Algorithmic Social Sciences Research Unit

ASSRU

Department of Economics
University of Trento
Via Inama 5
381 22 Trento Italy

Discussion Paper Series

11 – 01

Economics, Control Theory, and the Phillips Machine

Brian Hayes

March 2011

Economics, Control Theory, and the Phillips Machine

Brian Hayes
Senior Writer, *American Scientist*
brian@bit-player.org

Abstract

Can the same mathematical control laws that smooth out oscillations in the flight of an airplane also moderate economic cycles of boom and bust? Attempts to bring together the intellectual traditions of control engineering and economics go back at least as far as the hydraulic analog computer of A. W. H. Phillips, circa 1950. Today, economic policymakers remain committed to the ideal of controlling business cycles; it remains an open question whether tools from control theory might help to refine their strategies.[1]

1 Engineering the Economy

In November 1949, faculty, students and guests at the London School of Economics gathered to observe a demonstration. At the front of the room was a two-meter-tall contraption assembled out of transparent plastic pipes, tanks, valves and other plumbing hardware. The device, later dubbed the MONIAC, was a hydraulic analog computer for modeling the flow of money through a national economy. When the machine was powered up, colored water gurgled through the tubes and sloshed into reservoirs. Various streams represented consumption, investment, taxes, savings, imports and exports. Crank-wheels and adjustable cams and weirs allowed the water levels and flows to be regulated—the hydraulic equivalent of setting monetary and fiscal policies. This was real trickle-down economics!

The principal architect—and plumber—of the MONIAC was A. W. H. Phillips, a New Zealander who had been an electrical engineer before he launched a new career in the social sciences and economics; at the time of the 1949 debut of the MONIAC he was a Ph.D. candidate in economics (Leeson 2000, Bissell 2007). Very likely it was Phillips's engineering background that led him to choose such a mechanistic way of illustrating economic principles. The hydraulic simulation emphasizes that the circulation of money through a society obeys deterministic, mathematical laws, like those that govern fluids and other physical systems. And the crank-wheels and cams on the MONIAC imply that the behavior of an economy is not only predictable but also controllable. If we twiddle the knobs and nudge the levers in just the right way,

[1]This article is based in part on an essay published in *American Scientist* as (Hayes 2009). Another version was presented as a talk at the conference Celebrating the 60th Anniversary of the Phillips National Income Electro-Hydraulic Analogue Machine, held at the University of Trento in December 2010. Conversations with the organizers and participants in that conference inspired substantial revisions, but the views expressed herein remain the authors' own.

all the streams will flow smoothly and the various basins where wealth accumulates will never run dry or overflow.

In later publications, Phillips made the engineering context of his economic ideas more explicit (Phillips 1954, 1957). He interpreted the Keynesian methodology for moderating business cycles as an exercise in control theory, the branch of engineering and applied mathematics concerned with feedback systems. Thus unwanted fluctuations in economic activity were to be brought under control by the same kinds of techniques that smooth out oscillations in the flight of an airplane or regulate the operation of a petroleum refinery.

A number of others also took up the notion of applying control theory to economic problems. Arnold Tustin, an engineer and control theorist at the University of Birmingham, had published a book on the theme in 1953 (Tustin 1953). A decade later the control theorists David Livesey of Cambridge University and Robert Pindyck of M.I.T. began working on macroeconomic problems. There were crossovers in the other direction as well: In 1968 half a dozen graduate students and young faculty members in the economics department at Harvard University enrolled in a control-theory course taught by Arthur Bryson (Athans and Kendrick 1974, Kendrick 2005); two of those economists, Lance Taylor and David Kendrick, subsequently published a macroeconomic development model based on control-theory methods. By the 1970s control theorists and economists were coming together in joint workshops and conferences, and even the U.S. Federal Reserve announced that it was testing control-theory methods in setting monetary policy.

The budding romance between control theory and economics did not last long, for reasons I shall mention below. On the other hand, efforts to assert control over the wilder excursions of business cycles certainly have *not* ended. In the aftermath of the financial crisis of 2008, governments and central banks enacted stimulus programs of staggering magnitude, using every available instrument of monetary and fiscal policy to spur production and consumption. As a society, it seems we are committed to the idea that cycles of boom and bust are subject to control. The question I raise here is whether those efforts at control might be more effective if they were informed by a theoretical infrastructure in which elements of control theory and engineering would have a place.

The author of this essay is neither a control theorist nor an economist but a student and observer of both fields. Like many others, I have taken note of close parallels between the basic principles of control theory and those of Keynesian macroeconomics (Hayes 1989, 2009). And I have wondered why control methods have not proved more useful at the practical level in setting macroeconomic policies. What follows should be regarded as the record of an ongoing personal inquiry into this question rather than an attempt to formulate a definitive answer.

2 Feedback in Machines and Markets

Control theory and economics have their roots in the same time and place: the early stirrings of the Industrial Revolution in 18th-century Britain.

The centrifugal flyball governor for steam engines, invented by James Watt in 1788, is the iconic control device, often serving as an emblem for the entire field of control engineering. Two heavy balls are attached by a hinge-like linkage to a shaft that rotates with the main power shaft of the engine; another lever arm connects the spinning balls to a steam valve. When the engine speeds up, the balls spin faster and hence fly outward; this change in configuration is communicated to the steam throttle valve, closing it slightly and thereby reducing the engine speed. Likewise, if the engine slows, the governor opens the valve, and the higher steam flow restores the machine to its set-point speed. Through this "feedback" mechanism the governor maintains a constant operating speed even as the load on the engine varies.

Watt's governor was by no means the first application of the feedback principle to a mechanical device; there were precedents in antiquity, and Watt himself was probably inspired by innovations in the design of windmills. But Watt's adoption of the idea brought wide attention to the idea of feedback control and led to a flowering of other applications (Mayr 1970).

Feedback also had an essential role in the greatest landmark of the classical economics literature, Adam Smith's *Inquiry into the Nature and Causes of the Wealth of Nations*, published in 1776. According to Otto Mayr, at least three key mechanisms in Smith's account of economic affairs can be described in terms of feedback loops (Mayr 1971).

First, Smith introduced a feedback loop regulating the overall size of the laboring population:

> The demand for men, like that for any other commodity, necessarily regulates the production of men; quickens it when it goes on too slowly, and stops it when it advances too fast. (Smith 1776, p. 80)

(The mechanism by which the working-class population is to be regulated should not be passed over in silence: Smith posited a linkage between wage levels and the rate of infant mortality.)

Second, Smith invoked a more elaborate feedback process for allocating labor resources to various trades:

> If in the same neighbourhood, there was any employment evidently either more or less advantageous than the rest, so many people would crowd into it in the one case, and so many would desert it in the other, that its advantages would soon return to the level of other employments. (Smith 1776, p. 99)

And finally there is Smith's grand conception of the market mechanism—the law of supply and demand—which can also be cast in the form of a feedback loop:

> The market price of every particular commodity is regulated by the proportion between the quantity which is actually brought to market, and the demand of those who are willing to pay the natural price of the commodity. (Smith 1776, p. 56)

In the subsequent paragraphs Smith makes clear that price is both a dependent variable and an independent variable in this relation. That is to say, price is determined by the ratio of supply to demand, but at the same time both supply and demand are dependent on price. Circularities of this kind are unavoidable in closed feedback systems.

Watt and Smith were contemporaries and Scottish compatriots, and at least one source suggests they were also friends, though of different social stations. During the period when Smith was professor of moral philosophy at the University of Glasgow, Watt opened an instrument-making shop on the university grounds. Smith's

19th-century biographer, John Rae, reports that: “Watt’s workshop was a favourite resort of Smith’s during his residence at Glasgow College, for Watt’s conversation, young though he was, was fresh and original, and had great attractions for the stronger spirits about him. Watt on his side retained always the deepest respect for Smith.... ” (Rae 1895, p. 74). Whether feedback systems ever came up as a topic of conversation remains a fascinating but probably unanswerable question.

3 The Keynesian Program

The feedback loops in Smith’s *Wealth of Nations* differ in a subtle but important way from those of Watt’s apparatus for controlling steam engines. The Watt governor compares the actual speed of the engine with a desired or commanded speed, the “set point,” and makes corrections to the actual speed to bring it into closer accord with the set point. But there is nothing resembling a set point in Smith’s classical vision of market economics. A feedback process guides buyers and sellers toward an equilibrium price, but there is no knob or dial for setting a desired target price. Nor is there any way to control the aggregate production of sellers or the aggregate demand of buyers.

One hundred sixty years later John Maynard Keynes argued for the addition of just such controlling knobs and dials to the macroeconomic machinery. The aim was not to intervene in the market process that sets prices or production priorities; those details were to be left to Smith’s feedback loops. But Keynes held that merely ensuring an equilibrium between production and consumption did not in fact guarantee economic well-being. A depressed economy, with a low rate of production balancing an equally low rate of consumption, is clearly in equilibrium. But almost everyone would prefer a *different* equilibrium, with both production and consumption closer to their maximum sustainable levels.

These ideas were put forward in the context of the Great Depression of the 1930s. Keynes’s agenda was to provide tools for mitigating such cyclical fluctuations in business activity. Under the Keynes prescription, when an economy is growing at an unsustainable pace, leading to excessive inflation, the central bank raises interest rates and thereby restricts the money supply. At the same time, governments raise taxes or reduce spending, which also cools the economy. Conversely, when business slumps, the aim is to spur growth by lowering interest rates and by letting the government run a deficit, spending more than it takes in through taxes.

Keynes’s ideas have gone in and out of fashion, and they remain a subject for continual debate and reinterpretation among dozens of disputatious factions within the world of professional and academic economics. Out in the real-world economy, however, that debate is large beside the point. Leaders of governments and central banks find it a political necessity to respond actively to any severe or prolonged unfavorable economic circumstances. The tools available for such interventions tend to have a Keynesian flavor, whatever the avowed ideology of the party putting them to work. As Milton Friedman and Richard Nixon are said to have said, “We’re all Keynesians now” (Reich 2010, p. 44). In the U.S., the 2008 economic crisis arrived at a moment of transition between two regimes grounded in quite different political and economic philosophies. But the outgoing and the incoming administrations adopted very similar strategies (both monetary and fiscal) for dealing with the recession. Both expressed confidence that stimulus plans would reinvigorate a swooning economy.

Reliance on these policies would seem to imply an underlying faith in the basic proposition that an economy is subject to control, at least in principle, and that we have adequate means to maintain control over it. These are just the kinds of assumptions that the engineering discipline of control theory is equipped to test.

4 Everything is Under Control

On first acquaintance, the idea of feedback control seems straightforward enough (Åström and Murray 2008). Consider the design of a cruise-control system for an automobile, designed to keep the car moving at constant speed.

A zeroth-order version of such a device simply clamps the throttle valve (or the gas pedal) at a fixed position. This is an "open loop" controller, and it doesn't perform very well. The car slows as soon as it begins climbing an upgrade, and speeds up on every downhill.

The simplest closed-loop controller measures the instantaneous speed of the car, compares it with the desired speed, then adjusts the throttle by an amount proportional to the difference (taking proper care that the sign of the correction acts to reduce the error). If the car slows somewhat on an upgrade, the controller senses the discrepancy and opens the throttle wider, so that the car regains some of the lost speed.

A drawback of pure proportional control is that the car never quite attains the requested speed; as the error diminishes, so does the feedback signal, and the system settles into a state with some nonzero offset from the correct velocity. The offset can be eliminated by another form of feedback, based not on the error itself but on the integral of the error with respect to time. In effect, the integral measures the cumulative error, which keeps growing if the speed differs even slightly from the set point. Thus integral control ensures that over the long term the net error approaches zero and the average speed converges on the set-point speed.

Yet integral control has drawbacks of its own. Suppose the car lacks the power to maintain a commanded speed on an upgrade, even with the throttle wide open; an integral controller would compensate by going faster than the commanded speed on the next downhill, which is probably not desirable behavior. More generally, integral control has a tendency to overshoot and oscillate around the set point. A remedy is to add still another form of feedback, based on the time derivative of the error signal. Derivative feedback opposes rapid changes in speed and thus tends to damp out oscillations.

Proportional, integral and derivative control (together known as PID) are the basic tools of "classical" control theory. The mathematical apparatus supporting these methods was developed in the 1920s and 30s, most notably at Bell Telephone Laboratories, where the objective was to build low-distortion amplifiers for long-distance speech transmission (Bode 1960). The Bell workers (including Harry Nyquist, Harry S. Black and Hendrick W. Bode) developed frequency-domain methods for designing feedback controllers and for evaluating their performance and stability.

Stability is a vital issue. In many instances the system to be brought under control—usually called the "plant," whether it is an oil refinery, an airplane or a stock market—is inherently unstable when observed in isolation. Most aircraft, for example, have a pitch-axis instability known as porpoising or phugoid oscillation, in which the airplane alternately dives and ascends, trading airspeed for altitude, kinetic for potential energy. A PID controller can damp out those oscillations. But it's also possible for a controller to introduce instabilities to a plant that would be stable on its own.

The hazard of controller-induced instability is most acute when there are delays built into the feedback circuit. The nature of this problem is familiar in everyday life. You step into the shower and find that the water is too cool, so you twist the temperature-control valve counterclockwise. Nothing happens for a few seconds, and so you turn the valve a little more. When the hot water finally makes its way to the shower head, you find you've gone too far. You dial the valve back a little, but the water continues to get hotter, so you turn the control further clockwise. Soon you're shivering. The temperature oscillations can keep growing in amplitude until the shower is alternately emitting the hottest and the coldest water available.

An analysis of this behavior based on the frequency-domain methods of Nyquist and Bode would show that the controller has too much gain at high frequencies. The built-in lag in the response of the plant puts constraints on the performance of the controller. Unfortunately, similar lags are likely to be present in economic systems as well.

5 Modern Control Theory

Labeling something "classical"—whether in music, physics, economics or control theory—carries the suggestion that some newer and spiffier version has superseded it. In the case of control theory there is indeed a "modern" variant, developed mainly since the 1960s (Doyle et al. 1992).

Modern control theory favors time-domain analysis over the frequency-domain design methods of classical theory. In other words, the aim is to directly trace the evolution of the state of a dynamical system as a function of time, rather than studying the response of the system to cyclical signals of various frequencies. Time-domain methods can more readily handle multiple inputs and multiple outputs. For example, with the classical approach it's easy to create separate controllers for an aircraft's roll, pitch and yaw axes, but harder to account for coupling between the axes; modern methods yield a unified control law for all three axes.

There have been other innovations. Classical control theory is framed in terms of differential equations with continuous variables and continuous time; the modern theory can handle difference equations with variables that take discrete values and evolve in discrete time steps. Tools of computational optimization have been introduced to search for the best control law within some universe of candidates. And there are methods for coping with various kinds of uncertainty. Stochastic control tolerates noise or errors in the measurements of the system's state. Robust control finds laws that deliver reasonable performance even if the real system differs somewhat from the mathematical model that represents it. Adaptive control applies the feedback principle to the control laws themselves, allowing the controller to continue working as the system evolves.

The new techniques have transformed the practice of control engineering, but what has changed most is the process an engineer follows in designing a control system; at a deeper level, the operation of the controller itself is much the same.

"Under the hood," it's still about feedback loops—comparing the observed state of the system (present, past and predicted) with the desired state, and calculating appropriate corrections.

The entire apparatus of modern control theory stands at the ready for any economist who might want to apply it to problems in macroeconomics. In the brief flowering of such experiments in the 1960s and 70s there was particular interest in optimal control, and attention later turned to stochastic control. Yet the most pressing questions about control-theoretic methods in macroeconomics are at such a basic level that the choice between classical and modern design methodologies is almost beside the point. When one tries to translate the problem of stabilizing global business cycles into the language of control engineering, three fundamental questions come up immediately. First, what are the process variables, and do we have sensors that can measure them reliably? Second, what are the control variables, and do we have actuators that can manipulate them effectively? And, finally, do the dynamics of the system—most notably the delay between input and response—admit a stable control solution?

6 Economic Sensors and Actuators

In designing a control system for an economy, the first step is to identify economic equivalents of engineering concepts such as sensors and actuators. The sensors are meant to measure levels of economic activity, such as employment and savings rates, income, spending, imports, exports and dozens of other variables. The actuators in a Keynesian program are monetary and fiscal policies. Through monetary policy a central bank controls the supply of money (mainly by regulating the terms of credit). Fiscal policy is usually summarized as the balance between government revenues and expenditures, although in fact tax and spending policies often include many kinds of incentives and disincentives for particular kinds of economic activity.

On the sensor side, the number that most often serves as a proxy for overall economic health is gross domestic product (GDP). This measure (or its predecessor GNP) was in fact developed in the context of the original Keynesian program in the 1930s, mainly by Richard Stone in the U.K. and by Simon Kuznets in the U.S. As an indicator of general economic welfare it is subject to many criticisms— failure to capture certain sectors of economic activity such as unpaid work and indifference to wealth distribution—but as a candidate for a process variable in a control system the main failing of GDP is timeliness. In the U.S., the framework for GDP calculation is built on an economic census conducted at intervals of five years, supplemented by various monthly, quarterly and annual data sources (Landefeld et al. 2008). Final quarterly estimates of GDP are released three months after a quarter ends (and the figures are still subject to further revision, up to two years later). Thus if we choose to adopt GDP as the process variable for a control algorithm, we must cope with a time lag of at least three to six months.

With the spread of electronic commerce and communication networks, vast quantities of economic data are becoming available in close to real time. Banking and retail credit transactions are processed in seconds; stock-market trades can be tracked with millisecond precision. Unfortunately, non-monetary flows and stocks—such as inventory levels—cannot be gauged as quickly. Still, with enough effort perhaps it will be possible to devise some trustworthy measure of overall economic wellbeing that can be computed with a delay of days or weeks rather than months or years.

On the actuator side of the control process, monetary policy is typically revised every month or two, and there is no fundamental reason it could not be updated more frequently if the need arose. Fiscal policy is another matter. Tax codes and many kinds of government spending are captives of an annual budget process, which is intensely politicized. After an economic upset, reaching agreement on the need for corrective action can take months, with more months to negotiate the nature and magnitude of the response, and still more months before the effects of any legislation filter through to the public. Furthermore, the entire process is driven more by political imperatives than by a scientific assessment of economic needs.

The one branch of fiscal policy that responds promptly to economic events consists of entitlement payments such as unemployment compensation, which are triggered automatically, without the need for legislative intervention. Such "social safety net" mechanisms presumably do help to damp business-cycle oscillations, but from a control-theory perspective their inflexible, automatic character eliminates them as tools for imposing active control. From this point of view they are really part of the plant—the system to be controlled—rather than part of the controller.

7 Running Hot and Cold

Given the delays in both sensing and responding to economic events, the characteristic time scale for a macroeconomic feedback control loop would appear to be at least a year, and possibly longer. (This is the time required for a signal to make a full circuit of the loop: from the moment the controller first detects a departure from desired behavior, then calculates and applies a corrective action, until it finally detects the first effects of the correction.) A feedback system with such extended delays gives rise to worries about controller-induced instability—about the risk of being alternately scalded and frozen in the shower.

Phillips explored this issue with great care in his 1957 paper on "stablisation policy and the time-forms of lagged responses." This work was based on simulations run on electronic and electromechanical analog computers at the National Physical Laboratory and at the aviation firm Short Brothers and Harland. (The machines were conceptually similar to the hydraulic MONIAC, but faster, more versatile and more accurate.) In his analysis, Phillips distinguished between "lags," in which a control action begins immediately but takes effect gradually over an interval of time, and "delays," in which the action does not even begin to make itself felt until the end of the interval. His strategy was to set up a model with specified lags and delays, then search for controller parameters—gains for proportional, integral and derivative feedback loops—that would most effectively stabilize the economy. He found that both lags and delays could make the simulated economy harder to control, but delays had particularly disastrous consequences. In some cases, with delays of six months, no combination of PID coefficients could suppress the growth of undamped oscillations. Worse still, the unwanted oscillations were introduced by the action of the controller itself; they appeared even when the underlying economic model had no tendency to generate business cycles.

The result of Phillips's experiment could be taken as a sobering lesson and warning, given that the lags and delays involved in recent efforts to stabilize business cycles lie deep in the danger zone that Phillips identified, out beyond the six month threshold. At this moment, banks and governments might be trying to stimulate a sluggish economy when it is already on the verge of overheating. Then in the next phase of the cycle, they could react too late again and douse the smoldering ruins of an economy that by then would need renewed kindling.

In early 2009, thoughts along these lines led me to revive a longstanding personal interest in the connections between control theory and economics, and then to write the essay that appeared as (Hayes 2009). Looking back from two years later, the hazard of controller-induced oscillations has not gone away; on the other hand, there is no convincing evidence of such overshooting or instability. If anything, the economic predicament of many nations suggests an *excess* of stability: The stimulus programs enacted in 2008 and 2009 amounted to the largest deliberate economic intervention in history, and yet they have had little impact on unemployment or other measures of malaise.

This stubborn resistance suggests still another question: Could it be that governments and central banks lack the capacity to tame the business cycle? No cruise-control algorithm can keep a car up to speed if the engine cannot produce the torque needed to climb a hill. Likewise, agencies trying to correct a severe economic downturn may lack the resources to restore prosperity. It is a measure of our profound ignorance about the true present state of the economy that two diametrically opposite explanations—over- and under-correction—can plausibly be entertained at the same time.

8 Second Guessing

When economic applications of control theory fell from fashion in the later 1970s and 1980s, it was not because of any dramatic controller-induced instabilities like those that worried Phillips in 1957. In fact most of the attempts to apply optimal control theory in economics never got as far as examining closed-loop systems; they merely used optimization procedures to shape policy trajectories in an open-loop context. Under those circumstances, runaway out-of-phase feedback is not a concern.

One major cause of disillusionment was a critique published in 1976 by Robert E. Lucas, Jr., of the University of Chicago (Lucas 1976). Lucas argued that using any algorithm or mathematical model as a tool for setting economic policy was futile. Models, he said, predict the effect of policy changes without acknowledging that rational agents will alter their behavior under the new policies in ways that invalidate the assumptions on which the models were built. For example, if everyone knows (or can infer) the rule by which the Federal Reserve sets interest rates, borrowers and lenders will anticipate any changes in rates, adjusting their behavior in ways that tend to neutralize the effect of the policy change.

David Kendrick (2005) has described the effect of the Lucas critique on control theory in economics: "[W]ork on control theory models in general and stochastic control models in particular went into rapid decline and remained that way for a substantial time." This outcome is curious, because the Lucas critique doesn't really address the mechanism of feedback control, much less prove its futility. The critique describes only half of a feedback loop: The plant (represented by the agent with rational expectations) gets to react to the strategy of the controller, but the loop is never closed so that the controller can react in turn to the moves made by the plant.

The control strategy is assumed to be static, or at least always a time step behind the plant. A true closed-loop system has symmetrical dynamics, where the state of the plant and the controller continually co-evolve.

If the Lucas critique is to be taken seriously, it attacks not just control theory and mathematical modeling but *any* reasoned strategy for managing an economy—or a corporation, or a nation. Given the hypothesis that any predictable policy will always be undermined by a determined adversary, the only effective rulers would be those who are utterly capricious. Empirical evidence suggests otherwise.

Evidence also suggests that real-world economic agents are not as strongly motivated to "game the system" as Lucas supposed. Ray C. Fair of Yale University, using a model of the U.S. economy based on decades of empirical data, tested variations of the model in which agents could look ahead and base their behavior on predictions of future regulatory policies. The results suggest that such activity is not common in the real economy (Fair 2004). Another series of studies by Glenn D. Rudebusch of the Federal Reserve Bank of San Francisco reached a similar conclusion (Rudebusch 2002).

9 Taking Control

With or without the aid of control theory, governments will continue to make their best efforts to ameliorate the suffering caused by business cycles. Doing so is a humanitarian as well as a political necessity. One might imagine a leader who would stand aside and say: "I have no idea what's the matter with the economy, and since I don't know how to fix it, I'm going to keep my mitts off it." That would be refreshingly honest, but inexcusable all the same.

Informal feedback principles have probably helped to frame some of the countercyclical policies of recent years. In 1993 John B. Taylor of Stanford University proposed mathematical rules for setting central-bank interest rates in response to inflation and economic growth—rules that have almost surely influenced policy decisions over the past decade (Taylor 1993). The Taylor rules constitute a rudimentary feedback mechanism. On the other hand, they do not exploit the full power of control theory. Taylor's formulas ignore the question of controller-induced instability, and they take no account of the uncertainties that enter into stochastic and robust control methods.

Needless to say, the cost of miscalculation is high. Ben S. Bernanke, a student of the Great Depression, has argued that timid and misguided policies of the Federal Reserve were partly to blame for the length and severity of the 1930s depression in the U.S. (Bernanke 2004). (As chairman of the Fed, Bernanke has presided over a monetary response that is anything but timid.) On the other side the fence, Athanasios Orphanides, also of the Federal Reserve, has argued that overaggressive corrections in the 1970s contributed to the "stagflation" of that decade (Orphanides and Williams 2005). And still another faction questions whether monetary adjustments really have much effect at all. Through an ingenious computer simulation, Christopher A. Sims of Yale imposed the policies of the modern Fed on the economy of the 1930s, and vice versa. Swapping the strategies had little effect on the outcome (Sims 1998).

In a matter of such grave consequence it's unnerving to find so little consensus on basic principles. If the designers of an airplane disagreed so vehemently about the engineering of the flight-control system, the airplane would not leave the ground until the conflict had been resolved. But the design methodology and the engineering culture for aviation systems is quite different from that of economics. In control theory there is an unshakeable reliance not just on science and mathematics but at an even deeper level on cause and effect. The airplane controller may have to cope with model uncertainties and noisy measurements, and yet at the root of the system is a deterministic kernel for which the same input always yields the same output. Some level of predictability and repeatability in patterns of stimulus and response is a prerequisite for understanding and controlling one's environment. Have we yet reached that level in economic policy?

At the very beginning of the Great Depression, Keynes wrote essay that struggles bravely but vainly to maintain a chin-up message (Keynes 1930):

> This is a nightmare, which will pass away with the morning. For the resources of nature and men's devices are just as fertile and productive as they were. The rate of our progress towards solving the material problems of life is not less rapid. We are as capable as before of affording for everyone a high standard of life.... But to-day we have involved ourselves in a colossal muddle, having blundered in the control of a delicate machine, the working of which we do not understand. The result is that our possibilities of wealth may run to waste for a time—perhaps for a long time.

The best I can offer by way of a cheerful response is to suggest that it's not too late to come to an understanding of that delicate machine.

10 Acknowledgments

For tutoring on the rudiments of control theory, and for a glimpse into the culture of control engineering, I am indebted to Dale Enns, Frank J. Przybylski and Gunter Stein, who took time to talk with me in the 1980s, when they were all with Honeywell, Inc. More recently, David A. Kendrick of the University of Texas at Austin has provided valuable guidance and insight. At the Trento conference celebrating the Phillips Machine, in December 2010, David Colander and Robert Leeson offered particularly acute comments, criticisms and references. Finally, Kumaraswamy Velupillai and Stefano Zambelli and their colleagues and students in the Algorithmic Social Sciences Research Unit at the University of Trento have provided perfectly calibrated measures of enthusiastic support and uncompromising intellectual engagement.

References

[Åström and Murray 2008] Åström, Karl Johan, and Richard M. Murray. 2008. *Feedback Systems: An Introduction for Scientists and Engineers.* Princeton: Princeton University Press.

[Athans and Kendrick 1974] Athans, Michael, and David Kendrick. 1974. Control theory and economics: A survey, forecast, and speculations. *IEEE Transactions on Automatic Control* 19:518–524.

[Barr 2000] Barr, Nicholas. 2000. The history of the Phillips machine. In Leeson 2000, pp. 89–114.

[Bernanke 2004] Bernanke, Ben S. 2004. *Essays on the Great Depression.* Princeton: Princeton University Press.

[Bissell 2007] Bissell, Chris. 2007. Historical perspectives: The Moniac: A hydromechanical analog computer of the 1950s. *IEEE Control Systems Magazine* 27(1):69–74.

[Bode 1960] Bode, H. W. 1960. Feedback: The history of an idea. In *Symposium on Active Networks and Feedback Systems*, Polytechnic Institute of Brooklyn, 1960. Reprinted in *Selected Papers on Mathematical Trends in Control Theory* (R. E. Bellman and R. Kalaba, eds.), New York: Dover, 1964.

[Doyle et al. 1992] Doyle, John C., Bruce A. Francis and Allen R. Tannenbaum. 1992. *Feedback Control Theory.* New York: Macmillan Publishing Co.

[Fair 2004] Fair, Ray C. 2004. *Estimating How the Macroeconomy Works.* Cambridge, Mass.: Harvard University Press. See also `http://fairmodel.econ.yale.edu/main2.htm`

[Fair 2010] Fair, Ray C. 2010. Estimated macroeconomic effects of the U.S. stimulus bill. *Contemporary Economic Policy* 28:439452.

[Hayes 1989] Hayes, Brian. 1989. Everything is under control. *Scientific Honeyweller* 10(1):4–5. See `http://bit-player.org/bph-publications/SciHon-1989-Hayes-control.pdf`

[Hayes 2009] Hayes, Brian. 2009. Everything is under control. *American Scientist* 97(3):186–191. See `http://amsciadmin.eresources.com/libraries/documents/2009491133257238-2009-05Hayes.pdf`

[Kendrick 1988] Kendrick, David A. 1988. *Feedback: A New Framework for Macroeconomic Policy.* Boston: Kluwer Academic Publishers.

[Kendrick 2005] Kendrick, David Andrew. 2005. Stochastic control for economic models: Past, present and the paths ahead. *Journal of Economic Dynamics and Control* 29:3–30.

[Keynes 1930] Keynes, John Maynard. 1930. The great slump of 1930. *The Nation & Athenæum*, December 20 and December 27, 1930.

[Landefeld et al. 2008] Landefeld, J. Steven, Eugene P. Seskin and Barbara M. Fraumeni. 2008. Taking the pulse of the economy: Measuring GDP. *Journal of Economic Perspectives* 22(2):193–216.

[Leeson 2000] Leeson, Robert (editor). 2000. *A. W. H. Phillips: Collected Works in Contemporary Perspective.* Cambridge: Cambridge University Press.

[Lucas 1976] Lucas, Robert E., Jr. 1976. Econometric policy evaluation: A critique. Carnegie-Rochester Conference Series on Public Policy, Vol. 1, pp. 19–46.

[Mayr 1970] Mayr, Otto. 1970. *The Origins of Feedback Control.* Cambridge, Mass.: M.I.T. Press.

[Mayr 1971] Mayr, Otto. 1971. Adam Smith and the concept of the feedback system: Economic thought and technology in 18th-century Britain. *Technology and Culture* 12(1):1–22.

[Orphanides and Williams 2005] Orphanides, Athanasios, and John C. Williams. 2005. The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations. *Journal of Economic Dynamics and Control* 29:1927–1950.

[Phillips 1954] Phillips, A. W. 1954. Stabilisation policy in a closed economy. *The Economic Journal* 64(254):290–323. Reprinted in Leeson 2000.

[Phillips 1957] Phillips, A. W. 1957. Stabilisation policy and the time-forms of lagged responses. *The Economic Journal* 67(266):265–277. Reprinted in Leeson 2000.

[Rae 1895] Rae, John. 1895. *Life of Adam Smith*. London: Macmillan and Co.

[Reich 2010] Reich, Robert B. 2010. *Aftershock: The Next Economy and Americas Future*. New York: Alfred A. Knopf.

[Rudebusch 2002] Rudebusch, Glenn D. 2002. Assessing the Lucas critique in monetary policy models. Federal Reserve Bank of San Francisco, Working Paper 2002–02. See `http://www.frbsf.org/publications/economics/papers/2002/wp02-02bk.pdf`

[Sims 1998] Sims, Christopher A. 1998. The role of interest rate policy in the generation and propagation of business cycles: What has changed since the '30s? In *Beyond Shocks: What Causes Business Cycles?*, edited by Jeffrey C. Fuhrer and Scott Schuh. Boston: Federal Reserve Bank of Boston, Conference Series No. 42.

[Smith 1776] Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. Edited by Edwin Cannon, based on the 1798 fifth edition. London: Methuen & Co., 1904.

[Taylor 1993] Taylor, John B. 1993. Discretion versus policy rules in practice. Carnegie-Rochester Conference Series on Public Policy, Vol. 39, pp. 195–214.

[Tustin 1953] Tustin, Arnold. 1953. *The Mechanism of Economic Systems: An Approach to the Problem of Economic Stabilization from the Point of View of Control-System Engineering*. Cambridge, Mass.: Harvard University Press.