1 Targets, Instruments, and All That

1 Introduction

I realize that these are the Robbins lectures, not the Ricardo lectures. But please pardon a momentary digression on comparative advantage nonetheless, for I have long believed that one true test of whether a person is an economist is how devoutly he or she lives by the principle of comparative advantage. And I don't mean just *preaching* it, but actually *practicing* it. For example, I always harbor doubts about my economist friends who tell me that they mow their own lawns, rather than hiring a gardener, because they actually enjoy cutting grass. Such a claim is suspect on its face. But, more to the point, a true believer in comparative advantage should be constitutionally incapable of enjoying such activities; the David Ricardo inside him should make him feel too guilty.

As a devotee of comparative advantage, the topic of these lectures virtually chose itself. Our profession boasts greater economic theorists and more skilled econometricians than I. But there must be relatively few people on earth who have been as deeply immersed in monetary policy from both the

academic and central banking sides as I have. Therein, I presume, lies my comparative advantage and the topic of these three lectures: the theory and practice of central banking.

To keep things manageable, I have pared the topic beyond what a literal reading of the title may suggest. First, central bankers, I can assure you, are busy with many matters that are related tangentially if at all to monetary policy—such as managing the payments system and supervising banks. But I will stick to monetary policy proper. Second, I will deal much more with the behavior of central banks than with the monetary transmission mechanism. In these lectures, short-term interest rates are more often left-hand than right-hand variables.

The various issues covered in these three lectures are parts of a mosaic that could be taken up in many different orders. But a lecturer must draw some boundary lines, artificial though they may be, in order to divide the subject matter into lecture-sized portions. This I have done in the following way. The first lecture deals mainly with a variety of complications that a practical central banker must confront in trying to implement the classical targets-instruments approach. This presentation begs two important questions that I will address in the second lecture: What policy instrument should the central bank use? And should it attempt discretionary policy at all, rather than just relying on a simple rule? Finally, the last lecture is devoted to various positive and, especially, normative aspects of central bank independence.

As these lectures progress, it will become apparent that central banking looks rather different in practice than it does in theory. Having seen it from both sides now, I deeply

believe that both theory and practice could benefit from greater contact with and greater understanding of the other. Hence, I will periodically point out opportunities for cross-fertilization—places where central bankers have more to learn from academic research, and places where academic economists could profit from greater awareness of the practical world of central banking. Arbitrage should take care of the rest!

2 Targets and Instruments: The Rudiments

Monetary policymakers have certain objectives—such as low inflation, output stability, and perhaps external balance—and certain instruments to be deployed in meeting their responsibilities, such as bank reserves or short-term interest rates. Unless it has only a single goal,[1] the central bank is forced to strike a balance among competing objectives, that is, to face up to various *trade-offs*. Unless your education in economics is very thin (or very recent!), these two sentences immediately bring to mind Tinbergen (1952) and Theil (1961). So let us begin there, at the beginning.

In theory, it works like this. There is a known model of the macroeconomy, which I write in structural form as:

$$y = F(y,x,z) + e \tag{1}$$

and in reduced form as:

$$y = G(x,z) + e. \tag{2}$$

Here y is the vector of endogenous variables (a few of which are central bank objectives), x is the vector of policy instruments (which may be of size one), and z is the vector of

nonpolicy exogenous variables. The vector e of stochastic disturbances will fade in importance once I assume, with Tinbergen and Theil, that $F(.)$ is linear and the policymaker's objective function,

$$W = W(y), \tag{3}$$

is quadratic. In principle, the policymaker maximizes the expected value of (3) subject to the constraint (2) to derive an optimal policy "rule":

$$x^* = H(z). \tag{4}$$

All very simple.

What's wrong with this simple framework? Both nothing and everything. Starting with "nothing," I do believe that—once you add a host of complications, several of which I will discuss in this lecture—this is the right way for a central banker to think about monetary policy. You have an economy. Except for the policy instruments you control, you must accept it as it is. You also have multiple objectives—your own, or those assigned to you by the legislature—and you must weigh them somehow, though perhaps not quadratically. To a significant extent, though usually quite informally and to my mind not quite sufficiently, central bankers do think about policy this way.

But, as is well known, there are many complications. Let me list just a few, some of which I will dwell on at length in the balance of this lecture and the next:

1. *Model uncertainty:* In practice, of course, we do not know the model but must estimate it econometrically. Since econo-

mists agree neither on the "right" model nor on the "right" econometric techniques, this is a nontrivial problem. It means, among other things, that policy multipliers—the derivatives of $G(.)$ with respect to x—are subject to considerable uncertainty.

2. *Lags:* Any reasonable macroeconometric model will have a complex lag structure that is ignored by (1). This is not much of a problem in principle because, as all graduate students learn, this complication can be accommodated formally simply by appending further equations for lagged variables (see Chow 1975). However, in practice it creates serious difficulties that bedevil policymakers.

3. *Need for forecasts:* Because of lags, execution of the Tinbergen–Theil framework requires forecasts of the future paths of the exogenous variables—in principle, the entire z vector, which may be quite long. Such forecasts are neither easy to generate nor particularly accurate.

4. *Choice of instrument:* The Tinbergen–Theil framework takes as given that some variables are endogenous and others are policy instruments. In most cases, however, the central bank has at least some latitude, and maybe quite a lot, in choosing its instrument(s). One way of thinking about this is that some of the xs and ys can trade places at the discretion of the central bank. For example, the short-term interest rate can be the policy instrument and bank reserves an endogenous variable; or the central bank can do things the other way around. Some economists take this idea a little too far and write models in which central bankers can control, say, nominal GDP, the inflation rate, or the

unemployment rate on a period-by-period basis. Believe me, they cannot. If they could, monetary policy would be a good deal simpler than it is.

5. *The objective function:* The next problem can be framed as a question: Who supplies the objective function? The answer, typically, is: no one. The political authorities, who after all should decide such things, rarely if ever give such explicit instructions to their central banks. So central bankers must—in a figurative, not literal, sense—create their own social welfare function based on their legal mandate, their own value judgments, and perhaps their readings of the political will. This last thought brings up the independence of the central bank—a subject I will take up in depth in the third lecture.

Adding it all up, a curmudgeon could summarize the problems with applying the Tinbergen–Theil program as follows: We do not know the model, and we do not know the objective function, so we cannot compute the optimal policy rule. To some critics of "impractical" or "theoretical" economics, including some central bankers, this criticism is a show-stopper. But, speaking now as a former central banker, I think such know-nothingism is not a very useful attitude. In fact, in my view, we must use the Tinbergen–Theil approach—with as many of the complications as we can handle—even if in a quite informal way. An analogy will explain why.

Consider your role as the owner of an automobile. You have various objectives toward which the use of your car contributes, such as getting to work, shopping, and going

on pleasure trips. You do not literally "know" the utility function that weighs these objectives, but you presumably wish to maximize it nonetheless. The care and feeding of your car entails considerable expense, and you have great uncertainty about the "model" that maps inputs like gasoline, oil, and tires into outputs like safe, uneventful trips. Furthermore, there are substantial, stochastic lags between maintenance expenditures (e.g., frequent oil changes) and their payoff (e.g., greater engine longevity).

What do you do? One alternative is the "putting out fires" strategy: Do nothing for your car until it breaks down, then fix whatever is broken and continue driving until something else breaks down. I submit that few of us follow this strategy because we know it will produce poor results.[2] Instead, we all follow something that approximates—philosophically, if not mathematically—the Tinbergen–Theil framework. Central banks do, too. Or at least they should, for they will surely fail in their stabilization-policy mission if they simply "put out fires" as they observe them. Let me review briefly how the Tinbergen–Theil framework is used in practice.

To begin with, there must be a macro model. It need not be a system of several hundred stochastic difference equations, though that is not a bad place to start. In fact, no central bank that I know of, and certainly not the Federal Reserve, is wed to a single econometric model of its economy. Some banks have such models, and some do not. But, even if they do not, or do not use it, *some* kind of a model—however informal—is necessary to do policy, for otherwise how can you even begin to estimate the effects of changes in policy instruments?

Some central bankers scoff at large-scale macroeconometric models, as do some academic economists. And their reasons are not all that dissimilar. Many point, for example, to the likelihood of structural change in any economy over a period of several decades, which casts doubt on the stationarity assumptions that underlie standard econometric procedures and thus on the bedrock notion that the past is a guide to the future. Others express skepticism that something as complex as an entire economy can be captured in any set of equations. Still other critics emphasize a host of technical problems in time series econometrics that cast doubt on any set of estimated coefficients. Finally, some central bankers simply do not understand these ungainly creatures at all and doubt that they should be expected to.

Leaving aside the last, there is truth in each of these criticisms. Every model is an oversimplification. Economies do change over time. Econometric equations often fail subsample stability tests. Econometric problems like simultaneity, common trends, and omitted variables are ubiquitous in nonexperimental data. The Lucas critique warns us that some parameters may change when policy does.[3] Yet what are we to do about these problems? Be skeptical? Of course. Use several methods and models instead of just one? Certainly. But abandon all econometric modeling? I think not. The criticisms of macroeconometrics are not wrong, but their importance is often exaggerated and their implications misunderstood. These criticisms should be taken as warnings—as calls for caution, humility, and flexibility of mind—not as excuses to retreat into econometric nihilism. It is foolish to make the best the enemy of the moderately useful.

Indeed, I would go further. I don't see that central bankers even have the luxury of ignoring econometric estimates. Monetary policymaking requires more than just the qualitative information that theory provides—e.g., that if short-term interest rates rise, real GDP growth will subsequently fall. They must have quantitative information about magnitudes and lags, even if that information is imperfect.

I often put the choice this way: You can get your information about the economy from admittedly fallible statistical relationships, or you can ask your uncle. I, for one, have never hesitated over this choice. But I fear there may be altogether too much uncle-asking in government circles in general, and in central banking circles in particular. For example, a long line of politicians will assure you that a lower capital gains tax rate will spur investment. The only trouble is that no evidence supports this contention. Similarly, central bankers often take it as axiomatic that long-term interest rates are good forecasters of either (a) inflation or (b) future short-term interest rates. Unfortunately, the data refute both claims.

3 Uncertainties: Models and Forecasts

Let me now turn to the first of three important amendments to the Tinbergen–Theil framework, beginning with the obvious fact that no one knows the "true model." It would hardly have been news to Tinbergen and Theil that both models and forecasts of exogenous variables are subject to considerable uncertainties. And subsequent developments by economists have provided ways of handling or finessing

these gaps in our knowledge.[4] Let us consider, very briefly, three types of uncertainty.

Uncertainty about forecasts: In the linear-quadratic case, uncertainty about the values of future exogenous variables is no problem *in principle;* you need only replace unknown future variables with their expected values (the "certainty equivalence" principle). But here is one case in which the gap between theory and practice is huge because the task of generating unbiased forecasts of dozens or even hundreds of exogenous variables is a titanic practical problem. It is, for example, a major reason why large-scale econometric models are not terribly useful as forecasting tools.[5]

Skeptics often object to certainty equivalence on the grounds that (a) the economy is nonlinear and (b) there is no particular reason to think that the objective function is quadratic. Both objections are undoubtedly true and, if taken literally, invalidate the certainty-equivalence principle. But I think the importance of this point is often exaggerated by those who would denigrate the usefulness—and thereby escape the discipline—of formal econometric models. Policymakers almost always will be contemplating changes in policy instruments that can be expected to lead to small changes in macroeconomic variables. For such changes, any model of an economy is approximately linear and any convex objective function is approximately quadratic.[6] So this problem of principle is, in my view, of great practical importance only on those rare occasions when large changes in policy are contemplated.

Uncertainty about parameters: Uncertainty about parameters, and hence about policy multipliers, is much more

difficult to handle, even at the conceptual level. Certainty equivalence certainly does not apply. While there are some fairly sophisticated techniques for dealing with parameter uncertainty in optimal control models with learning, those methods have not attracted the attention of either macroeconomists or policymakers. There is good reason for this inattention, I think: You don't conduct experiments on a real economy solely to sharpen your econometric estimates.

There is, however, one oft-forgotten principle that I suspect practical central bankers can—and in a rough way do—rely on. Many years ago, William Brainard (1967) demonstrated that, under certain conditions,[7] uncertainty about policy multipliers should make policymakers *conservative* in the following specific sense: They should compute the direction and magnitude of their optimal policy move in the way prescribed by Tinbergen–Theil and then do less.

Here is a trivial adaptation of Brainard's simple example. Simplify equation (2) to:

$$y = Gx + z + e, \tag{2'}$$

and suppose that G and z are independent random variables with means g and $\bar{z}$ respectively. The policymaker wishes to minimize $E(y - y^*)^2$. Interpret $z + e$ as the value of y in the absence of any further policy move ($x = 0$) and x as the contemplated change in policy.[8] If G is nonrandom, the optimal policy adjustment is certainty equivalence:

$$x = (y^* - \bar{z})/G,$$

that is, fully closing the expected gap between y^* and $\bar{z}$. But if G is random with mean g and standard deviation σ, the loss function is minimized by setting:

$$x = \frac{y^* - \bar{z}}{g + \frac{\sigma^2}{g}},$$

which means that policy aims to fill only part of the gap.

My intuition tells me that this finding is more general—or at least more wise—in the real world than the mathematics will support.[9] And I certainly hope it is, for I can tell you that it was never far from my mind when I occupied the Vice Chairman's office at the Federal Reserve. In my view as both a citizen and a policymaker, a little stodginess at the central bank is entirely appropriate.

Uncertainty over model selection: Parameter uncertainty, while difficult, is at least a relatively well defined problem. Selecting the right model from among a variety of nonnested alternatives is another matter entirely. While there is some formal literature on this problem,[10] I think it is safe to say that central bankers neither know nor care much about this literature. I leave it as an open question whether they are missing much.

My approach to this problem while on the Federal Reserve Board was relatively simple: Use a wide variety of models and don't ever trust any one of them too much. So, for example, when the Federal Reserve staff explored policy alternatives, I always insisted on seeing results from (a) our own quarterly econometric model, (b) several alternative econometric models, and (c) a variety of vector autoregressions (VARs) that I developed for this purpose. My usual procedure was to simulate a policy on as many of these models as possible, throw out the outlier(s), and average the rest to get a point estimate of a dynamic multiplier path.

This can be viewed as a rough—make that very rough—approximation to optimal information processing.[11] As they say: Good enough for government work!

4 Lags in Monetary Policy

It is a commonplace that monetary policy operates on the economy with "long and variable lags." As I noted previously, the formalism of the Tinbergen–Theil framework can readily accommodate distributed lags. The costs are two-fold. First, the dimensionality of the problem increases; but with modern computing power this is not much of a problem. Second, the optimization problem changes from one of calculus to one of dynamic programming.[12] This latter point is significant in practice and, I think, inadequately appreciated by practitioners.

A dynamic programming problem is typically "solved backward," that is, if T is the final period and x is the policy instrument, you first solve a one-period optimization problem for period T, thereby deriving ${}_tx_T$ conditional on a past history. (The postscript denotes calendar time and the prescript denotes the date at which the expectation is taken.) Then, given your solution for ${}_tx_T$, which most likely depends *inter alia* on ${}_tx_{T-1}$, you solve a two-period problem for ${}_tx_T$ and ${}_tx_{T-1}$ jointly. Proceeding similarly, by a process of backward induction you derive an entire *solution path*:

$$x_t, {}_tx_{t+1}, {}_tx_{t+2}, \ldots, {}_tx_T.$$

Don't get me wrong. I do not believe it is important for central bankers to acquire any deep understanding of Bellman's principle, still less of the computational techniques

used to implement it. What really matters for sound decisionmaking is the way dynamic programming teaches us to think about intertemporal optimization problems—and the discipline it imposes. It is essential, in my view, for central bankers to realize that, in an dynamic economy with long lags in monetary policy, today's monetary policy decision must be thought of as the first step along a path. The reason is simple: Unless you have thought through your expected future actions, it is impossible to make today's decision rationally. For example, when a central bank begins a cycle of either tightening or easing, it should have some idea about where it is going before it takes the first step.

Of course, by the time period $t + 1$ rolls around, the policymaker will have new information and may wish to change his or her mind about the earlier tentative decision ${}_{t}x_{t+1}$. That is fine. In fact, given the information then available, the policymaker will want to plan an entirely new path:

$$x_{t+1},\ {}_{t+1}x_{t+2},\ {}_{t+1}x_{t+3},\ \ldots,\ {}_{t+1}x_{T}$$

But that realization in no way obviates the need to think ahead in order to make today's decision—which is the important lesson of dynamic programming. It is an intensely practical lesson and, I believe, one that is inadequately understood.[13]

Too often decisions on monetary policy—and, indeed, on other policies—are taken "one step at a time" without any clear notion of what the next several steps are likely to be. In central banking circles, it is often claimed that such one-step-at-a-time decisionmaking is wise because it maintains "flexibility" and guards against getting "locked in" to decisions the central bank will later regret. I often heard sentiments like this expressed both at FOMC meetings and at international meetings of central bankers.

But this attitude reflects a fundamental misunderstanding of the way dynamic programming teaches us to think. It is absolutely correct that flexibility should be maintained and that locking yourself in should be avoided. But both of these notions are inherent in dynamic programming. If there are any surprises at all, the decisions that you actually carry out in the future will differ from the ones you originally planned. That's flexibility. Ignoring your own likely future actions is myopia.

These matters are really quite intuitive. Despite their lack of understanding of the fine points of the calculus of variations, ordinary rational people do not deem it wise to ignore the admittedly unknown future in order to "maintain flexibility." Think, for example, about students formulating educational and career plans. In choosing a major, and sometimes even in choosing a college, many undergraduates are looking ahead to their ultimate career objectives. They know their crystal ball is cloudy, and they realize that they may have many reasons to change their minds along the way. But they nonetheless find it rational to plan ahead when making the initial decision. And they are right.

Applying this abstract discussion to a concrete problem in monetary policy may help resolve a long-standing issue in central banking. Policymakers in the United States and elsewhere have often been accused of making a particular type of systematic error in the timing of policy changes. Specifically, it is alleged that they overstay their policy stance—whether it is tightening or loosening—thereby causing overshoots in both directions.[14] I believe this criticism

may be correct, although I know of no systematic study that demonstrates it. I furthermore believe that the error, if it exists, may be due to following a strategy I call "looking out the window."

The error is well illustrated by what I call the parable of the thermostat. The following has probably happened to you; it has certainly happened to me. You arrive at night in an unfamiliar hotel and find the room temperature too cold. So you turn up the heat and take a shower. Emerging 10 minutes later, you still find the room too cold. So you turn the heat up another notch and go to sleep. At about 2 a.m. you wake up in a pool of sweat in a room that is oppressively hot.

By analogy, a central bank following the "looking out the window" strategy proceeds as follows. For concreteness, suppose it is in the process of tightening. At each decision-making juncture, the bank takes the economy's temperature and, if it is still too hot, tightens monetary conditions another notch. Given the long lags in monetary policy, you can easily see how such a strategy can keep the central bank tightening for too long.

Now compare "looking out the window" to proper dynamic optimization. Under dynamic programming, at each stage the bank would project an entire path of future monetary policy actions, with associated paths of key economic variables. It would, of course, act only on today's decision. Then, if things evolved as expected, it would keep following its projected path, which would be likely (given the lags in monetary policy) to tell it to stop tightening while the economy was still "hot." Of course, economies rarely evolve as expected. Surprises are the norm, not the exception, and

they would induce the central bank to alter its expected path in obvious ways. If the economy steamed ahead faster than expected, the bank would tighten more. If the economy slowed down sooner than expected, the bank would tighten less or even reverse its stance.

Do central banks actually behave this way? Yes and no. Like a skilled billiards player who does not understand the laws of physics, a skilled practitioner of monetary policy may follow a dynamic-programming-type strategy intuitively and informally. In the last few years, for example, the notion that it is wise to pursue a strategy of "preemptive strikes" against inflation seems to have caught on among central bankers. The main impetus for this change in fashion was, I believe, the leadership and perceived success of the Federal Reserve in first tightening monetary policy "preemptively" in early 1994 and then achieving the fabled "soft landing." By now, a variety of other central banks are talking the same talk. But the very fact that this style of decision-making was perceived to be a great advance suggests that the dynamic programming way of thinking has not yet permeated central banking circles.

A preemptive strategy implies a certain amount of confidence in both your forecast and your model of how monetary policy affects the economy, both of which are hazardous. But preemption does not require too much confidence. Remember the flexibility principle of dynamic programming and the Brainard conservatism principle. Taken together, they lead to the following sort of strategy:[15]

Step 1. Estimate how much you need to tighten or loosen monetary policy to "get it right." Then do less.

Step 2. Watch developments.

Step 3a. If things work out about as expected, increase your tightening or loosening toward where you thought it should be in the first place.

Step 3b. If the economy seems to be evolving differently from what you expected, adjust policy accordingly.

Two final points about preemptive monetary policy are worth making. First, a successful stabilization policy based on preemptive strikes will appear to be misguided and may therefore leave the central bank open to severe criticism. The reason is simple. If the monetary authority tightens so early that inflation never rises, the preemptive strike is a resounding success, but critics of the central bank will wonder—out loud, no doubt—why the bank decided to tighten when the inflationary dragon was nowhere to be seen. Similarly, a successful preemptive strike against economic slack will prevent unemployment from rising and leave critics complaining that the authorities were hallucinating about rising unemployment. Precisely these criticisms of the Fed's tightening in 1994–1995 and subsequent easing in 1995–1996 were heard in the United States in recent years.

Second, the logic behind the preemptive strike strategy is symmetric. The same reasoning that tells a central bank to get a head start against inflation says it should also strike preemptively against rising unemployment. That is why Chairman Alan Greenspan told Congress in February 1995, just after the Fed had completed a year-long tightening cycle that raised short-term interest rates 300 basis points, that: "There may come a time when we hold our policy stance unchanged, or even ease, despite adverse price data, should we see signs that underlying forces are acting ultimately to

reduce inflationary pressures."[16] In fact, the statement itself amounted to a monetary easing, since it fueled a bond-market rally well before the Fed started cutting interest rates (which did not occur until July 1995). Notably, both Greenspan's statement and the Fed's interest-rate cut in July 1995 came while the unemployment rate was below contemporary estimates of the natural rate.

Under what circumstances might the preemptive strike strategy apply more to fighting inflation than to fighting unemployment?

First, if the short-run Phillips curve is distinctly nonlinear in the way Phillips originally drew it, so that low unemployment raises inflation more than high unemployment lowers it. But, with due apologies to those notable curve-fitting exercises done at the London School of Economics in the 1950s, the U.S. evidence is decidedly against this hypothesis. A linear Phillips curve fits the data extremely well,[17] and tests for nonlinearity suggest, if anything, a concave (to the origin) Phillips curve rather than with a convex one.[18]

Second, the central bank's loss function could attach much more weight to inflation than to unemployment—as some observers of central banking have suggested, and as some central bank charters (but not the Fed's) mandate.

Third, lags in monetary policy could be longer for inflation fighting than for unemployment fighting, calling for earlier preemption in the former case. This last circumstance appears to obtain, and may be the main justification for acting more preemptively against inflation than against unemployment.

Note, however, that political considerations most likely point in the opposite direction. In most situations, the central bank will take far more political heat when it tightens pre-

emptively to avoid higher inflation than when it eases preemptively to avoid higher unemployment.

5 Central Banking by Committee

So far, I have offered one explanation for the alleged tendency of central banks to overstay their stance—remaining tight for too long, thereby causing recessions, and remaining easy for too long, thereby allowing inflation to take root: a failure to internalize the dynamic-programming way of thinking. But a prominent institutional feature of some central banks (including the Federal Reserve) may also contribute to this problem. Specifically, in many countries monetary policy is made not by a single individual but by a *committee*.

While serving on the FOMC, I was vividly reminded of a few things all of us probably know about committees: that they laboriously aggregate individual preferences; that they need to be led; that they tend to adopt compromise positions on difficult questions; and—perhaps because of all of the above—that they tend to be inertial. Had Newton served on more faculty committees at Cambridge, his first law of motion might have read: A decisionmaking body at rest or in motion tends to stay at rest or in motion in the same direction unless acted upon by an outside force.

Inertial behavior has its virtues, as I will explain shortly. But it also has some vices. In particular, decisionmaking by committee may contribute to the systematic policy errors I have mentioned already by inducing the central bank to maintain its policy stance too long.

While the Federal Open Market Committee has not been immune to this ailment over the years, there is at least one

tradition at the Federal Reserve that tends to minimize it: that of the powerful chairman. The law says that each of the 12 voting members of the FOMC has one vote. But no one has ever doubted that Alan Greenspan, or Paul Volcker, or Arthur Burns were "more equal" than the others. The Chairman of the Federal Reserve Board is virtually never on the losing side of a monetary policy vote. So, to a significant extent, FOMC decisions are *his* decisions, as tempered by the opinions of the other members. Nonetheless, a chairman who needs to build consensus may have to move more slowly than if he were acting alone.

Now for the positive side. America is the land of checks and balances. Our political traditions harbor great fear of unbridled, centralized power. It is an anti-government form of government—the little government that couldn't because it was too tied up in knots. Yet the Federal Open Market Committee has virtually total freedom to do as it pleases with monetary policy—without asking permission from any other branch of government and with little fear of being countermanded. So long as FOMC decisions are done by the book and remain within the Fed's legal authority, the committee is neither checked nor balanced—at least not externally.

But the group nature of FOMC decisions creates what amounts to an *internal* system of checks and balances. No chairman can deviate too far from the view that prevails in his committee. Decisionmaking by committee, especially when there is a strong tradition of consensus, makes it very difficult for idiosyncratic views to prevail.[19] So monetary policy decisions tend to regress toward the mean and to be inertial—and hence biased in just the same way that adap-

tive expectations are biased relative to rational expectations. But errors like that, while systematic, will generally be small and will tend to shrink over time. And, in return, the system builds in natural safeguards against truly horrendous mistakes.

I leave it to some clever theorist to prove that the FOMC is an example of optimal institutional design. My own hunch is that, on balance, the additional monetary policy inertia imparted by group decisionmaking provides a net benefit to society. It does, at least, provide something of a check against an overzealous Fed chairman. But my main point is simpler: My experience as a member of the FOMC left me with a strong feeling that the theoretical fiction that monetary policy is made by a single individual maximizing a well-defined preference function misses something important. In my view, monetary theorists should start paying some attention to the nature of decisionmaking by committee, which is rarely mentioned in the academic literature.[20]

6 Conclusion

Overall, however, the message of this lecture is rather cheerful. Working in their cloistered universities, Tinbergen, Theil, Brainard, and others taught valuable abstract lessons that turned out to be of direct practical use in central banking. So did other scholars who developed their ideas further, pointed out additional complexities, and brought more powerful technical tools to bear—such as macroeconometric models and dynamic programming. Their ideas do not provide pat answers for central bankers, and their techniques cannot be applied mechanically. The real world is much too

complicated for that. So there must be both art and science in central banking. Nonetheless, the science is still useful; at least I found it so while on the Federal Reserve Board.

In fact, I think central bankers could learn a good deal more from the academics. For example, I emphasized that the dynamic programming way of thinking is not sufficiently ingrained into the habits of monetary policymakers, who too often just "look out the window" and base policy judgments on present circumstances. I believe this is a fundamental mistake and is one reason why central banks often overstay their policy stance.

There also seems to be much too much reliance on "uncle asking," relative to econometric evidence, among practical central bankers. Skepticism about econometric estimates is one thing, and is highly appropriate. But healthy skepticism should not be allowed to devolve into econometric nihilism, which is too often an excuse for wishful thinking and an escape from the discipline of the data.

But please don't think I believe that all wisdom resides in universities, from where it flows, somewhat impeded, to central banks. The next lecture should dispel that notion completely. After dealing briefly with one more issue where academic thinking was both correct and triumphant—the choice of monetary instrument—I will turn in detail to the rules versus discretion debate, where I will argue that much recent academic research has been barking up the wrong—or, rather, nonexistent—trees. On that issue, I believe, the academics must learn from the central bankers, and the sooner the better.