point would be exactly one mile. But he does not know where the road lies, and because the forest is dense with trees, he cannot see the road from afar. So, he must physically reach the road in order to find it. What strategy should he follow in searching for the road? A geometric description of the problem is provided in box 1, although these details are not essential for following the remainder of this discussion.

Solving this problem requires establishing a criterion by which a strategy can qualify as "best" among all possible strategies. In principle, if the hiker knew his propensity to lie down in a particular orientation

relative to the direction he travelled from, he could assign a probability that his starting point could be found in any given direction. In that case, an obvious candidate for the optimal strategy is the one that minimizes the expected distance to reach some point on the road. But most people would be unlikely to know their likelihood of lying down in any particular direction or the odds they don't turn in their sleep. While it might seem tempting to simply treat all locations as equally likely, this effectively amounts to making assumptions on just these likelihoods. It is therefore arguable that we cannot assign probabilities that the starting point lies in any particular direction, implying we cannot calculate an expected travel distance for each strategy and choose an optimum. As an alternative criterion, Bellman (1956) proposed choosing the strategy that minimizes the amount of walking required to ensure reaching the road regardless of where it is located. That is, for any strategy, we can compute the longest distance one would have to walk to make sure he reaches the main road regardless of where it is located. We then pick the strategy for which this distance is shortest. This rule ensures we do not have to walk any more than is absolutely necessary to reach the road. While other criteria have been proposed for the lost in a forest problem, many have found the criterion of walking no more than is absolutely necessary to be intuitively appealing. But this is precisely the robust control approach. The worst-case scenario for any search strategy involves exhaustively searching through every wrong location before reaching the true location. Bellman's suggestion thus amounts to using the strategy whose worst-case scenario requires less walking than the worst-case scenario of any other strategy. In other words, the "best" strategy is the one that minimizes the amount of walking needed to run through the gamut of all possible locations for the hiker's original starting point.

Although Bellman (1956) first proposed this rule as a way of solving the lost in a forest problem, it was Isbell (1957) who derived the strategy that meets this criterion. His solution is presented in box 1. The hiker can ensure he finds the road by walking out one mile and, if he doesn't reach the road, continue walking along the circle of radius one mile around where he woke up. While this strategy ensures finding the road eventually, it turns out that deviating from this scheme in a particular way still ensures finding the road eventually, but with less walking.

Note that in the lost in a forest problem, the set of possibilities the hiker must consider to compute the worst-case scenario is an objective feature of the environment: The main road must lie somewhere along a circle of radius one mile around where the hiker fell asleep (we just don't know exactly where). By contrast, in most economic applications, the region that a decision-maker is uncertain about is not an objective feature of the environment but an artificial construct. In particular, the decision-maker is assumed to contemplate the worst-case scenario from a restricted set of economic models that he believes can capture his environment. This setup has the decision-maker ruling out some models with certainty even as he admits other arbitrarily close models that would be hard to distinguish empirically from those he rejected. Unfortunately, changing the admissible set of models often affects the worst-case scenario and thus the implied policy recommendation. The lost in a forest problem provides a relatively clean motivating example in which we can apply the robust control approach, although this problem obscures important issues that arise in economic applications, such as how to construct the set of scenarios from which a decision-maker calculates the worst case.

As noted previously, many mathematicians regard the robust strategy as a satisfactory solution for the lost in a forest problem. However, this strategy has been criticized in ways that mirror the criticisms of robust control applications in economics. One such critique is that the robust policy is narrowly tailored to do well in particular scenarios rather than in most scenarios. This critique is sometimes described as "perfection being the enemy of the good": The robust strategy is chosen because it does well in the one state of the world that corresponds to the worst-case scenario, even if that state is unlikely and even if the strategy performs much worse than alternative strategies in most if not all remaining states of the world.[3] In the lost in a forest problem, the worst-case scenario for any search strategy involves guessing each and every one of the wrong locations first before finding the road. Arguably, guessing wrong at each possible turn is rather unlikely. But the robust policy is tailored to this scenario, and because of this, the hiker does not take advantage of shortcuts that allow him to search through many locations without having to walk a great distance. As discussed in box 1, such shortcuts exist, but they would involve walking a longer distance if the spot on the road where the hiker started from happened to be situated at the last possible location he tries, and so the robust strategy avoids them. Viewed this way, the robust strategy might seem less appealing.

The problem with this critique is that the lost in a forest problem assumes it is not possible to assign a probability distribution to which direction the nearest point on the road lies. Absent such a distribution, one

The lost in a forest problem formally amounts to choosing a path starting from an initial point (the hiker's location when he wakes up) that must ultimately intersect with an infinite straight line (the road that cuts through the forest) whose closest distance to the initial point is one mile. That is, we need to choose a path starting from the center of a circle of radius one mile to any point on a particular straight line that is tangent to this circle. The fact that the hiker forgot where he came from corresponds to the stipulation that the location of the tangency point on the circle is unknown.

This situation is illustrated graphically in panel A of figure B1, which shows three of the continuum of possible locations for the road.

Because the forest is dense with trees, the hiker is assumed not to know where the main road is until he actually reaches it. Bellman (1956) was the first to suggest choosing the path that minimizes the longest distance needed to reach the line with absolute certainty regardless of the location of the tangency point on the unit circle. To better appreciate this criterion, consider the strategy of walking a straight path for
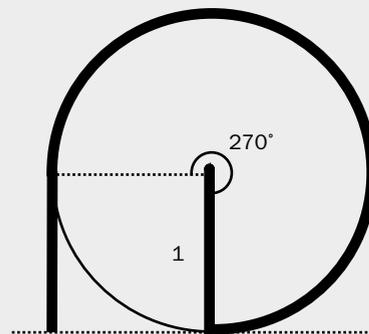
---

**FIGURE B1**
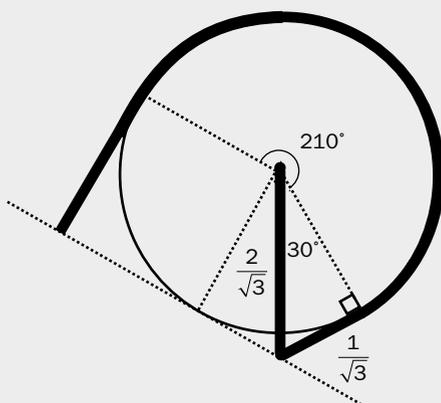
**The geometry of the lost in a forest problem**

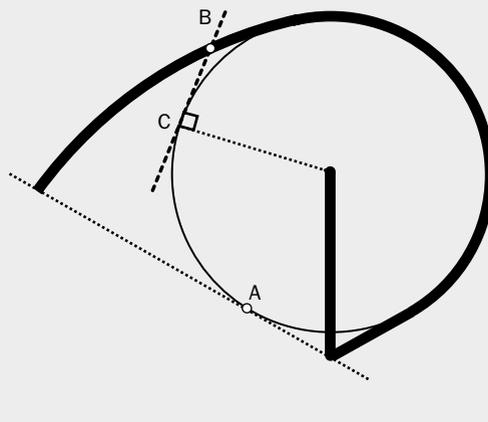**A. Graphical representation of the problem**
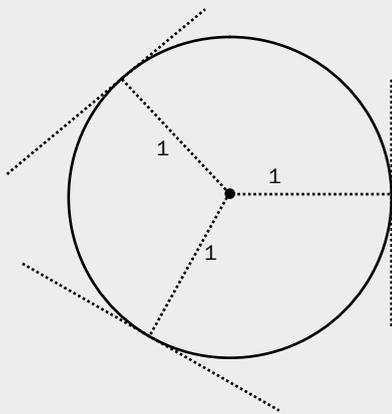
**B. Path that improves on travelling the circle**

**C. Minimax path, derived by Isbell (1957)**

**D. Conjectured min-mean path by Gluss (1961)**

Note: In panel D, the path up to B covers all locations in arc AC.

**BOX 1** (CONTINUED)

### The lost in a forest problem

one mile until reaching the circle along which the tangency point must be located, then travelling counterclockwise along the circle until reaching the road. This strategy will reach the road with certainty, and the longest distance needed to ensure reaching the road regardless of its location corresponds to walking a mile plus the entire distance of the circle, that is, $1 + 2\pi \approx 7.28$ miles. But it is possible to ensure we will reach the road regardless of its location with an even shorter path. To see this, suppose we walk out a mile and proceed to walk counterclockwise along the circle as before, but after travelling for three-fourths of the circle, rather than continuing to walk along the circle, we instead walk straight ahead for a mile, as shown in panel B of figure B1. This approach also ensures we will reach the road with certainty regardless of where it is located, but it involves walking at most $1 + \frac{3}{2}\pi + 1 \approx 6.71$ miles. It turns out that it is possible to do even better than this. The optimal path, derived by Isbell (1957), is illustrated in panel C of figure B1. This procedure involves walking out $\frac{2}{\sqrt{3}} \approx 1.15$ miles, then turning clockwise 60 degrees and walking back toward the circle for another $\frac{1}{\sqrt{3}} \approx 0.58$ miles until reaching the perimeter of the circle of radius one mile around where the hike started, walking 210 degrees along the circle, and then walking a mile along the tangent to the circle at this point. The strategy requires walking at most $\frac{1+2}{\sqrt{3}} + \frac{7}{6}\pi + 1 \approx 6.40$ miles. This is the absolute minimum one would have to walk and still ensure reaching the road regardless of its location.

When he originally posed the lost in a forest problem, Bellman (1956) suggested as an alternative strategy the path that minimizes the expected distance of reaching the road, assuming the tangency point was distributed uniformly over the unit circle. To the best of my knowledge, this problem has yet to be solved analytically. However, Gluss (1961) provided some intuition as to the nature of this solution by numerically solving for the optimal path among a parameterized set of possible strategies. He showed that the robust path in panel C of figure B1 does not minimize the expected distance, and he demonstrated various strategies that improve upon it. The general shape Gluss found to perform well among the paths he considered is demonstrated in panel D of figure B1. In order to minimize the expected distance, it turns out that it will be better to eventually stray outside the circle rather than always hewing close to it as Isbell's (1957) path does. The reason is that one can cover more possibilities walking outside the circle and reaching the road at a nontangency point than hewing to the circle and searching for tangency points. This can be seen in panel D of figure B1, where walking along the proposed path up to point B covers all possible locations for the tangency point in the arc AC, whereas walking along the arc AC would have required walking a significantly longer distance. The drawback of straying from the circle this way is that if the road happens to be located at the last possible location, the hiker would have to cover a much greater distance to reach that point. But since the probability that the road will lie in the last possible location to be searched is small under the uniform distribution, it is worth taking this risk if the goal is to minimize the *expected* travel time rather than the *maximal* travel time.

cannot argue that exhaustively searching through all other paths is an unlikely scenario, since to make this statement precise requires a probability distribution as to where the road is located. One might argue that, even absent an exact probability distribution, we can infer from common experience that we do not often run through all possibilities before we find what we are searching for, so we can view this outcome as remote even without attaching an exact probability to this event. But such intuitive arguments are tricky. Consider the popularity of the adage known as Murphy's Law, which states that whatever can go wrong will go wrong. The fact that people view things going wrong at every turn as a sufficiently common experience to be humorously compared with a scientific law suggests they might not view looking through all of the wrong locations first as such a remote possibility. Moreover, in neither the lost in a forest problem nor many economic applications is it common

that the robust strategy performs poorly in *all* scenarios other than the worst-case one. By continuity, the strategy that is optimal in the worst-case scenario will be approximately optimal in similar situations—for example, exhausting most but not all possible locations before reaching the road in the lost in a forest problem. Such continuity is common to many economic applications. Hence, even if the probability of the worst-case scenario is low, there may be other nearby states that are not as infrequent where the policy remains approximately optimal. In the lost in a forest problem, it also turns out that the robust strategy does well if the road lies in one of the regions to be explored first. But if the road lies in neither the first nor last regions to be explored, the robust strategy involves walking an unnecessarily long distance. Criticizing a policy because it performs poorly in some states imposes an impossible burden on policy. Even a policy designed

to do well on average (assuming the distribution of outcomes is known) may perform poorly in certain states of the world.

Another critique of robust control holds that, rather than choosing a policy that is robust, decision-makers should act like Bayesians; that is, they should assign subjective beliefs to the various possibilities they contemplate, compute an implied expected loss for each strategy, and then choose the strategy that minimizes the expected loss. For example, Sims (2001) argued decision-makers should avoid rules that violate the sure-thing principle, which holds that if one action is preferred to another action regardless of which event is known to occur, it should remain preferred if the event were unknown. The robust control approach can violate this principle, while subjective expected utility does not. The notion of assigning subjective probabilities to different scenarios is especially compelling in the lost in a forest problem, where assigning equal probabilities to all locations seems natural given there is no information to suggest that any one direction is more likely than the other. In fact, when Bellman (1956) originally posed his question, he suggested both minimizing the longest path (minimax) and minimizing the expected path assuming a uniform prior (min-mean) as ways of solving this problem. This approach need not contradict the policy recommendation that emerges from the robust control approach. In particular, if the cost of effort involved in walking rises steeply with the distance one has to walk, a policy that eliminates the possibility of walking very long distances would naturally emerge as desirable. But at shorter distances, assigning a probability distribution to the location of the road might lead to a different strategy from the robust one. This critique is not aimed at a particular strategy per se, but against using robustness as a criterion for choosing which policy to pursue.[4]

The problem with this critique is that it is not clear that decision-makers would always agree with the recommendation that they assign subjective probabilities to scenarios whose likelihood they do not know. As an example, assigning a distribution to the location of the road in the lost in a forest problem is incompatible with the notion of Murphy's Law. Inherent in Murphy's Law is the notion that the location of the road depends on where the hiker chooses to search. But the Bayesian approach assumes a fixed distribution regardless of what action the hiker chooses. Thus, to a person who finds Murphy's Law appealing, proceeding like a Bayesian would ring false. As another example, consider the "Ellsberg paradox," which is due to Ellsberg (1961). This paradox is based on a thought experiment in which people are asked to choose between a lottery with a known probability of winning and another lottery featuring identical prizes but with an unknown probability of winning. Ellsberg argued that most people would prefer to avoid the lottery whose probability of winning they do not know and would not choose as if they assigned a fixed subjective probability to the lottery with an unknown probability of winning. In other words, the preferences exhibited by most people would seem paradoxical to someone who behaved like a Bayesian. Subsequent researchers who conducted experiments offering these choices to real-life test subjects, starting with Becker and Brownson (1964), confirmed this conjecture. The saliency of these findings suggests that the failure to behave like a Bayesian may reflect genuine discomfort by test subjects with the Bayesian approach of assigning subjective probabilities to outcomes whose probabilities they do not know. But if this is so, we cannot objectively fault decision-makers for not adopting the Bayesian approach, since any recommendation we make to them would have to respect their preferences.

Of course, even accepting that policymakers may not always find the Bayesian approach appealing, it does not automatically follow that they should favor the robust control approach in particular. The relevant question is whether there is a compelling reason for decision-makers to specifically prefer the robust policy. One result often cited by advocates of robust control is the work of Gilboa and Schmeidler (1989). They show that if decision-makers' preferences over lotteries satisfy a particular set of restrictions, it will be possible to represent their choices as if they chose the action that minimizes the largest possible expected loss across a particular set of probability distributions. However, this result is not an entirely satisfactory argument for why policymakers should adopt the robust control approach. First, there is little evidence to suggest that standard preferences obey the various restrictions derived by Gilboa and Schmeidler (1989). While the Ellsberg paradox suggests many people have preferences different from those that would lead them to behave like Bayesians, it does not by itself confirm that preferences accord with each one of the restrictions in Gilboa and Schmeidler (1989). Second, Gilboa and Schmeidler (1989) show that the set of scenarios from which the worst case is calculated depends on the preferences of the decision-makers. This is not equivalent to arguing that policymakers, once they restrict the set of admissible models that could potentially account for the data they observe, should always choose the action that minimizes the worst-case outcome from this set.

In the lost in a forest problem, Gilboa and Schmeidler's (1989) result only tells us that if an

individual exhibited particular preferences toward lotteries whose outcomes dictate distances he would have to walk, he would in fact prefer the minimax solution to this problem. It does not say that whenever he faces uncertainty more generally—for example, if he also forgot how far away he was from the road when he lay down—that he would still choose the strategy dictated by the robust control approach. In short, Gilboa and Schmeidler (1989) show that opting for a robust strategy is coherent in that we can find well-posed preferences that rationalize this behavior, but their analysis does not imply such preferences are common or that robustness is a desirable criterion whenever one is at a loss to assign probabilities to various possible scenarios.[5]

The theme that runs through the discussion thus far is that if the decision-makers cannot assign probabilities to scenarios they are uncertain about, there is no inherently correct criterion on how to choose a policy. As Manski (2000, p. 421) put it, "there is no compelling reason why the decision maker should or should not use the maximin rule when [the probability distribution] is a fixed but unknown objective function. In this setting, the appeal of the maximin rule is a personal rather than normative matter. Some decision makers may deem it essential to protect against worst-case scenarios, while others may not."[6] One can point to unappealing elements about robust control, but these do not definitively rule out this approach. Conversely, individuals with particular preferences toward lotteries might behave as if they were following a minimax rule, but this does not imply that they will adopt such a rule whenever they are unable to assign probabilities to possible scenarios. Among engineers, the notion of designing systems that minimize the worst-case scenario among the set of possible states whose exact probability is unknown has carried some appeal. Interestingly, Murphy's Law is also an export from the field of engineering.[7] The two observations may be related: If worst-case outcomes are viewed not as rare events but as common experiences, robustness would naturally seem like an appealing criterion. Policymakers who are nervous about worst-case outcomes would presumably find appeal in the notion of keeping the potential risk exposure to the bare minimum. More generally, studying robust policies can help us to understand the costs and benefits of maximally aggressive risk management so that policymakers can contemplate their desirability.

## Recasting the Brainard model as a robust control problem

Now that I have described what it means for a policy to be robust, I can return to the question of how a policymaker concerned about robustness should act when trying to target a variable in an uncertain environment. In particular, I will now revisit the environment that Brainard (1967) considered, but with one key difference: The policymaker is assumed to be unable to assign probabilities to the scenarios he is uncertain about. In what follows, I introduce uncertainty in a way that Hansen and Sargent (2008) and Williams (2008) describe as structured uncertainty; that is, I assume the policymaker knows the model but is uncertain about the exact value of one of its parameters. More precisely, he knows that the parameter lies in some range, but he cannot ascribe a probability distribution to the values within this range. By contrast, unstructured uncertainty corresponds to the case where a model is defined as a probability distribution over outcomes, and where the policymaker is unsure about which probability distribution from some set represents the true distribution from which the data are drawn.[8]

Once again, I begin by assuming that the variable in question, $y$, is affected linearly by a policy variable, $r$; various factors that the policymaker can observe prior to setting policy, $x$; and other factors that the policymaker cannot observe prior to setting his policy but whose distribution is known, $\varepsilon_u$:

5) $\quad y = x - kr + \varepsilon_u.$

As before, I assume $\varepsilon_u$ has mean 0 and variance $\sigma_u^2$. To capture uncertainty about the effect of policy, I modify the coefficient on $r$ to allow for uncertainty:

6) $\quad y = x - (k + \varepsilon_k)\, r + \varepsilon_u.$

In contrast to Brainard's (1967) setup, I assume that rather than knowing the distribution of $\varepsilon_k$, the policymaker only knows that its support is restricted to the interval $\left[\underline{\varepsilon}, \bar{\varepsilon}\right]$ that includes 0, that is, $\underline{\varepsilon} < 0 < \bar{\varepsilon}$. In other words, the effect of $r$ on $y$ can be less than, equal to, or higher than $k$. Beyond this, he will not be able to assign probabilities to particular values within this interval.

Since the support of $\varepsilon_k$ will figure prominently in formulating the robust strategy, it is worth commenting on where it might come from. In practice, information about $k + \varepsilon_k$ is presumably compiled from past data. That is, given time-series data on $y$, $x$, and $r$, we can estimate $k + \varepsilon_k$ by using standard regression techniques. With a finite history, our estimate would necessarily be noisy due to variation from $\varepsilon_u$. However, we might still be able to reject some values for $k + \varepsilon_k$ as implausible—for example, values that are several standard errors away from our point estimate. Still, there is something seemingly arbitrary in classifying some values of $k$ as possible while treating virtually identical

values as impossible. Although this may be consistent with the common practice of "rejecting" hypotheses whose likelihood falls below some set cutoff, it is hard to rationalize such dogmatic rules for including or omitting possible scenarios in constructing worst-case outcomes. In what follows, I will treat the support for $\varepsilon_k$ as given, sidestepping these concerns.[9]

The robust control approach in this environment can be cast as a two-step process. First, for each value of $r$, we compute its worst-case scenario over all values $\varepsilon_k \in \left[\underline{\varepsilon}, \bar{\varepsilon}\right]$, or the largest expected loss the policymaker could incur. Define this expected loss as $W(r)$; that is,

$$W(r) \equiv \max_{\varepsilon_k \in \left[\underline{\varepsilon}, \bar{\varepsilon}\right]} E\left[y^2\right] = \max_{\varepsilon_k \in \left[\underline{\varepsilon}, \bar{\varepsilon}\right]} \left\{\left(x - \left(k + \varepsilon_k\right)r\right)^2 + \sigma_u^2\right\}.$$

Second, we choose the policy $r$ that implies the smallest value for $W(r)$. The robust strategy is defined as the value of $r$ that solves $\min_r W(r)$; that is,

7) $\quad \min_r \max_{\varepsilon_k \in \left[\underline{\varepsilon}, \bar{\varepsilon}\right]} \left\{\left(x - \left(k + \varepsilon_k\right)r\right)^2 + \sigma_u^2\right\}.$

I explicitly solve equation 7 in appendix 2. In what follows, I limit myself to describing the robust strategy and providing some of the intuition behind it. It turns out that the robust policy hinges on the lowest value that $\varepsilon_k$ can assume. If $\underline{\varepsilon} < -k$, which implies that the coefficient $k + \varepsilon_k$ can assume either positive or negative values, the solution to equation 7 is given by

8) $\quad r = 0.$

If instead $\underline{\varepsilon} > -k$, so the policymaker is certain that the coefficient $k + \varepsilon_k$ is positive (but is still unsure of its value), the solution to equation 7 is given by

9) $\quad r = \dfrac{x}{k + \left(\underline{\varepsilon} + \bar{\varepsilon}\right)/2}.$

Thus, if the policymaker knows the sign of the effect of $r$ on $y$, he will respond to changes in $x$ in a way that depends on the extreme values $\varepsilon_k$ can assume, that is, the endpoints of the interval $\left[\underline{\varepsilon}, \bar{\varepsilon}\right]$. But if the policymaker is unsure about the sign of the effect of policy on $y$, he will not respond to changes in $x$ at all.

To better understand why concerns about robustness lead to this rule, consider first the result that if the policymaker is uncertain about the sign of $k + \varepsilon_k$, he should altogether abstain from responding to $x$.

This is related to Brainard's (1967) original attenuation result: There is an inherent asymmetry in that a passive policy where $r = 0$ leaves the policymaker unexposed to risk from $\varepsilon_k$, while a policy that sets $r \neq 0$ leaves him exposed to such risk. When the policymaker is sufficiently concerned about the risk from $\varepsilon_k$, which turns out to hinge on whether he knows the sign of the coefficient on $r$, he is better off resorting to a passive policy that protects him from this risk than trying to offset nonzero values of $x$. However, the attenuation here is both more extreme and more abrupt than what Brainard found. In Brainard's formulation, the policymaker will always act to offset $x$, at least in part, but he will moderate his response to $x$ continuously with $\sigma_k^2$. By contrast, robustness considerations imply a threshold level for the lower support of $\varepsilon_k$, which, if crossed, leads the policymaker to radically shift from actively offsetting $x$ to passively not responding to it at all.
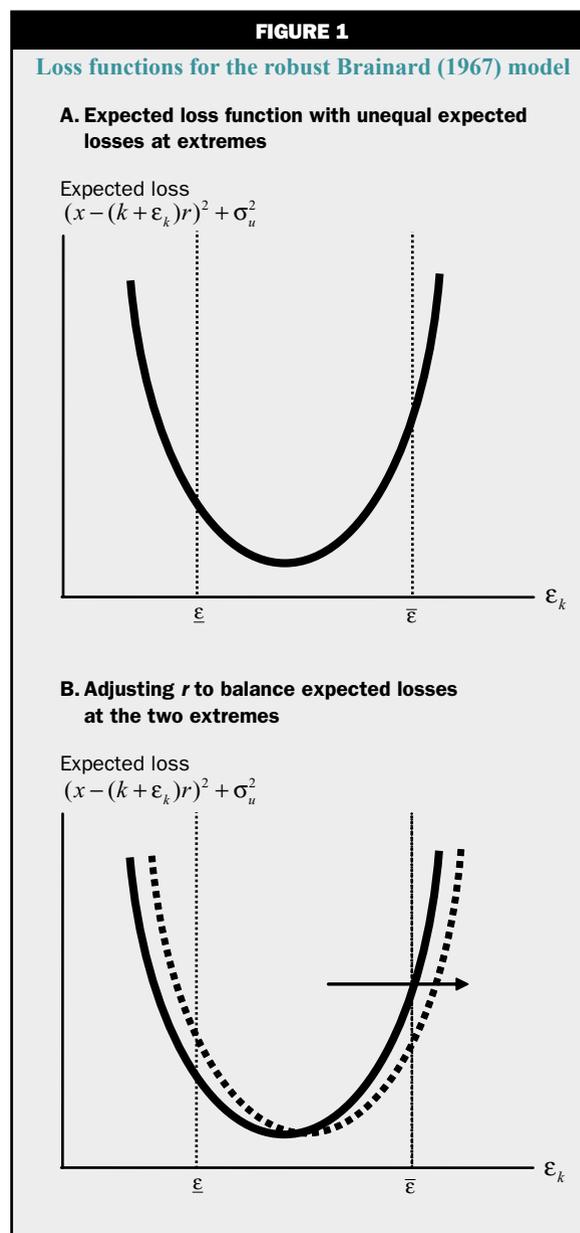
The abrupt shift in policy in response to small changes in $\varepsilon_k$ demonstrates one of the criticisms of robust control cited earlier—namely, that this approach formulates policy based on how it performs in specific states of the world rather than how it performs in general. When $\underline{\varepsilon}$ is close to $-k$, it turns out that the policymaker is almost indifferent among a large set of policies that achieve roughly the same worst-case loss. When $\underline{\varepsilon}$ is just below $-k$, setting $r = 0$ performs slightly better under the worst-case scenario than setting $r$ according to equation 9. When $\underline{\varepsilon}$ is just above $-k$, setting $r$ according to equation 9 performs slightly better under the worst-case scenario than setting $r = 0$. When $\underline{\varepsilon}$ is exactly equal to $-k$, both strategies perform equally well in the worst-case scenario, as does any other value of $r$. However, the two strategies lead to different payoffs in scenarios other than the worst case, that is, for values of $\varepsilon_k$ that are between $\underline{\varepsilon}$ and $\bar{\varepsilon}$. Hence, concerns for robustness might advocate dramatic changes in policy to eke out small gains under the worst-case scenario, even if these changes result in substantially larger losses in most other scenarios. A dire pessimist would feel perfectly comfortable guarding against the worst-case scenario in this way. But in situations such as this, where the policymaker chooses his policy based on minor differences in how the policies perform in one particular case even when the policies result in enormous differences in other cases, the robust control approach has a certain tail-wagging-the-dog aspect to it that makes it seem less appealing.

Next, consider what robustness considerations dictate when the policymaker knows the sign of $k + \varepsilon_k$ but not its precise magnitude. To see why $r$ depends on the endpoints of the interval $\left[\underline{\varepsilon}, \bar{\varepsilon}\right]$, consider figure 1. This

figure depicts the expected loss $[(x-(k+\varepsilon_k)r)^2+\sigma_u^2]$ for a fixed $r$ against different values of $\varepsilon_k$. The loss function is quadratic and convex, which implies the largest loss will occur at one of the two extreme values for $\varepsilon_k$. Panel A of figure 1 illustrates a case in which the expected losses at $\varepsilon_k=\underline{\varepsilon}$ and $\varepsilon_k=\overline{\varepsilon}$ are unequal: The expected loss is larger for $\varepsilon_k=\overline{\varepsilon}$. But if the losses are unequal under some rule $r$, that value of $r$ fails to minimize the worst-case scenario. This is because, as illustrated in panel B of figure 1, changing $r$ will shift the loss function to the left or the right (it might also change the shape of the loss function, although this can effectively be ignored). The policymaker should thus be able to reduce the largest possible loss over all values of $\varepsilon_k$ in $[\underline{\varepsilon},\overline{\varepsilon}]$. Although shifting $r$ would lead to a greater loss if $\varepsilon_k$ happened to equal $\underline{\varepsilon}$, since the goal of a robust policy is to reduce the largest possible loss, shifting $r$ in this direction is desirable. Robustness concerns would therefore lead the policymaker to adjust $r$ until the losses at the two extreme values were balanced, that is, until the loss associated with the policy being maximally effective was exactly equal to the loss associated with the policy being minimally effective.

When there is no uncertainty, that is, when $\underline{\varepsilon}=\overline{\varepsilon}=0$, the policymaker would set $r=x/k$, since this would set $y$ exactly equal to its target. When there is uncertainty, whether the robust policy will respond to $x$ more or less aggressively than this benchmark depends on how the lower and upper bounds are located relative to 0. If the region of uncertainty is symmetric around 0 so that $\underline{\varepsilon}=-\overline{\varepsilon}$, uncertainty has no effect on policy. To see this, note that if we were to set $r=x/k$, the expected loss would reduce to $(x/k)^2\varepsilon_k^2+\sigma_u^2$, which is symmetric in $\varepsilon_k$. Hence, setting $r$ to offset $x$ would naturally balance the loss at the two extremes. But if the region of uncertainty is asymmetric around 0, setting $r=x/k$ would fail to balance the expected losses at the two extremes, and $r$ would have to be adjusted so that it either responds more or less to $x$ than in the case of complete certainty. In particular, the response to $x$ will be attenuated if $\overline{\varepsilon}>-\underline{\varepsilon}$, that is, if the potential for an overly powerful stimulus is greater than the potential for an overly weak stimulus, and will be amplified in the opposite scenario.

This result begs the question of when the support for $\varepsilon_k$ will be symmetric or asymmetric in a particular direction. If the region of uncertainty is constructed using past data on $y$, $x$, and $r$, any asymmetry would have to be driven by differences in detection probabilities across different scenarios—for example, if it is more difficult to detect $k$ when its value is large than



**FIGURE 1**

**Loss functions for the robust Brainard (1967) model**

**A. Expected loss function with unequal expected losses at extremes**

Expected loss
$(x-(k+\varepsilon_k)r)^2+\sigma_u^2$

**B. Adjusting _r_ to balance expected losses at the two extremes**

Expected loss
$(x-(k+\varepsilon_k)r)^2+\sigma_u^2$

when it is small. This may occur if the distribution of $\varepsilon_u$ were skewed in a particular direction. But if the distribution of $\varepsilon_u$ were symmetric around 0, policymakers who rely on past data should find it equally difficult to detect deviations in either direction, and the robust policy would likely react to shocks in the same way as if $k$ were known with certainty.

Deriving the robust strategy in Brainard's (1967) setting reveals two important insights. The first is that the robustness criterion does not inherently imply that policy should be more aggressive in the face of uncertainty. Quite to the contrary, the robust policy exhibits

a more extreme form of the same attenuation principle that Brainard demonstrated, for essentially the same reason: The asymmetry between how passive and active possibilities leave the policymaker exposed to risk tends to favor passive policies. More generally, whether facing uncertainty about the economic environment leads to a more gradual policy or a more aggressive policy depends on asymmetries in the underlying environment. If the policymaker entertains the possibility that policy can be far too effective but not that it will be very ineffective, he will naturally tend to attenuate his policy. But if his beliefs are reversed, he will tend to magnify his response to news about potential deviations from the target level.[10]

The second insight is that, at least in some circumstances, the robust strategy can be described as one that balances the losses from different risks. Consider the case where $\underline{\varepsilon} > -k$: In that case, the robust strategy will be the one that equates the loss from the risk of policy being overly effective with the loss from the risk of policy being insufficiently effective. This suggests that, in some cases, robustness amounts to a recommendation of keeping opposing risks in balance, in line with what monetary authorities often cite as the principle that guides their policies in practice. That said, the notion that concerns for robustness amount to balancing losses in this way is not common to all environments. The next section presents an example in which robustness considerations would recommend proceeding as if the policymaker knew the worst-case scenario to be true rather than to keep different risks in balance. Moreover, the notion of balancing losses or risks is implicit in other approaches to modeling decision-making under uncertainty. For example, choosing a policy to minimize expected losses will typically call on the policymaker to equate expected marginal losses across states of the world or to otherwise balance expected costs and benefits of particular policies. Hence, robust control is neither uniquely nor fundamentally a recommendation to balance risks. Nevertheless, in some circumstances it involves balancing opposing risks in a way that mirrors some of the stated objectives of monetary authorities.

## Robustness and aggressive rules

Since robustness considerations can lead to policies that are either more gradual or more aggressive, depending on the underlying asymmetry of the environment, it seems natural to ask which asymmetries tended to favor aggressive policies in the original work on robust monetary policy. The papers cited earlier consider different environments, and their results are not driven by one common feature. I now offer two examples inspired

by these papers to illustrate some of the relevant asymmetries. The first assumes the policymaker is uncertain about the persistence of shocks, following Sargent (1999). The second assumes the policymaker is uncertain about the trade-off between competing objectives, following Giannoni (2002). As I show next, both of these features could tilt a policymaker who is uncertain about his environment in the direction of overreacting to news about missing a particular target, albeit for different reasons.

### Uncertain persistence

One of the first to argue that concerns for robustness could dictate a more aggressive policy under uncertainty than under certainty was Sargent (1999). In discussing a model proposed by Ball (1999), Sargent asked how optimal policy would be affected when we account for the possibility that the model is misspecified—in particular that the specification errors are serially correlated. To gain insight into this question, I adapt the model of trying to meet a target I described earlier to allow for the possibility that the policymaker is uncertain about the persistence of the shocks he faces, and I examine the implied robust policy. I show that there is an asymmetry in the loss from underreacting to very persistent shocks and the loss from overreacting to moderately persistent shocks. Other things being equal, this tilts the policymaker toward reacting more to past observable shocks when he is uncertain about the exact degree of persistence.

Formally, consider a policymaker who wants to target a variable that is affected by both policy and other factors. Although Sargent (1999) considers a model in which the policymaker is concerned about multiple variables, it will be simpler to assume there is only one variable he cares about. Let $y_t$ denote the value at date $t$ of the variable that the policymaker wishes to target to 0. As in equation 5, I assume $y_t$ is linear in the policy variable $r_t$ and in an exogenous shock term $x_t$:

10) $y_t = x_t - kr_t$.

Here, I no longer assume the policymaker is uncertain about the effect of his policy on $y$. As such, it will be convenient to normalize $k$ to 1. However, I now assume he is uncertain about the way $x_t$ is correlated over time. Suppose

11) $x_t = \rho x_{t-1} + \varepsilon_t$,

where $\varepsilon_t$ are independent and identically distributed over time with mean 0 and variance $\sigma_\varepsilon^2$. At each date $t$, the policymaker can observe $x_{t-1}$ and condition his

policy on its realization. However, he must set $r_t$ before observing $x_t$. He will be uncertain about $x_t$ for two reasons: He must act before getting to observe $\varepsilon_t$, and he may not know the value of $\rho$ with certainty.

I assume the policymaker discounts future losses at rate $\beta < 1$ so that his expected loss is given by

$$E\left[\sum_{t=0}^{\infty}\beta^t y_t^2\right] = E\left[\sum_{t=0}^{\infty}\beta^t\left(x_t - r_t\right)^2\right].$$

If the policymaker knew $\rho$ with certainty, his optimal strategy would be to set $r_t = \rho x_{t-1}$, which is the expected value of $x_{t+1}$. Suppose instead that he knew $\rho$ fell in some interval $\left[\underline{\rho}, \bar{\rho}\right]$. Let $\rho^*$ denote the midpoint of this interval; that is,

$$\rho^* = \left(\underline{\rho} + \bar{\rho}\right)/2.$$

To emphasize asymmetries inherent to the loss function as opposed to the region of uncertainty, suppose the interval of uncertainty is symmetric around the certainty benchmark; that is, in assessing whether the robust policy is more aggressive, we will compare it to the policy the monetary authority would pursue if it knew $\rho = \rho^*$. An important and empirically plausible assumption in what follows is that $\rho^* > 0$; that is, the beliefs of the monetary authority are centered around the possibility that shocks are positively correlated.

Once again, we can derive the robust strategy in two steps. First, for each rule $r_t$, define $W(r_t)$ as the biggest loss possible among the different values of $\rho$; that is,

$$W\left(r_t\right) = \max_{\rho \in \left[\underline{\rho}, \bar{\rho}\right]} E\left[\sum_{t=0}^{\infty}\beta^t y_t^2\right].$$

We then choose the policy rule $r_t$ that minimizes $W(r_t)$; that is, we solve

12) $\min_{r_t} \max_{\rho \in \left[\underline{\rho}, \bar{\rho}\right]} E\left[\sum_{t=0}^{\infty}\beta^t y_t^2\right].$

Following Sargent (1999), I assume the policymaker is restricted in the type of policies $r_t$ he can carry out: The policymaker must choose a rule of the form $r_t = ax_{t-1}$, where $a$ is a constant that cannot vary over time. This restriction is meant to capture the notion that the policymaker cannot learn about the parameters about which he is uncertain and then change the way policy reacts to information as he observes $x_t$ over time and potentially infers $\rho$. I further assume the expectation

in equation 12 is the unconditional expectation of future losses; that is, the policymaker calculates his expected loss from the perspective of date 0. To simplify the calculations, I assume $x_0$ is drawn from the stationary distribution for $x_t$.

The solution to equation 12, subject to the constraint that $r_t = ax_{t-1}$, is derived in appendix 3. The key result shown in that appendix is that as long as $\rho^* > 0$, the robust policy would set $a$ to a value in the interval $\left[\underline{\rho}, \bar{\rho}\right]$ that is strictly greater than the midpoint $\rho^*$. In other words, starting with the case in which the policymaker knows $\rho = \rho^*$, if we introduce a little bit of uncertainty in a symmetric fashion, so the degree of persistence can deviate equally in either direction, the robust policy would react more to a change in $x_{t-1}$ in the face of uncertainty than it would react to such a change if the degree of persistence were known with certainty.

To understand this result, suppose the policymaker instead set $\rho = \rho^*$. As in the previous section, the loss function is convex in $\rho$, so the worst-case scenario will occur when $\rho$ assumes one of its two extreme values, that is, either when $\rho = \bar{\rho}$ or $\rho = \underline{\rho}$. It turns out that when $a = \rho^*$, setting $\rho = \bar{\rho}$ imposes a bigger cost on the policymaker than setting $\rho = \underline{\rho}$. Intuitively, for any given $\rho$, setting $a = \rho^*$ will imply $y_t = (\rho - \rho^*) x_{t-1} + \varepsilon_t$. The expected deviation of $y_t$ from its target given $x_{t-1}$ will have the same expected magnitude in both cases; that is, $\left|(\rho - \rho^*) x_{t-1}\right|$ will be the same when $\rho = \bar{\rho}$ and $\rho = \underline{\rho}$, given $\left[\underline{\rho}, \bar{\rho}\right]$ is symmetric around $\rho^*$. However, the process $x_t$ will be more persistent when $\rho$ is higher, and so deviations from the target will be more persistent when $\rho = \bar{\rho}$ than when $\rho = \underline{\rho}$. More persistent deviations imply more volatile $y_t$ and hence a larger expected loss. Since the robust policy tries to balance the losses at the two extreme values of $\rho$, the policymaker should choose a higher value for $a$ to reduce the loss when $\rho = \bar{\rho}$.

The basic insight is that, while the loss function for the policymaker is symmetric in $\rho$ around $\rho = 0$, if we focus on an interval that is centered in either direction of $\rho = 0$, the loss function will be asymmetric. This asymmetry will tend to favor policies that react more to past shocks. In fact, this feature is not unique to policies guided by the robustness criterion: If we assumed the policymaker assigned symmetric probabilities to values of $\rho$ in the interval $\left[\underline{\rho}, \bar{\rho}\right]$ and acted to minimize expected losses, the asymmetry in the loss function would tilt his policy toward being more aggressive; that is, the value of $a$ that solves

$\min_{a} E\left[\sum_{t=0}^{\infty}\beta^t\left(\rho x_{t-1} - ax_{t-1} + \varepsilon_t\right)^2\right]$ would also exceed

$\rho^*$ if $\rho^* > 0$.

While my example highlights a force that generally favors more aggressive policies, it should be emphasized that my exercise is not quite equivalent to the one in Sargent (1999). Sargent allowed the policymaker to entertain stochastic processes that are correlated as in my example, but he also allowed the policymaker to entertain the possibility that the mean of the process is different from zero. In addition, the region of uncertainty Sargent posited was not required to be symmetric around the certainty benchmark; rather, it was constructed based on which types of processes are easier to distinguish from the certainty benchmark case. The analysis here reveals an asymmetry in the loss function that magnifies concerns about more persistent processes and thus encourages reacting more to past shocks, but the nature of the robust policy depends crucially on the set of scenarios from which the worst-case is constructed.

### Uncertain trade-off parameters

Following Sargent (1999), other papers also argued that robust policies tended to be aggressive. While these papers reached the same conclusion as Sargent (1999), they considered different environments. For example, Giannoni (2002) assumed that the policymakers know the persistence of shocks but are uncertain about parameters of the economic model that dictate the effect of these shocks on other economic variables. This leaves policymakers uncertain about the trade-off between their competing objectives. In this environment, Giannoni too found that the robust policy is more aggressive in the face of uncertainty than when policymakers know the trade-off parameters with certainty.

To provide some insight behind these results, consider the following simplified version of Giannoni's (2002) model. Suppose the monetary authority cares about two variables, denoted $y$ and $\pi$; in Giannoni's model, $y$ and $\pi$ are the output gap and inflation, respectively. The monetary authority has a quadratic loss function:

13) $\alpha y^2 + \pi^2$.

The variables $y$ and $\pi$ are related linearly

14) $\pi = \lambda y + x$,

where $x$ is an observable shock.[11] This relationship implies a trade-off between $\pi$ and $y$. If we set $\pi = 0$ as desired, then $y$ would vary with $x$ and deviate from 0. If we set $y = 0$, then $\pi$ would vary with $x$ and deviate from 0. Substituting equation 14 into the loss

function allows us to express the policymaker's problem as choosing $\pi$ to minimize the loss

$$\alpha \frac{(\pi - x)^2}{\lambda^2} + \pi^2.$$

Taking the first-order condition with respect to $\pi$ gives us the optimal choices for $y$ and $\pi$ as

15) $\pi = \dfrac{\alpha x}{\alpha + \lambda^2}, y = -\dfrac{\lambda x}{\alpha + \lambda^2}.$

Suppose the policymaker were uncertain about $\lambda$, knowing only that it lies in some interval $\left[\underline{\lambda}, \bar{\lambda}\right]$. Given a choice of $\pi$, the worst-case scenario over this range of $\lambda$ is given by

$$\max_{\lambda \in \left[\underline{\lambda}, \bar{\lambda}\right]} \alpha \frac{(\pi - x)^2}{\lambda^2} + \pi^2.$$

The worst case always corresponds to $\lambda = \underline{\lambda}$, except when $\pi = x$, in which case the value of $\lambda$ has no effect on the loss function. The robust strategy, therefore, is to set $\pi$ and $y$ to their values in equation 15 as if the policymaker knew $\lambda = \underline{\lambda}$, the lowest value $\lambda$ can assume. Thus, as long as the certainty benchmark $\lambda$ lies in the interior of the uncertainty interval $\left[\underline{\lambda}, \bar{\lambda}\right]$, concerns for robustness will lead the policymaker to have $y$ respond less to $x$, as well as $\pi$ respond more to $x$. The robustness criterion leads the policymaker to stabilize $y$ more aggressively against shocks to $x$ and stabilize $\pi$ less aggressively against these same shocks. The reason is that when a shock $x$ causes $\pi$ to deviate from its target, a lower value of $\lambda$ implies that pushing $\pi$ back to its target would require $y$ to deviate from its target by a greater amount. The worst-case scenario is if $y$ deviates to the largest extent possible, and so the robust policy advocates stabilizing $y$ more aggressively while loosening up on $\pi$.

Figure 2 illustrates this result graphically. The problem facing the policymaker is to choose a point from the line given by $\pi = \lambda y + x$. Ideally, it would like to move toward the origin, where $\pi = y = 0$. Changing $\lambda$ will rotate the line from which the policymaker must choose as depicted in the figure. A lower value of $\lambda$ corresponds to a flatter curve. Given the policymaker prefers to be close to the origin, a flatter curve leaves the policymaker with distinctly worse options that are farther from the origin, since one can show that the policymaker would only choose points

in the upper left quadrant of the figure. This explains why the worst-case scenario corresponds to the flattest curve possible. If we assume the policymaker must choose his relative position on the line before knowing the slope of the line (that is, before knowing λ), then the flatter the line could be, the greater his incentive will be to locate close to the π-axis rather than risk deviating from his target on both variables, as indicated by the path with the arrow. This corresponds to more aggressively insulating *y* from *x*.

Robustness concerns thus encourage the policymaker to proceed as if he knew λ was equal to its lowest possible value. Note the difference from the two earlier models, in which the robust policy recommended balancing losses associated with two opposing risks. Here, by contrast, the policy equates two *marginal* losses (the loss from letting *y* deviate a little more from its target and the loss from letting π deviate a little more) for a particular risk, namely, that λ will be low. Thus, robustness does not in general amount to balancing losses from different risks as in my two previous examples.

## Conclusion

In recent years, economists have paid increasing attention to the problem of formulating policy under uncertainty, particularly when it is not possible to attach probabilities to the scenarios that concern policymakers. One recommendation for policy in these circumstances, often attributed to Wald (1950), is the robust control approach, which argues for choosing the policy that achieves the most favorable worst-case outcome. Recent work has applied this notion to economic questions, especially in dynamic environments. However, there seem to be only limited references to this literature in monetary policy circles.

One reason for the limited impact of this literature appears to be that early applications of robust control to monetary problems focused on applications in which policymakers should amplify rather than attenuate their responses to shocks, in contrast with the theme emphasized by Brainard (1967) in his model. One of the points of my article is that comparing these applications of robust control to Brainard's result is somewhat misleading, since the policymaker is guarding against different types of uncertainty in the two models. Brainard examined the case of a policymaker who was unsure as to the effect of his policy on the variable he wanted to target; the policymaker found that attenuating his response to shocks would be optimal. But applying a robustness criterion in the same environment would suggest attenuating policy even more drastically, not responding at all to shocks when the range of possibilities that the policymaker is uncertain about is large.



**FIGURE 2**

**The uncertain trade-off model**
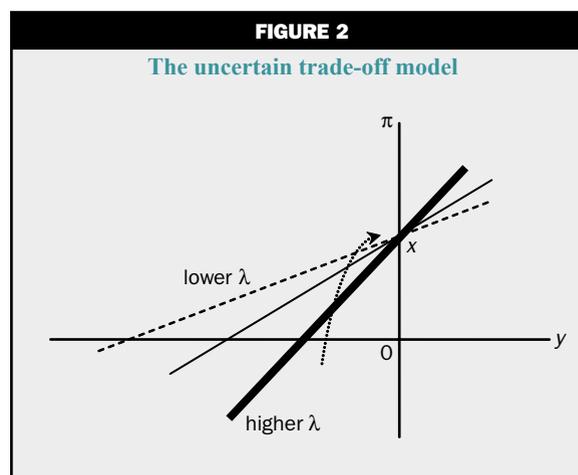
By contrast, early applications of robust control were concerned with uncertainty over the persistence of shocks or parameters that govern the trade-off between conflicting policy objectives. In that case, robustness concerns suggest amplifying the response of policy to shocks. But so would the expected utility criterion that Brainard considered. Whether the policymaker should change his policies in the face of uncertainty depends on the nature of that uncertainty—specifically whether it involves any inherent asymmetries that would tilt policy from its benchmark when the policymaker is fully informed. These considerations can be as important as the criterion by which a policy is judged to be optimal.

Ultimately, whether policymakers will find the robust control approach appealing depends on their preferences and on the particular application at hand. If policymakers are pessimistic and find an appeal in the dictum of Murphy's Law, which holds that things that can go wrong quite often do go wrong, they will find minimizing the worst-case scenario appealing. But in an economic environment where monetary policy has little impact on the worst-case scenario and has substantial impact on other scenarios, as in one of the examples presented here, even relatively pessimistic policymakers might find alternative approaches to robust control preferable for guiding the formulation of monetary policy.

## NOTES

[1]Onatski and Stock (2002) also argue that aggressiveness is not a generic feature of robust policies, but they note that in their framework, aggressiveness will arise for many classes of perturbations.

[2]The problem is sometimes referred to as the "lost at sea" problem. Richard Bellman, who originally posed the problem, is well-known among economists for his work on dynamic programming, a tool used by economists to analyze sequential decision problems that unfold over time. I refer to this example only because of its intuitive appeal and not to draw any analogies between the practice of conducting monetary policy and the situation of being lost.

[3]An example of this critique can be found in Svensson (2007), who writes "if a Bayesian prior probability measure were to be assigned to the feasible set of models, one might find that the probability assigned to the models on the boundary are exceedingly small. Thus, highly unlikely models can come to dominate the outcome of robust control." Bernanke (2007) also refers to (but does not invoke) this critique of robust control.

[4]In this spirit, Sims (2001) suggested using robust control to aid Bayesian decision-makers who for convenience rely on procedural rules rather than explicit optimization. The idea is to figure out under what prior beliefs about the models it would be optimal to pursue the robust strategy, and then reflect on whether this prior distribution seems sensible. As Sims notes, deriving the robust strategy "may alert decision-makers to forms of prior that, on reflection, do not seem far from what they actually might believe, yet imply decisions very different from that arrived at by other simple procedures." Interestingly, Sims' prescription cannot be applied to the lost in a forest problem: There is no distribution over the location of the road for which the minimax path minimizes expected distance. However, an analogous argument can be made regarding the cost of effort needed to walk various distances. If the hiker relies on procedural rules, he can back out how steeply the cost of effort must rise with distance to justify the minimax path. This will alert him to paths that are different from recommendations arrived at by other simple procedures but rely on plausible cost of effort functions.

[5]Strictly speaking, Gilboa and Schmeidler (1989) only consider a static one-shot decision, while the lost in a forest problem is a dynamic problem in which the hiker must constantly choose how to search given the results of his search at each point in time. However, the problem can be represented as a static decision in which the hiker chooses his search algorithm before starting to search. This is because he would never choose to revise his plans as a result of what he learns; if he did, he could have designed his search algorithm that way before starting to search. This will not be true in many economic applications, where there may be problems with time inconsistency that complicate the task of how to extend the minimax notion to such choices. For a discussion of axiomatic representations of minimax behavior in dynamic environments, see Epstein and Schneider (2003) and Maccheroni, Marinacci, and Rustichini (2006).

[6]For the purposes of this article, the terms "minimax rule" and "robust strategy" can be viewed as interchangeable with the term "maximin rule" that Manski (2000) uses.

[7]According to Spark (2006), the law is named after aerospace engineer Edward Murphy, who complained after a technician attached a pair of sensors in a precisely incorrect configuration during a crash test Murphy was observing. Engineers on the team Murphy was working with began referring to the notion that things will inevitably go wrong as Murphy's Law, and the expression gained public notoriety after one of the engineers used it in a press conference.

[8]For example, the set of distributions the policymaker is allowed to entertain might be those whose relative entropy to the distribution of some benchmark model falls below some threshold. There are other papers that similarly focus on structured uncertainty because of its simplicity—for example, Svensson (2007).

[9]One way to avoid this issue is to model concerns for robustness in a different way. In particular, rather than restrict the set of scenarios policymakers can entertain to some set, we can allow the set to be unrestricted but introduce a penalty function that punishes scenarios continuously depending on how much they differ from some benchmark scenario—for example, the point estimates from an empirical regression. This formulation is known as multiplier preferences, since the penalty function is scaled by a multiplicative parameter that captures concern for model misspecification. See Hansen and Sargent (2008) for a more detailed discussion.

[10]Rustem, Wieland, and Zakovic (2007) also consider robust control in asymmetric models, although they do not discuss the implications of this for gradual policy versus aggressive policy.

[11]This relationship is a simplistic representation of the New Keynesian Phillips curve Giannoni (2002) uses, in which inflation at date $t$ depends on the output gap at date $t$, expected inflation at date $t + 1$, and a shock term; that is, $\pi_t = \lambda y_t + \beta E_t \pi_{t+1} + x_t$. If $x_t$ is independent and identically distributed over time, expected inflation would just enter as a constant, and the results would be identical to those in the simplified model I use.

## APPENDIX 1. DERIVING THE OPTIMAL RULE IN BRAINARD'S (1967) MODEL

This appendix derives the optimal policy that solves equation 3 (on p. 39). Using the fact that $\frac{d}{dr}E\left[y^2\right]$ $= E\left[\frac{d}{dr}y^2\right]$, we can write the first-order condition for the problem in equation 3 as $E[2y \times dy/dr] = 0$. This implies

$$E\left[-2\left(x - (k + \varepsilon_k)\,r + \varepsilon_u\right)(k + \varepsilon_k)\right] = 0.$$

Expanding the term inside the expectation operator, we get

$$E\left[2\left(-x(k + \varepsilon_k) + (k + \varepsilon_k)^2 r - \varepsilon_u(k + \varepsilon_k)\right)\right] = 0.$$

Using the fact that $E[\varepsilon_k] = E[\varepsilon_u] = 0$ and that $\varepsilon_k$ and $\varepsilon_u$ are independent so $E[\varepsilon_k \varepsilon_u] = E[\varepsilon_u]\,E[\varepsilon_k] = 0$, the preceding equation reduces to

$$-2xk + 2r\left(k^2 + \sigma_k^2\right) = 0.$$

Rearranging and setting this expression to 0 yields the value of $r$ described in the text (on p. 39):

$$r = \frac{x}{k + \sigma_k^2\,/\,k}.$$

Note that $\frac{dr}{dx}$ is decreasing in $\sigma_k^2$; that is, greater uncertainty leads the policymaker to attenuate his response to shocks.

## APPENDIX 2. DERIVING THE ROBUST STRATEGY IN BRAINARD'S (1967) MODEL

This appendix derives the policy that solves equation 7 (on p. 46), that is,

$$\min_r \max_{\varepsilon_k \in \left[\underline{\varepsilon}, \bar{\varepsilon}\right]}\left\{\left(x - (k + \varepsilon_k)r\right)^2 + \sigma_u^2\right\}.$$

For ease of exposition, let us rewrite this problem as

$$\min_r W(r),$$

where

$$W(r) = \max_{\varepsilon_k \in \left[\underline{\varepsilon}, \bar{\varepsilon}\right]}\left\{\left(x - (k + \varepsilon_k)r\right)^2 + \sigma_u^2\right\}.$$

Since the second derivative of $(x - (k + \varepsilon_k)\,r)^2$ with respect to $\varepsilon_k$ is just $2r^2 \geq 0$, the function $W(r)$ is convex in $\varepsilon_k$. It follows that the maximum value must occur at one of the two endpoints of the support, that is, either when $\varepsilon_k = \underline{\varepsilon}$ or when $\varepsilon_k = \bar{\varepsilon}$.

Next, I argue that if $r^*$ solves equation 7, then we can assume without loss of generality that $W(r^*)$ takes on the same value when $\varepsilon_k = \bar{\varepsilon}$ as when $\varepsilon_k = \underline{\varepsilon}$. Suppose instead that at the value $r^*$ that solves equation 7, $W(r^*)$ is unequal at these two values. Let us begin first with the case where $W(r^*)\big|_{\varepsilon_k = \underline{\varepsilon}} > W(r^*)\big|_{\varepsilon_k = \bar{\varepsilon}}$. This implies

$$\left(x - (k + \underline{\varepsilon})r^*\right)^2 > \left(x - (k + \bar{\varepsilon})r^*\right)^2.$$

I obtain a contradiction by establishing there exists an $r \neq r^*$ that achieves a lower value of $W(r)$ than $W(r^*) = \left(x - (k + \underline{\varepsilon})r^*\right)^2 + \sigma_u^2$. If this is true, then $r^*$ could not have been the solution to equation 7, since the solution requires that $W(r^*) \leq W(r)$ for all $r$.

Differentiate $\left(x - (k + \underline{\varepsilon})r\right)^2$ with respect to $r$ and evaluate this derivative at $r = r^*$. If this derivative,

$2(k + \underline{\varepsilon})\left((k + \underline{\varepsilon})r^* - x\right)$, is different from 0, then we can change $r$ in a way that lowers $\left(x - (k + \underline{\varepsilon})r\right)^2$. By continuity, we can find an $r$ close enough to $r^*$ to ensure that

$$\left(x - (k + \underline{\varepsilon})r\right)^2 > \left(x - (k + \bar{\varepsilon})r\right)^2.$$

It follows that there exists an $r \neq r^*$ such that $W(r)$ $= \left(x - (k + \underline{\varepsilon})r\right)^2 + \sigma_u^2 < W(r^*)$; this is a contradiction. This leaves us with the case where $2(k + \underline{\varepsilon}) \times \left((k + \underline{\varepsilon})r^* - x\right)$ is equal to 0. If $(k + \underline{\varepsilon})r^* = x$, then we have

$$0 = \left(x - (k + \underline{\varepsilon})r\right)^2 > \left(x - (k + \bar{\varepsilon})r\right)^2 \geq 0,$$

which once again is a contradiction. The last remaining case involves $k + \underline{\varepsilon} = 0$. In that case, $W(r) = x^2$. But we can achieve this value by setting $r = 0$, and since $W(0)$ does not depend on $\varepsilon_k$, the statement follows trivially. That is, when $k + \underline{\varepsilon} = 0$ and $\varepsilon_k = \underline{\varepsilon}$ solves the maximization problem that underlies $W(r)$, there will be multiple solutions to equation 7, including one that satisfies the desired property.

The case where $W(r^*)\big|_{\varepsilon_k = \bar{\varepsilon}} > W(r^*)\big|_{\varepsilon_k = \underline{\varepsilon}}$ leads to a similar result, without the complication that $k + \bar{\varepsilon}$ might vanish to 0.

Equating the losses at $\varepsilon_k = \underline{\varepsilon}$ and $\varepsilon_k = \bar{\varepsilon}$, we have

$$\left(x - (k + \underline{\varepsilon})r\right)^2 = \left(x - (k + \bar{\varepsilon})r\right)^2,$$

$$x^2 + (k + \underline{\varepsilon})^2 r^2 - 2x(k + \underline{\varepsilon})r = x^2 + (k + \bar{\varepsilon})^2 r^2 - 2x(k + \bar{\varepsilon})r,$$

$$\left(k^2 + 2\underline{\varepsilon}k + \underline{\varepsilon}^2\right)r^2 - 2x(k + \underline{\varepsilon})r = \left(k^2 + 2\bar{\varepsilon}k + \bar{\varepsilon}^2\right)r^2 - 2x(k + \bar{\varepsilon})r,$$

$$\left(2k(\bar{\varepsilon} - \underline{\varepsilon}) + \bar{\varepsilon}^2 - \underline{\varepsilon}^2\right)r^2 = 2x(\bar{\varepsilon} - \underline{\varepsilon})r,$$

$$\left(2k + \bar{\varepsilon} + \underline{\varepsilon}\right)r^2 = 2xr.$$

The roots of this quadratic equation are $r = 0$ and $r = \dfrac{x}{k + (\underline{\varepsilon} + \bar{\varepsilon})/2}$. Substituting in each of these two values yields

$$W(r) = \begin{cases} x^2 + \sigma_u^2 & \text{if } r = 0 \\[2ex] \left[\dfrac{\bar{\varepsilon} - \underline{\varepsilon}}{2k + \underline{\varepsilon} + \bar{\varepsilon}}\right]^2 x^2 + \sigma_u^2 & \text{if } r = \dfrac{x}{k + (\underline{\varepsilon} + \bar{\varepsilon})/2} \end{cases}.$$

Define

$$\Phi = \frac{\bar{\varepsilon} - \underline{\varepsilon}}{2k + \underline{\varepsilon} + \bar{\varepsilon}}.$$

Of the two candidate values for $r$, the one where $r = 0$ will minimize $W(r)$ if $|\Phi| > 1$ and the one where

$r = \dfrac{x}{k + (\underline{\varepsilon} + \bar{\varepsilon})/2}$ will minimize $W(r)$ if $|\Phi| < 1$.

Since $\underline{\varepsilon} < 0 < \bar{\varepsilon}$, the numerator for $\Phi$ is always positive. Hence, $\Phi$ will be negative if and only if $2k + \underline{\varepsilon} + \bar{\varepsilon} < 0$. But if $2k + \underline{\varepsilon} + \bar{\varepsilon} < 0$, then $\Phi < -1$. The solution to equation 7 will thus set $r = 0$ when $\Phi < 0$. Since $\bar{\varepsilon} > 0$, a necessary condition for $\Phi < 0$ is for

$$\underline{\varepsilon} < -2k.$$

The only case in which the optimal policy will set $r \neq 0$ is if $0 \leq \Phi < 1$. This in turn will only be true if $2k + \underline{\varepsilon} + \bar{\varepsilon} > \bar{\varepsilon} - \underline{\varepsilon}$, which reduces to

$$\underline{\varepsilon} > -k.$$

This implies the robust strategy will set $r = 0$ whenever $\underline{\varepsilon} > -k$ but not otherwise.

---

## APPENDIX 3. DERIVING THE ROBUST RULE WITH UNCERTAIN PERSISTENCE

This appendix derives the policy that solves equation 12 (on p. 49). Substituting in for $y_t$, $x_t$, and $r_t = ax_{t-1}$, we can rewrite equation 12 as

A1) $\quad \min\limits_{a} \max\limits_{\rho} E\left[\sum\limits_{t=0}^{\infty} \beta^t \left[(\rho - a)x_{t-1} + \varepsilon_t\right]^2\right].$

Since $x_0$ is assumed to be drawn from the stationary distribution, the unconditional distribution of $x_t$ is the same for all $t$. In addition, $x_{t-1}$ and $\varepsilon_t$ are independent. Hence, equation A1 can be rewritten as

$$\min\limits_{a} \max\limits_{\rho} \sum\limits_{t=0}^{\infty} \beta^t \left[\frac{(\rho - a)^2 \sigma_\varepsilon^2}{1 - \rho^2} + \sigma_\varepsilon^2\right]$$

or just

$$\min\limits_{a} \max\limits_{\rho} \frac{1}{1 - \beta}\left[\frac{(\rho - a)^2 \sigma_\varepsilon^2}{1 - \rho^2} + \sigma_\varepsilon^2\right].$$

By differentiating the function $\dfrac{(\rho - a)^2}{1 - \rho^2}$, one can show that it is convex in $\rho$ for $\rho \in [-1,1]$ and $a \in [-1,1]$. Hence, the biggest loss will occur at the extreme values, that is, when $\rho$ is equal to either $\underline{\rho}$ or $\bar{\rho}$. By the same argument as in appendix 2, the robust strategy must set $a$ in order to equate the losses at these two extremes, which are given by $\dfrac{(\bar{\rho} - a)^2}{1 - \bar{\rho}^2}$ when $\rho = \bar{\rho}$ and to $\dfrac{(\underline{\rho} - a)^2}{1 - \underline{\rho}^2}$ when $\rho = \underline{\rho}$. If we equate these two expressions

and rearrange, the condition $a$ would have to satisfy

A2) $\quad \left(\dfrac{\bar{\rho} - a}{\underline{\rho} - a}\right)^2 = \dfrac{1 - \bar{\rho}^2}{1 - \underline{\rho}^2}.$

Given that $\underline{\rho}$ and $\bar{\rho}$ are symmetric around $\rho^* > 0$, the expression on the right-hand side of equation A2 is less than 1. Then the left-hand side of equation A2 must be less than 1 as well. Using the fact that $a \in [\underline{\rho}, \bar{\rho}]$, it follows that this in turn requires

$$\bar{\rho} - a < a - \underline{\rho}$$

or, upon rearranging,

$$a > (\underline{\rho} + \bar{\rho})/2 = \rho^*.$$

By contrast, if the policymaker knew $\rho = \rho^*$ with certainty, he would set $a = \rho^*$. The policymaker thus responds more aggressively to past shocks with uncertainty than if he knew $\rho$ was equal to the midpoint of the set with certainty.

Finally, note that if $\rho^* = 0$, the symmetry requirement would imply $\underline{\rho} = -\bar{\rho}$, in which case the solution to equation A2 would imply $a = 0 = \rho^*$. Hence, the aggressive response stems from the fact that $\rho^*$ is assumed to be strictly positive; that is, the region of uncertainty is parameterized to be asymmetric with respect to no autocorrelation. This asymmetry drives the result that the robust strategy is more aggressive under uncertainty than under certainty.

REFERENCES

**Ball, Laurence,** 1999, "Policy rules for open economies," *Monetary Policy Rules*, John Taylor (ed.), National Bureau of Economic Research Studies in Business Cycles, Vol. 31, Chicago: University of Chicago Press, pp. 127–144.

**Becker, Selwyn, and Fred Brownson,** 1964, "What price ambiguity? Or the role of ambiguity in decision-making," *Journal of Political Economy*, Vol. 72, No. 1, February, pp. 62–73.

**Bellman, Richard,** 1956, "Minimization problem," *Bulletin of the American Mathematical Society*, Vol. 62, No. 3, p. 270.

**Bernanke, Ben,** 2007, "Monetary policy under uncertainty," speech at the 32nd Annual Economic Policy Conference, Federal Reserve Bank of St. Louis, October 19, available at www.federalreserve.gov/newsevents/speech/bernanke20071019a.htm.

**Brainard, William,** 1967, "Uncertainty and the effectiveness of policy," *American Economic Review*, Vol. 57, No. 2, May, pp. 411–425.

**Ellsberg, Daniel,** 1961, "Risk, ambiguity, and the Savage axioms," *Quarterly Journal of Economics*, Vol. 75, No. 4, November, pp. 643–669.

**Epstein, Larry, and Martin Schneider,** 2003, "Recursive multiple-priors," *Journal of Economic Theory*, Vol. 113, No. 1, November, pp. 1–31.

**Finch, Steven, and John Wetzel,** 2004, "Lost in a forest," *American Mathematical Monthly*, Vol. 111, No. 8, October, pp. 645–654.

**Giannoni, Marc,** 2002, "Does model uncertainty justify caution? Robust optimal monetary policy in a forward-looking model," *Macroeconomic Dynamics*, Vol. 6, No. 1, February, pp. 111–144.

**Gilboa, Itzhak, and David Schmeidler,** 1989, "Maxmin expected utility with nonunique prior," *Journal of Mathematical Economics*, Vol. 18, No. 2, April, pp. 141–153.

**Gluss, Brian,** 1961, "An alternative solution to the 'lost at sea' problem," *Naval Research Logistics Quarterly*, Vol. 8, No. 1, March, pp. 117–121.

**Hansen, Lars, and Thomas Sargent,** 2008, *Robustness*, Princeton, NJ: Princeton University Press.

**Isbell, John,** 1957, "An optimal search pattern," *Naval Research Logistics Quarterly*, Vol. 4, No. 4, December, pp. 357–359.

**Maccheroni, Fabio, Massimo Marinacci, and Aldo Rustichini,** 2006, "Dynamic variational preferences," *Journal of Economic Theory*, Vol. 128, No. 1, May, pp. 4–44.

**Manski, Charles,** 2000, "Identification problems and decisions under ambiguity: Empirical analysis of treatment response and normative analysis of treatment choice," *Journal of Econometrics*, Vol. 95, No. 2, April, pp. 415–442.

**Onatski, Alexi, and James Stock,** 2002, "Robust monetary policy under model uncertainty in a small model of the U.S. economy," *Macroeconomic Dynamics*, Vol. 6, No. 1, February, pp. 85–110.

**Rustem, Berc, Volker Wieland, and Stan Zakovic,** 2007, "Stochastic optimization and worst-case analysis in monetary policy design," *Computational Economics*, Vol. 30, No. 4, November, pp. 329–347.

**Sargent, Thomas,** 1999, "Comment on 'Policy rules for open economies' by Laurence Ball," in *Monetary Policy Rules*, John Taylor (ed.), National Bureau of Economic Research Studies in Business Cycles, Vol. 31, Chicago: University of Chicago Press, pp. 144–154.

**Sims, Christopher,** 2001, "Pitfalls of a minimax approach to model uncertainty," *American Economic Review*, Vol. 91, No. 2, May, pp. 51–54.

**Spark, Nick,** 2006, *A History of Murphy's Law*, Los Angeles: Periscope Film.

**Svensson, Lars,** 2007, "Robust control made simple: Lecture notes," Princeton University, February.

**Wald, Abraham,** 1950, *Statistical Decision Functions*, New York: Wiley.

**Williams, Noah,** 2008, "Robust control," in *The New Palgrave Dictionary of Economics*, Steven Durlauf and Lawrence Blume (eds.), 2nd ed., 8 vols., New York: Palgrave Macmillan.