Nobel Prize Lecture

QUANTITATIVE AGGREGATE THEORY

by

Finn E. Kydland

2004 Memorial Prize in Economic Sciences

QUANTITATIVE AGGREGATE THEORY

by

Finn E. Kydland

Professor, Jeff Henley Endowed Chair in Economics, University of California, Santa Barbara
and Research Associate, Federal Reserve Bank of Dallas

I'm delighted to stand before so many people. I'm also very happy when I get to work with *models* with many people. That is the key to the framework for which Ed Prescott and I were cited by the Nobel committee: The people are introduced explicitly in the models. Their decision problems are fully dynamic—people are forward-looking. That is one of the prerequisites for what we ultimately seek, a framework in which we can evaluate economic policy.

The eminent researcher and 1995 Nobel laureate in economics, Bob Lucas, from whom I've learned a lot, wrote (Lucas 1980): "One of the functions of theoretical economics is to provide fully articulated, artificial economic systems that can serve as laboratories in which policies that would be prohibitively expensive to experiment with in actual economies can be tested out at much lower cost….(696) Our task, as I see it…is to write a FORTRAN program that will accept specific economic policy rules as 'input' and will generate as 'output' statistics describing the operating characteristics of time series we care about, which are predicted to result from these policies." (709–10) The desired environments to which Lucas refers would make use of information on "*individual* responses [that] can be documented relatively cheaply…by means of…censuses, panels [and] other surveys…." (710) Lucas seems to suggest that economic researchers place people in desired model environments and record how they behave under alternative policy rules.

In practice, that is easier said than done. The key tool macroeconomists use is the *computational experiment.* Using it, the researcher performs exactly what I just described—places the model's people in the desired environment and records their behavior. But the purpose of the computational experiment is wider than simply to evaluate policy rules, as Lucas suggests. The computational experiment is useful for answering a host of questions, particularly quantitative questions, that is, those for which we seek numerical answers. When evaluating government policy, the policy is in the form of a rule that specifies how the government will behave—what action to take under vari-

ous contingencies—today and in the indefinite future. That's why it would be so difficult and prohibitively expensive to perform the alternative Lucas mentions, namely, to test the policies in actual economies.

THE COMPUTATIONAL EXPERIMENT

These models, as I already implied, contain millions of people. My tiny laptop contains several such models. People are characterized by their preferences over goods and leisure into the indefinite future. Their budget constraints are explicit. They receive income from working and from owning capital, and they must remain within their budget constraints, given the prices they face—wage rates and interest rates, for example. In other words, these models are explicit about people's dynamic decision problems.

The models also contain thousands of businesses. Implied, then, is a description of aggregate production pos-

sibilities—say, in the form of an aggregate production function. It describes the technology for converting inputs of capital and labor into output of goods, which can be used for consumption or to add to future productive capital—for investment.

A key aspect of the production function is its description of the technology level and its change over time. It's a broad concept at this level of abstraction. Technological change encompasses anything that affects the transformation, implied by the aggregate production function, of aggregate inputs of capital and labor into goods and services. It includes, of course, the usual outcomes of innovative activity, but also could include, again at this level of abstraction, factors such as oil shocks, new environmental regulations, changes in the legal constraints affecting the nature of contracting between workers and firms, government provision of infrastructure, and the loss in financial intermediation associated with banking panics—all elements one might want to study in more detail, depending on the question. But, for many questions, it makes perfect sense to include them implicitly as part of the technology level.

I've described two elements of typical models used for computational experiments: the millions of model inhabitants and the thousands of businesses. An essential aspect, however, is the calibration of the model environment. In a sense, models are measuring devices: they need to be calibrated, or otherwise we would have little faith in the answers they provide. In this sense, they are like thermometers. We know what a thermometer is supposed to register if we dip it into water with chunks of ice, or into a pot of boiling water. In the same sense, the model should give approximately correct answers to questions whose answers we already know. Usually, there are many such questions. In the context of business-cycle analysis, we know a lot about the long run of the economy, or we may use the Panel Study of Income Dynamics, say, or similar panel studies from other nations to collect the data to calibrate the model. Thus, the calibration is part of the action of making the quantitative answer as reliable as possible.

A computational experiment yields time series of the aggregate decisions of the model economy's people. Through the model formulation and its calibration, we have determined what the economic environment should look like. Then, the millions of people and the thousands of businesses in the economy make their decisions over time, and the computer records their decisions. We obtain time series as if we were confronted with an actual economy. These time series may be described statistically and compared with analogous statistics from the data for the nation under study. In a business-cycle study, these statistics may include standard deviations of detrended aggregates, describing the amplitudes of their business-cycle movements, as well as correlation coefficients describing their comovements.

A SIMPLE EXAMPLE

Now I would like to walk you through a simple model—substantially simpler than that in Kydland and Prescott (1982), for example. It contains a household and a business sector. To make it as straightforward as possible, I'll abstract from the government; there will be neither government nor a foreign sector in this model. I will have two main goals: to discuss what it means to say that the model contains a household and a business sector, and to give examples of what's involved in calibrating the parameters (see Cooley and Prescott 1995 for a detailed description of the practice of calibration, and Kydland 1995 for a more elaborate example in which all the details have been worked out).

First, we have a description of the typical household's preferences in the form of a utility function to be maximized:

$$E\sum_{t=0}^{\infty} \beta^t \frac{\left(C_t^{\alpha} L_t^{1-\alpha}\right)^{1-\sigma} - 1}{1 - \sigma}$$

Business cycles involve uncertainty about the future, so what one aims to do is maximize *expected* (denoted by E) utility as a function of consumption, which I call C, and leisure, L, over the indefinite future. It may seem a little farfetched to be summing the utility from today (period zero, let's say) to infinity. I'll return to that assumption. The parameter β is some number slightly less than 1 and can be calibrated from knowledge of the long-run real interest. It simply describes the degree of people's impatience. Additional parameters are α and σ, also parameters we hope to calibrate. The σ is what we may call a risk aversion parameter, about which finance people know a lot. I'll return to α in a minute.

The model formulation as you see it is the statement of a planner's problem whose solution can be shown to be the equilibrium of an economy inhabited by millions of people with preferences such as this utility function. There is a resource constraint,

$$C_t + I_t = Z_t K_t^{\theta} N_t^{1-\theta} = r_t K_t + w_t N_t,$$

which says that the sum of consumption and investment cannot exceed what the economy produces. The right-hand side of the first equality says that the economy produces output using capital—factories, machines, office buildings—along with the labor input of workers, and the technology level is denoted by Z. In other words, this is total output—gross domestic product—as given by the production function, the specification of which is essential to all of macroeconomics these days. Moreover, GDP has to equal gross domestic income, the sum of capital and labor income, which appears on the right-hand side of the second equality.

In addition to this resource constraint, we have a constraint on time, which here can be devoted either to leisure

or to labor input:

$$L_t + N_t = 1.$$

The right-hand side is 1; that is, without loss of generality I've chosen units so that if we add all the discretionary time—total time net of sleep and personal care—across all people, it equals 1.

Then we have two relations that describe key aspects of what makes an economy dynamic:

$$K_{t+1} = (1 - \delta)\, K_t + I_t$$

and

$$Z_{t+1} = \rho Z_t + \varepsilon_{t+1}.$$

The first, where K_t denotes the capital stock at the beginning of period t, describes how the capital stock at any time depends on past investment decisions, where δ is the depreciation rate. Finally, the technology level is all-important because it's what, in this simple model, gives rise to uncertainty. If, as turns out to be realistic, the parameter ρ is close to 1, the relation says that if there are new technological innovations, given by ε, then they are long-lasting. One usually imagines that this random variable ε is drawn from a normal probability distribution, whose variance can be estimated from the data.

As we have seen, this simple economy already has a number of parameters we need to calibrate. One reason for presenting this model is so I can give two typical examples of calibration, namely of the parameters α in the utility functions and θ in the production function. Suppose we went to a panel of thousands of people and calculated the average of how much time they devote to market activity. It turns out that figure pins down, via a steady-state first-order condition, the value of α that makes this average identical in the model economy. Similarly, with regard to the parameter θ in the production function, a property of the model is that if we look up National Income and Product Accounts data and find that out of total gross domestic income, on the average 34 percent is compensation for capital input and 66 percent represents labor income, then that calibrates the parameter θ with a high degree of accuracy.

I've used this model as a vehicle for talking about the two key sectors of the economy. The household sector contains lots of people characterized by the utility function—a description of the preferences over consumption and leisure into the indefinite future. The business sector is described by the technology for producing goods and services from capital and labor inputs. I have talked about the key features that make this model dynamic, and about a key source of uncertainty. One could include many other such features. Ed Prescott mentioned in his lecture the so-called time-to-build assumption, which would make the model more detailed, as in the 1982 paper to which the Nobel committee refers. That model also contains inventories, as well as both permanent and temporary shocks. What to include depends on the question the model is designed to address. The question for which this framework was first put to use by Ed Prescott and me can be stated as follows: If technology shocks were the only source of impulse, what portion of business-cycle fluctuations would still remain? The model produced a preliminary answer to that question: on the order of 2/3, and that answer has pretty much been confirmed to be somewhere around 70 percent. The model has provided measurement.

DOES BEING DIFFERENT MATTER?

Returning to the utility function, I assume in my prototype model above that preferences are given by some function that covers the entire future—goes to infinity. In other words, we have great power in setting up this economy: we can decide that people are immortal! That assumption turns out to be surprisingly innocuous for many questions. Of course it makes sense to check if it makes a difference and, as economists often conclude in many contexts, it depends. For many business-cycle questions, the answer is no. That's rather surprising because, if you think about mortal people and their life-cycle behavior, typically they earn relatively little labor income early in their lives, experience a substantial increase in income when they enter the middle stage, and then, for those who live long enough, enter a period in which they will have retired from market work. In other words, the labor-earnings profile is decidedly hump-shaped. But we also know that people prefer a consumption stream that's much more even over time. So there will be a period in which they spend more than their income, then spend less for two or three decades, and finally revert to spending more than their labor income toward the end of their lives. Moreover, the behavior in different ways typically is quite interesting at the beginning and end of one's working life.

Thus, it would seem that life-cycle behavior could matter substantially. For example, Víctor Ríos-Rull (1996) finds for a typical business-cycle question such as the one I mentioned above that when we employ an economy with mortal consumers in which realistic life-cycle behavior is included, as we aggregate across all of these people the time series in the computational experiments, we get the same answer as in the immortal-consumer economy. Of course, there are a lot of questions for which life-cycle behavior does make a difference. Among those are the economic impact on savings and interest rates of immigration, Social Security reform, and baby boomers' retirement, to mention a few.

To give you a sense of how different people are and emphasize the need for including them for addressing some questions, I'll show you some numbers. Figure 1 displays the average life-cycle profile of people's efficiency of

Figure 1

U.S. Life-Cycle Wage Profiles

Wage (normalized to 1 on average)
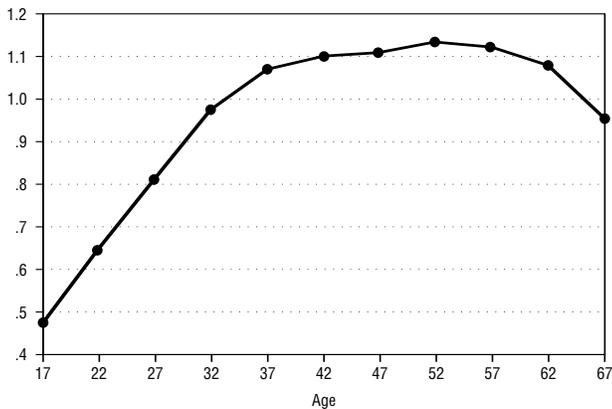
working in the market sector, as indicated by their real wage rates.

The graph shows a major reason for the hump-shaped profile of people's labor earnings depending on age. The curve is normalized so that it averages 1. It starts at around 0.5 and rises rapidly so that for a long time span later in people's working lives their efficiency is more than twice what it is when they enter the workforce. In addition to these life-cycle differences in workers' skills comes the fact that workers are of quite different abilities as they enter the

work force, depending on education and other factors. An interesting study of the aggregate implications of the interaction between, on the one hand, the labor input divided into low- and high-skilled workers and, on the other hand, the capital input divided into structures and equipment is in Krusell, Ohanian, Ríos-Rull, and Violante (2000). Their focus is on real-wage movements in particular. For a discussion of cyclical implications, especially as they pertain to measured labor-input fluctuations, see Kydland and Petersen (1997), on which some parts of this lecture are based.

Figure 2 displays the age distribution of the U.S. population in 1994 and also projected to 2020. The vertical axis shows the percentage of people of different ages. You see the noticeable hump in 1994 roughly in the 30-to-40 age range. Predictably, there will be a corresponding hump in 2020. Of course, a reason to worry about this empirical pattern is that in 2020 many, if not most, of these baby boomers will have retired, putting a major strain on the government budget constraint in general and the Social Security system in particular. A beautiful study of the effects the baby boomers in Spain (where immigration represents much less of a complication for the population dynamics than for the United States) may have on savings and real interest rates is in Ríos-Rull (2001).

Finally, Figure 3 tells us about the age distribution of immigrants to the United States. The curve for U.S. natives is the same as that for 1994 in Figure 2, except now each age group is five years wide and so the curve is smoother. The key message is that most immigrants to the United States are young.

These are all elements that one may wish to add to a

Figure 2

U.S. Age Distribution of Population, 1994 and 2020

Percentage of total population
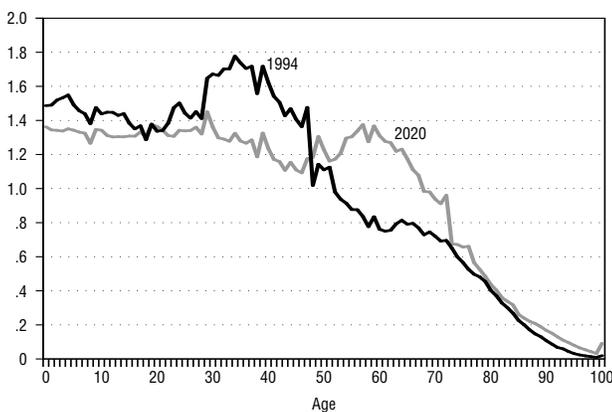
Figure 3

U.S. Age Distribution of Natives and New Immigrants

Percentage of total population
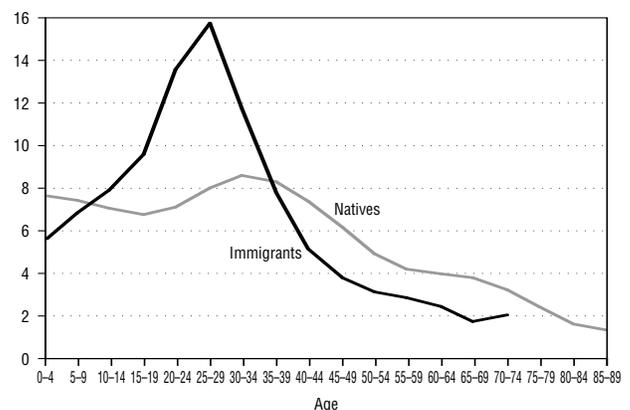
4

model of heterogeneous individuals—something we as economists have become adept at doing. When Víctor Ríos was my colleague at Carnegie Mellon University in the early 1990s, computers were not nearly as powerful as they are today. Víctor did some of the early pioneering research with such models. Some could take a long time—maybe a day or two—for the computer to calculate the model time series to analyze.

All of these features to which I've alluded—the age-dependent efficiency of working, population dynamics, and so on—can and have been added to models such as those used by Víctor Ríos and others in the past decade. A student of Víctor's and mine at Carnegie Mellon, Kjetil Storesletten, now at the University of Oslo, made an interesting study of the interaction of immigration with government fiscal policy. Some stark predictions have been made by people who do intergenerational accounting, suggesting that tax rates will have to rise in the not-so-distant future in order for the government budget constraint to be satisfied. The interesting question Storesletten (2000) asks is, To what extent can one avoid that tax increase by raising the rate of immigration, especially if one could be somewhat selective in the immigrants to admit?

Our ability to compute equilibriums for economies with very different people has expanded dramatically in recent years, with many studies heavily influenced by the pioneering paper by Per Krusell and Tony Smith (1998). Today, we see interesting research in which, for example, income and wealth distributions are allowed to vary and evolve over time, for example Storesletten, Telmer, and Yaron (2004). This exciting work is made possible through advances in our understanding of dynamic methodology, but also because of the power of today's computers.

NO MONEY?

A belief expressed by some is that this framework is used for analyzing real phenomena only. That's a huge misunderstanding. The same framework is used also to study monetary phenomena. For example, one could use it to ask the perennial question, Do monetary shocks cause business cycles?

[Before going on, I would like to say that there are two people whom I would have loved to see in Stockholm this week, but who will not be here because they have passed away. One is my father; the other is Scott Freeman, who died a few months ago. I've had the great fortune to work with the greatest economist in the world, Ed Prescott. But Scott Freeman was not far behind. He was a tremendous economist, with great insight and innovative ability. He and I did work on the interaction of monetary phenomena and real factors. In his memory, I've included two pictures. In the first, you see Scott in a pensive mood. In the second, he's enjoying himself at a party a couple of years ago.]

Here's a way to introduce money into a framework

such as the one I've described to you. Suppose people purchase a whole variety of sizes of goods. We might as well say there's a continuum of goods, from tiny to large. People make small purchases and large purchases. Because of the cost of carrying out transactions using means of exchange (checks, for example) backed by interest-earning assets, it has to be optimal to make the small purchases using currency and the large purchases using these other means of exchange. The extent to which you want to use either becomes an economic decision, whose incentives change over the cycle. They change for the choice of the proportion of the two means of exchange one wishes to hold, as well as for the frequency with which one replenishes one's liquid balances. The finding from this study with Scott Freeman (2000) is that money fluctuates procyclically even when the central bank does nothing. In other words, if one finds, as was the case over extended periods of U.S. history, that money moves up and down with output, that fact by itself says nothing about money causing output.

Because these models are inhabited by people, we can evaluate the welfare cost of inflation. In a project with Scott Freeman and Espen Henriksen (forthcoming), a Carnegie Mellon Ph.D. student, we did exactly that. We are now pushing that project further, asking, for example, what will happen if transaction costs drop over time, which already has happened and likely will continue to do so.
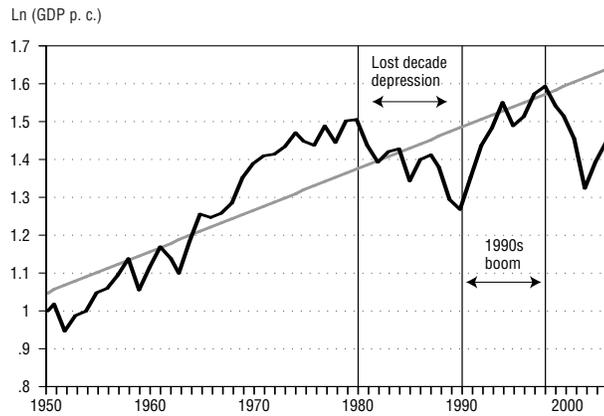
INTERNATIONAL BUSINESS CYCLES

I presented to you earlier a closed-economy model. In the past 10 or 15 years, however, economists have put this framework to use to study the interaction of many nations. This is a particularly interesting field because anomalies abound for bright young (and even old) researchers to try to account for. Here's an example that, on the face of it, may seem like an anomaly: For many nations, cyclically the trade balance is the worst when one's goods are cyclically the cheapest. It turns out that once you write down a model that allows for trade across nations, as, for example, Backus, Kehoe, and I did (1994), capital accumulation is important for the answer. Another factor is that there's "nonsynchronized" technological change in the different nations, which over time spills over from one nation to the next. The conclusion is that the empirical regularity to which I just referred is not an anomaly at all. It is exactly what the model suggests would happen.

Here's a cute application. I always loved to use it in my undergraduate course. I came across an article in the *Wall Street Journal* in April 1998 reporting that the International Monetary Fund dispatched representatives to Argentina, supposedly to convince the Argentine government to cool the economy. The reasons stated were threefold: (i) high growth rates, 6.5 to 7 percent annually, coming on top of strong growth that started in 1990, interrupted only by the Tequila crisis around 1995; (ii) export prices falling dramatically; and (iii) the trade deficit returning. Sound bad? As it turns out, these comovements are what a standard model would tell us to expect in an economy that's doing well. Our framework dictates that these three features, in combination, ought to be favorable. I should say that I have no way of knowing if the *Wall Street Journal* to some extent

misstated the IMF's basis for going to Argentina. It could be, for example, the IMF was worried about fiscal "over-stimulation," as one might call it.

THE CASE OF ARGENTINA

A number of studies of great depressions have been carried out recently. Many such studies were assembled for a conference at the Federal Reserve Bank of Minneapolis and will be collected in a volume edited by Tim Kehoe and Ed Prescott. The reasons I mention the great depression studies are twofold. The first is that people used to think great depressions are events of such magnitude that we need a separate framework to study them. I think this conference showed that any such suggestion is nonsense. The second reason is that this conference gave Carlos Zarazaga and me (2002) the excuse to study the case of Argentina, which had a great depression in the 1980s. To give you a sense of what has happened in Argentina in the last 50 years, Figure 4 displays the log of its real GDP per person of working age.

Logs are particularly useful because constant growth rate translates into a straight line, and whether Argentina is as small as it was in the 1950s, or half again larger in 1998, 1 cm deviation from trend, let's say, represents the same percent deviation from trend. So that's how to read this picture. You see the dramatic decline in the 1980s—over 20 percent—during Argentina's "Lost Decade," qualifying it as a great depression. An even larger and much faster decline took place after 1998.

As previously mentioned, Argentina's economy experienced an upturn in the 1990s. That episode, to Carlos Zarazaga and me (forthcoming), was even more interesting than the depression. Clearly, Argentina grew fast by most standards. The surprising thing was—and only the model could tell us this—when you put the numbers for total factor productivity growth into a standard model and calibrate it, the model says that investment should have been much larger in the 1990s. Of course, for that very reason, the capital stock should have been much larger, about 20 percent or so, by the end of the decade.

Figure 5 contains a picture of real GDP for Argentina, again in log scale. You can see the growth in the 1990s. Suppose we put the numbers into the model. We use the period up to 1980 to estimate statistically the process for the technology level. Then we enter into the model the actual numbers for total factor productivity measured by the same method that Robert Solow (1957) proposed for measuring them in a growth context. The model accounts well for the great depression of the 1980s, and it accounts well also for the downturn after 1999. The large discrepancy is for the 1990s where the model says that growth in the 1990s should have been much higher. The third curve is included to indicate what happens if we assume that the capital stock in 1999 is taken from the actual data for that

Figure 5

Argentina's GDP

Ln (GDP)

SOURCE: Kydland and Zarazaga (2004).

Figure 6

Argentina's Capital Input

Ln (K)

SOURCE: Kydland and Zarazaga (2004).
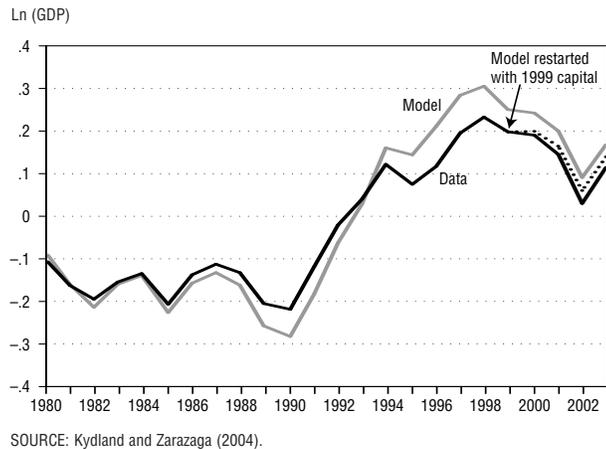
year and then we start the model up again in 1999. The model accounts well for the remaining years.

What if we look more closely at the capital input? I mentioned it as representing the key anomaly. That is borne out in Figure 6, which displays an even greater discrepancy between model prediction and data than in the case of GDP. The difference in 1999 is almost 20 percent. As in Figure 5, the third curve displays the model prediction if we start with the 1999 capital stock so as to account for the remaining five years.

For Argentina, the data in Figure 7 must be extremely depressing because they show the fall in capital stock per working-age person (which would look more or less the same in per-capita terms).

This represents the quantity of productive capacity in Argentina, given by the best measurements available. The capital stock in 2003, per capita, was much lower than in 1982. The neoclassical growth model then would imply, as the data show, much lower wage rates than those that would have prevailed in Argentina if the economy had grown the way other nations' economies did. This is bad news for the future of Argentina's poor (it certainly has been so far). Clearly Argentina would need to grow at a steady and much faster rate—not just 3 or 4 percent a year—to catch up. If it doesn't, the poor will surely stay poor for a long time. People with relatively high human capital are likely to do reasonably well, but the wealth and income disparities will keep getting wider.

What are possible explanations for the 1990s? The first thing to look for is any indication of measurement problems. In many nations like Argentina, the data are sometimes of poor quality. Moreover, aggregate series can be constructed from available data in different ways. A Ph.D.
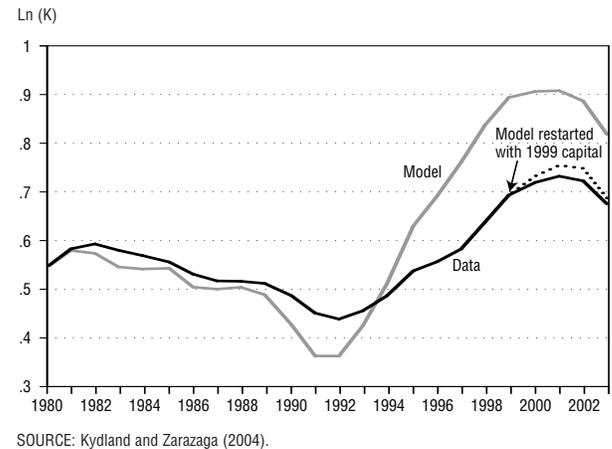
student at Carnegie Mellon, José de Anchorena (2004), tried an alternative way of constructing the capital series but reached the same conclusion.

One possibility, and I'd like to return to it because it relates to our 1977 paper that Ed Prescott talked about, is that the outcome for the 1990s in part is the result of what we may call the "time-inconsistency disease" due to bad policies in Argentina before 1990. People had their memories from the past, even if former President Carlos Menem and other politicians did their best to make Argentina a

Figure 7

Argentina's Capital Input Per Working Age Person

Lower Capital: Lower Real Wages, Worse Distribution of Income

Ln (K p. c.)

SOURCE: Kydland and Zarazaga (2004).

7
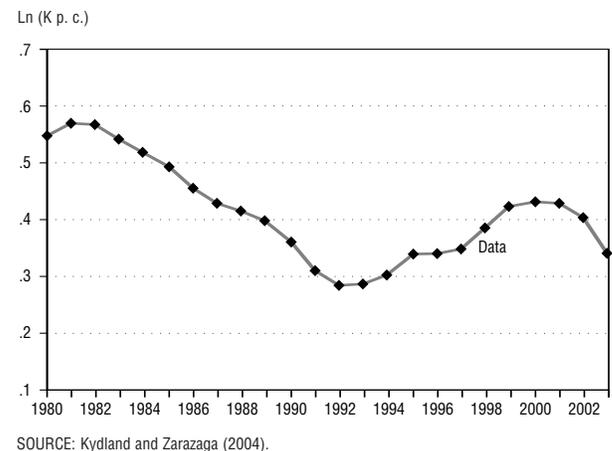
credible country in which to invest for the long run. Chances are Argentina still lacked the necessary credibility. There was growth but not nearly as much as Argentina should have experienced.

The Argentines have recovered in the past couple of years. I already mentioned that if it doesn't happen at a rapid speed, if the gap is not closed, the poor will stay that way for a long time. How will Argentina restore confidence? There's no easy answer. Once credibility has been lost, economists don't know much about how to restore it. What is needed is not a policy of patchwork for a year or two; Argentina needs a policy geared for the long run, with credible incentives for innovative activity and human and physical capital accumulation yielding returns far into the future.

CONCLUDING REMARKS

In this brief lecture, I've tried to give you a taste of the vast variety of questions, with the model details dictated accordingly, that have been addressed in macroeconomics in the past two decades, all within the framework that serves as the overall theme for this lecture: The decision problems of the models' people and businesses are explicit, and they are dynamic. I could have provided hundreds of references. Some of those that I chose to include are authored or co-authored by researchers with whom I've enjoyed tremendously to interact. I'm delighted to have them here in Stockholm as my guests.

As there are many students in the audience, I'd like to conclude with some remarks about learning macroeconomics. Almost all interesting phenomena in macroeconomics are dynamic; they are intertemporal. We need a theory of forward-looking people. Unfortunately, dynamic macro is difficult for beginners to learn; it's not easy to do dynamics on paper. Perhaps mainly for that reason, in the past 20 years the gap between research and textbooks has grown wider and wider. What to do?

There are some recent attempts to bridge the gap. I like many aspects of Steve Williamson's (2005) recent textbook, for example. It's amazing, however, that I've continued for so long to use (supplemented by my own notes) a textbook first published as early as 1974 by Merton Miller and Charles Upton (1986). It presents a dynamic framework with many of the features I have talked about, even life-cycle behavior. These two authors were simply great economists, and they included in their textbook the key elements they thought ought to be in basic dynamic models of macroeconomics.

One possible remedy for teaching macroeconomics is to use the computer for computational experiments (see Bjørnestad and Kydland 2004). This tool, which has been so influential in modern macroeconomic research, can also help the beginning and intermediate students learn dynamic macroeconomics. Students can compare model and real-economy cyclical statistics. The computer can generate plots of impulse responses. Shocks occur in every time period. It's hard in practice to disentangle the effect of each particular shock, as one occurs in every period, the shocks are not easy to observe and measure at the time they occur, and the effect of each is long-lasting. But model economies let us help the intuition. For example, with an impulse response, one pretends that there hasn't been a shock for a long time—that the economy is in its steady state. Then we give the model economy a single shock or impulse and record what happens over a number of time periods—a great aid to the intuition.

I would like to stop there and just say: Takk for at dere alle kom for å høre på meg. (Thank you all for coming to listen to me today.)

REFERENCES

Anchorena, José (2004): "Capital Accumulation, Sectoral Productivity and Real Exchange Rate." Carnegie Mellon University Working Paper.

Backus, David K., Patrick J. Kehoe, and Finn E. Kydland (1994): "Dynamics of the Trade Balance and the Terms of Trade: The J-Curve?" *American Economic Review*, 84(1): 84–103.

Bjørnestad, Solveig, and Finn E. Kydland (2004): "The Computational Experiment as an Educational Tool in Basic Macroeconomics." Working paper.

Cooley, Thomas F., and Edward C. Prescott (1995): "Economic Growth and Business Cycles." *Frontiers of Business Cycle Research*, T.F. Cooley (ed.), Princeton: Princeton University Press, 1–38.

Freeman, Scott, and Finn E. Kydland (2000): "Monetary Aggregates and Output." *American Economics Review*, 90(5): 1125–35.

Freeman, Scott, Espen Henriksen, and Finn E. Kydland (forthcoming): "The Welfare Cost of Inflation in the Presence of Inside Money," in *Low Inflation Economics*, Federal Reserve Bank of Cleveland.

Krusell, Per, and Anthony Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, 106(5): 867–96.

Krusell, Per, Lee Ohanian, José-Víctor Ríos-Rull, and Giovanni Violante (2000): "Capital-Skill Complementarity and Inequality." *Econometrica*, 68(5): 1029–53.

Kydland, Finn E. (1995): "Business Cycles and Aggregate Labor Market Fluctuations." *Frontiers of Business Cycle Research*, T.F. Cooley (ed.), Princeton: Princeton University Press, 126–56.

Kydland, Finn E., and D'Ann M. Petersen (1997): "Does Being Different Matter?" *Economic Review* (Third Quarter), Federal Reserve Bank of Dallas.

Kydland, Finn E., and Edward. C. Prescott (1977): "Rules Rather than Discretion: The Time Inconsistency of Optimal Plans," *Journal of Political Economy* 85(2) (June): 473–91.

Kydland, Finn E., and Edward. C. Prescott (1982): "Time to Build and Aggregate Fluctuations." *Econometrica*, 50: 1345–70.

Kydland, Finn E., and Carlos E. J. M. Zarazaga (2002): "Argentina's Lost Decade." *Review of Economic Dynamics*, 5(1): 152–65.

Kydland, Finn E., and Carlos E. J. M. Zarazaga (2004): "Argentina's Capital Gap Puzzle," Federal Reserve Bank of Dallas, Center for Latin American Economics Working Paper 0405, www.dallasfed.org/latin/papers/2004/lawp0405.pdf.

Kydland, Finn E., and Carlos E. J. M. Zarazaga (forthcoming): "Argentina's Lost Decade and the Subsequent Recovery Capital Gap Puzzle," in *Great Depressions of the Twentieth Century*, Federal Reserve Bank of Minneapolis, Tim J. Kehoe and Edward C. Prescott (ed.).

Lucas, Robert E., Jr. (1980): "Methods and Problems in Business Cycle Theory." *Journal of Money, Credit and Banking*, 12(4): 696−715.

Miller, Merton H., and Charles W. Upton (1986): *Macroeconomics, A Neoclassical Introduction*, Chicago: The University of Chicago Press.

Ríos-Rull, José-Víctor (1996): "Life-Cycle Economies and Aggregate Fluctuations," *Review of Economic Studies* 63 (July): 465−90.

Ríos-Rull, José-Víctor (2001): "Population Changes and Capital Accumulation: The Aging of the Baby Boom." *The B. E. Journals in Macroeconomics*, 1(1) (May).

Solow, Robert M. (1957): "Technical Change and the Aggregate Production Function." *Review of Economics and Statistics*, 39(3): 312−20.

Storesletten, Kjetil (1995): "On the Economics of Immigration," Ph.D. dissertation. Carnegie Mellon University.

Storesletten, Kjetil (2000): "Sustaining Fiscal Policy Through Immigration." *Journal of Political Economy*, 108(2): 300−23.

Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron (2004): "Consumption and Risk-Sharing Over the Life Cycle." *Journal of Monetary Economics*, 51(3): 609−33.

Williamson, Stephen D. (2005): *Macroeconomics*, 2nd ed., Boston: Pearson Addison Wesley.

Federal Reserve Bank of Dallas
www.dallasfed.org