By ROBERT E. LUCAS, JR.*

Macroeconomics was born as a distinct field in the 1940's, as a part of the intellectual response to the Great Depression. The term then referred to the body of knowledge and expertise that we hoped would prevent the recurrence of that economic disaster. My thesis in this lecture is that macroeconomics in this original sense has succeeded: Its central problem of depression prevention has been solved, for all practical purposes, and has in fact been solved for many decades. There remain important gains in welfare from better fiscal policies, but I argue that these are gains from providing people with better incentives to work and to save, not from better fine-tuning of spending flows. Taking U.S. performance over the past 50 years as a benchmark, the potential for welfare gains from better long-run, supply-side policies exceeds by far the potential from further improvements in short-run demand management.

My plan is to review the theory and evidence leading to this conclusion. Section I outlines the general logic of quantitative welfare analysis, in which policy comparisons are reduced to differences perceived and valued by individuals. It also provides a brief review of some examples—examples that will be familiar to many—of changes in long-run monetary and fiscal policies that consumers would view as equivalent to increases of 5–15 percent in their overall consumption levels.

Section II describes a thought-experiment in which a single consumer is magically relieved of all consumption variability about trend. How much average consumption would he be willing

* Department of Economics, University of Chicago, 1126 East 59th Street, Chicago, IL 60637. I am grateful for discussions with Fernando Alvarez, Gadi Barlevy, Lars Hansen, Per Krusell, Ellen McGrattan, Chris Phelan, Edward Prescott, Esteban Rossi-Hansberg, Tom Sargent, Matthew Shapiro, Tony Smith, Nancy Stokey, Kjetil Storesletten, and Tom Tallarini, and for the able assistance of Adrian Kats and Mikhail Golosov. to give up in return? About one-half of onetenth of a percent, I calculate. I will defend this estimate as giving the right order of magnitude of the potential gain to society from improved stabilization policies, but to do this, many questions need to be addressed.

How much of aggregate consumption variability should be viewed as pathological? How much can or should be removed by monetary and fiscal means? Section III reviews evidence bearing on these questions. Section IV considers attitudes toward risk: How much do people dislike consumption uncertainty? How much would they pay to have it reduced? We also know that business-cycle risk is not evenly distributed or easily diversified, so welfare cost estimates that ignore this fact may badly understate the costs of fluctuations. Section V reviews recently developed models that let us explore this possibility systematically. These are hard questions, and definitive answers are too much to ask for. But I argue in the end that, based on what we know now, it is unrealistic to hope for gains larger than a tenth of a percent from better countercyclical policies.

I. Welfare Analysis of Public Policies: Logic and Results

Suppose we want to compare the effects of two policies, A and B say, on a single consumer. Under policy A the consumer's welfare is $U(c_A)$, where c_A is the consumption level he enjoys under that policy, and under policy B it is $U(c_B)$. Suppose that he prefers c_B : $U(c_A) < U(c_B)$. Let $\lambda > 0$ solve

$$U((1 + \lambda)c_A) = U(c_B).$$

We call this number λ —in units of a percentage of all consumption goods—the *welfare gain* of a change in policy from A to B. To evaluate the effects of policy change on many different consumers, we can calculate welfare gains (perhaps losses, for some) for all of them, one at a time, and add the needed compensations to obtain the

⁺ Presidential Address delivered at the one-hundred fifteenth meeting of the American Economic Association, January 4, 2003, Washington, DC.

welfare gain for the group. We can also specify the compensation in terms of one or a subset of goods, rather than all of them: There is no single, right way to carry these comparisons out. However it is done, we obtain a method for evaluating policies that has comprehensible units and is built up from individual preferences.

There is a great tradition of quantitative public finance that applies this general framework using well-chosen Taylor expansions to calculate estimates of the compensation parameter λ , "welfare triangles" as Arnold C. Harberger called them. Today we use numerical simulation of general-equilibrium models, often dynamic and subject to unpredictable shocks, to carry out welfare analysis with the general logic that I have just sketched. Some examples will, I hope, convey the applicability of this approach and some of the estimates that have emerged.

Martin J. Bailey's (1956) thought-experiment of a perfectly predictable inflation at a constant rate, induced by sustained growth in the money supply, was a pioneering example of the quantitative evaluation of policy. In a replication of the Bailey study, I estimated the welfare gain from reducing the annual inflation rate from 10 to 0 percent to be a perpetual consumption flow of 1 percent of income.¹ Some economists take estimates like this to imply that inflation is a relatively modest problem, but 1 percent of income is a serious amount of money, and in any case, the gain depends on how much inflation there is. The gain from eliminating a 200percent annual inflation-well within the range of recent experience in several South American economies-is about 7 percent of income.

The development of growth theory, in which the evolution of an economy over time is traced to its sources in consumer preferences, technology, and government policies, opened the way for extending general-equilibrium policy analysis to a much wider class of dynamic settings. In the 1980's, a number of economists used versions of neoclassical growth theory to examine the effects of taxation on the *total* stock of capital, not just the composition of that stock.² The models used in these studies differ in their details, but all were variations on a one-good growth model in which consumers (either an infinitely lived dynasty or a succession of generations) maximize the utility of consumption and leisure over time, firms maximize profit, and markets are continuously cleared.

In general, these studies found that reducing capital income taxation from its current U.S. level to zero (using other taxes to support an unchanged rate of government spending) would increase the balanced-growth capital stock by 30 to 60 percent. With a capital share of around 0.3, these numbers imply an increase of consumption along a balanced growth path of 7.5 to 15 percent. Of course, reaching such a balanced path involves a period of high investment rates and low consumption. Taking these transition costs into account, overall welfare gains amount to perhaps 2 to 4 percent of annual consumption, in perpetuity.

Production per adult in France is about 70 percent of production per adult in the United States. Edward C. Prescott (2002) observes that hours worked per adult in France, measured as a fraction of available hours, are also about 70 percent of the comparable U.S. figure. Using estimates for France and the United States of the ratio $(1 + \tau_c)/(1 - \tau_h)$ that equals the marginal rate of substitution between consumption and leisure in the neoclassical growth model, he shows that tax differences can account for the entire difference in hours worked and, amplified by the indirect effect on capital accumulation, for the entire difference in production. The steady-state welfare gain to French households of adopting American tax rates on labor and consumption would be the equivalent of a consumption increase of about 20 percent. The conclusion is not simply that if the French were to work American hours, they could produce as much as Americans do. It is that the utility consequences of doing so would be equivalent to a 20-percent increase in consumption with no increase in work effort!

The gain from reducing French taxes to U.S. levels can in part be viewed as the gain from adopting a flat tax on incomes,³ but it is doubt-

¹ Lucas (2000). My estimates are based on the money demand estimates in Allan H. Meltzer (1963).

² For example, William A. Brock and Stephen J. Turnovsky (1981), Christophe P. Chamley (1981), Law-

rence H. Summers (1981), Alan J. Auerbach and Laurence J. Kotlikoff (1987), and Kenneth L. Judd (1987).

³ See also Robert E. Hall and Alvin Rabushka (1995).

ful that all of it can be obtained simply by rearranging the tax structure. It entails a reduction in government spending as well, which Prescott interprets as a reduction in the level of transfer payments, or in the government provision of goods that most people would buy anyway, financed by distorting taxes. Think of elementary schooling or day care. The gains from eliminating such fiscal "cross-hauling" (as Sherwin Rosen [1996] called the Swedish daycare system) involve more than eliminating "excess burden," but they may well be large.

The stakes in choosing the right monetary and fiscal policies are high. Sustained inflation, tax structures that penalize capital accumulation and work effort, and tax-financed government provision of private goods all have uncompensated costs amounting to sizeable fractions of income. We can see these costs in differences in economic performance across different countries and time periods. Even in the United States, which visibly benefits from the lowest excess burdens in the modern world, economic analysis has identified large potential gains from further improvements in long-run fiscal policy.

II. Gains from Stabilization: A Baseline Calculation

In the rest of the lecture, I want to apply the public finance framework just outlined to the assessment of gains from improved stabilization policy. Such an exercise presupposes a view of the workings of the economy in which short-run monetary and fiscal policies affect resource allocation in ways that are different from the supply side effects I have just been discussing.

One possibility is that instability in the quantity of money or its rate of growth, arising from government or private sources, induces inefficient real variability. If that were all there was to it, the ideal stabilization policy would be to fix the money growth rate. (Of course, such a policy would require the Federal Reserve to take an active role in preventing or offsetting instabilities in the private banking system.) But this cannot be all there is to it, because an economy in which monetary fluctuations induce real inefficiencies—indeed, any economy in which money has value—must be one that operates under missing markets and nominal rigidities that make changes in money into something other than mere units changes. Then it must also be the case that these same rigidities prevent the economy from responding efficiently to real shocks, raising the possibility that a monetary policy that reacts to real shocks in some way can improve efficiency.

If we had a theory that could let us sort these issues out, we could use it to work out the details of an ideal stabilization policy and to evaluate the effects on welfare of adopting it. This seems to me an entirely reasonable research goal—I have been thinking success is just around the corner for 30 years—but it has not yet been attained. In lieu of such a theory, I will try to get quantitative sense of the answer to the thought-experiment I have posed by studying a series of simpler thought-experiments.

In the rest of this section, I ask what the effect on welfare would be if *all* consumption variability could be eliminated.⁴ To this end, consider a single consumer, endowed with the stochastic consumption stream

(1)
$$c_t = A e^{\mu t} e^{-(1/2)\sigma^2} \varepsilon_t,$$

where $log(\varepsilon_t)$ is a normally distributed random variable with mean 0 and variance σ^2 . Under these assumptions

$$E(e^{-(1/2)\sigma^2}\varepsilon_t)=1$$

and mean consumption at t is $Ae^{\mu t}$. Preferences over such consumption paths are assumed to be

(2)
$$E\left|\sum_{t=0}^{\infty}\left(\frac{1}{1+\rho}\right)^{t}\frac{c_{t}^{1-\gamma}}{1-\gamma}\right|$$

٢

where ρ is a subjective discount rate, γ is the coefficient of risk aversion, and the expectation is taken with respect to the common distribution of the shocks ε_0 , ε_1 ,

Such a risk-averse consumer would obviously prefer a deterministic consumption path to a risky path with the same mean. We quantify this utility difference by multiplying the risky path by the constant factor $1 + \lambda$ in all dates and states, choosing λ so that the household is

⁴ This calculation replicates the one I carried out in Lucas (1987, Ch. III).

indifferent between the deterministic stream and the compensated, risky stream. That is, λ is chosen to solve

(3)
$$E\left\{\sum_{t=0}^{\infty}\beta^{t}\frac{\left((1+\lambda)c_{t}\right)^{1-\gamma}}{1-\gamma}\right\}$$
$$=\sum_{t=0}^{\infty}\beta^{t}\frac{\left(Ae^{\mu t}\right)^{1-\gamma}}{1-\gamma},$$

where c_t is given by (1). Canceling, taking logs, and collecting terms gives

(4)
$$\lambda \cong \frac{1}{2} \gamma \sigma^2.$$

This compensation parameter λ —the *welfare* gain from eliminating consumption risk—depends, naturally enough, on the amount of risk that is present, σ^2 , and the aversion people have for this risk, γ .

We can get an initial idea of the value to the economy as a whole of removing aggregate risk by viewing this agent as representative of U.S. consumers in general. In this case, to estimate λ we need estimates of the variance σ^2 of the log of consumption about its trend, and of the coefficient γ of risk aversion. Using annual U.S. data for the period 1947–2001, the standard deviation of the log of real, per capita consumption about a linear trend is 0.032.⁵ Estimates of the parameter γ in use in macroeconomics and public finance applications today range from 1 (log utility) to 4. Using log utility, for example, the formula (4) yields the welfare cost estimate

(5)
$$\lambda = \frac{1}{2}(0.032)^2 = 0.0005,$$

about one-twentieth of 1 percent of consumption.

Compared to the examples of welfare gains from fiscal and monetary policy changes that I cited above, this estimate seems trivially small: more than an order of magnitude smaller than the gain from ending a 10-percent inflation! Many questions have been raised about this estimate, and subsequent research on this issue has pursued many of them, taking the discussion deep into new scientific territory. In the next four sections, I will review some of the main findings.

III. Removeable Variance: Two Estimates

Even if we do not know exactly how much consumption risk would be removed by an optimal monetary and fiscal policy, it is clear that it would fall far short of the removal of all variability. The major empirical finding in macroeconomics over the past 25 years was the demonstration by Finn E. Kydland and Prescott (1982), replicated and refined by Gary D. Hansen (1985) and by many others since then, that technology shocks measured by the method of Robert M. Solow (1957) can induce a reasonably parameterized stochastic growth model to exhibit nearly the same variability in production and consumption as we see in postwar U.S. time series. In the basic growth model, equilibrium and optimal growth are equivalent, so that if technology shocks are all there is to postwar business cycles, resources are already being allocated efficiently and a variance-reducing monetary-fiscal policy would be welfare reducing. Even if the equilibrium is inefficient, due to distorting taxes, missing markets or the like, in the face of unavoidable technology and preference shocks an optimal monetary and fiscal policy will surely be associated with a positive level of consumption variance. We need to estimate the size of that part and remove it from the estimate of σ^2 used in (4).

Matthew D. Shapiro and Mark W. Watson's (1988) study is one of several relatively atheoretical attempts to break down the variance of production and other variables into a fraction due to what these authors call "demand" shocks (and which I will call "nominal" shocks) and fractions due to technology and other sources. Their study represents quarterly U.S. time series over the period 1951–1985 as distributed lags of serially independent shocks. The observables include first differences of a measure of hours worked, a log real GDP measure, and the corresponding implicit price deflator. To these three rates of change are added an *ex post* real interest rate (the three-month Treasury bill rate

⁵ The comparable figure using a Hodrick-Prescott trend with the smoothing parameter 400 is 0.022.

5

minus the inflation rate) and the change in the relative price of oil. The coefficients of an invertible vector autoregression are estimated, subject to several restrictions. This procedure yields time series of estimated shocks $\hat{\varepsilon}_t$ and decompositions of the variance of each of the five variables into the fractions "explained" by the most recent k values of each of the five shocks.

Shapiro and Watson apply a variety of theoretical principles to the interpretation of their estimates. They do not consistently follow the general-equilibrium practice of interpreting all shocks as shifts in preferences, technologies, or the behavior of policy variables, but they have in mind some kind of monetary growth model that does not have a long-run Phillips curve.⁶ Real variables, in the long run, are determined by real factors only. Nominal shocks can affect real variables and relative prices in the short run but not in the long run. This idea is not tested: Long-run neutrality is *imposed* on the statistical model. In return it becomes possible to estimate separately the importance of nominal shocks to the short- and medium-run variability of output, hours, and real interest rates.⁷

In the five-variable scheme that Shapiro and Watson use, there are two nominal variablesthe inflation rate and the nominal interest rate and three real ones-output, hours, and the relative price of oil. They assume as well five shocks, two of which are nominal in the sense of having no effect on real variables in the long run. They are not able to measure the effects of the two dimensions of nominal instability separately. The other three shocks are taken to be real. The assumed exogeneity of oil price shocks plus a long-run neutrality hypothesis on hours are used to estimate the importance of three distinct real shocks. This aspect of their identification seems to me questionable, and in any case it is of an entirely different nature from the neutrality of nominal shocks. I will just lump the effects of the real shocks together, as

TABLE 1—PERCENTAGE OF VARIANCE DUE TO NOMINAL SHOCKS AT DIFFERENT FORECAST HORIZONS

Quarter	Output	Hours	Inflation	Interest rate
1	28	36	89	83
4	28	40	82	71
8	20	31	82	72
12	17	27	84	74
20	12	20	86	79
36	8	12	89	85
00	0	0	94	94

Shapiro and Watson do with the two nominal shocks, and interpret their paper as partitioning the variance of output and hours into nominal and real sources. The resulting Table 1 is a condensation of their Table 2.

The two zeroes for output and hours in the last, long-run, row of Table 1 are there by the *definition* of a nominal shock. But the two 94percent entries in this row for inflation and the nominal interest rate could have come out any way. I take the fact that these values are so close to 1 as a confirmation of Shapiro and Watson's procedure for identifying nominal shocks. According to Table 1, these nominal shocks have accounted for something less than 30 percent of short-run production variability in the postwar United States. This effect decays slowly, with no change after one year, a reduction to 20 percent after two years, and so on.

One can ask whether a better estimate of the importance of nominal shocks could have obtained by using M1 or some other observable measure of monetary shocks. Many studies have proceeded in this more direct way,8 and much has been learned, but in the end one does not know whether the importance of monetary shocks has been estimated or just the importance of a particular, possibly very defective, measure of them. Information on future prices is conveyed to people by changes in monetary aggregates, of course, but it is also conveyed by interest-rate and exchange-rate movements, by changes in the fiscal situation that may lead to tighter or easier money later on, by changes in financial regulations, by statements of influential people, and by many other factors. Shapiro and Watson's method bypasses these hard

⁶ To remove any doubt on the latter point, they quote from Milton Friedman's (1968) Presidential Address.

⁷ A similar, and similarly motivated, identification procedure was used in Olivier J. Blanchard and Danny Quah (1989). Thomas J. Sargent and Christopher A. Sims (1977) is a predecessor in spirit, if not in detail.

⁸ For example, Lawrence J. Christiano, et al. (1996).

measurement questions and goes directly to an estimation of the importance of nominal shocks in general, those we know how to measure and those we do not, whatever they may be.

A second reason for preferring the procedure Shapiro and Watson used is that the effects of nominal shocks as they estimate them include the effects of real shocks that could have been offset by monetary policy but were not. Whatever it is that keeps prices from rising in proportion to a given increase in money must also keep relative prices from adjusting as neoclassical theory would predict they should to, say, an increase in the OPEC-set price of oil. Effects of either kind-those initiated by monetary changes and those initiated by real shocks-will last only as long as the rigidity or glitch that gives rise to them lasts, vanishing in the long run, and will be identified as arising from the "nominal," or "demand," shock under the Shapiro and Watson identification procedure. Thus I want to interpret the estimates in columns 2 and 3 of Table 1 as upper bounds on the variance that could have been removed from output and hours at different horizons under some monetary policy other than the one actually pursued. The table gives no information on what this variance-minimizing monetary policy might have been, and there is no presumption that it would have been a policy that does not respond to real shocks.

Shapiro and Watson applied the theoretical idea that nominal shocks should be neutral in the long run to obtain an estimate of the fraction of short-run output variability that can be attributed to such shocks. Prescott (1986a) proceeded in a quite different way to arrive at an estimate of the fraction of output variability that can be attributed to technology shocks. He used actual Solow residuals to estimate the variance and serial correlation of the underlying technology shocks. Feeding shocks with these properties into a fully calibrated real-business-cycle model resulted in output variability that was about 84 percent of actual variability.9 In a complementary study, S. Rao Aiyagari (1994) arrived at an estimate of 79 percent for the contribution of technology shocks, based on comovements of production and labor input over the cycle.

Shapiro and Watson find that at most 30 percent of cyclical output variability can be attributed to nominal shocks. Working from the opposite direction, Prescott and Aiyagari conclude that at least 75 percent of cyclical output variability must be due to technology shocks. These findings are not as consistent as they may appear, because there are important real factors besides technological shocks-shocks to the tax system, to the terms of trade, to household technology, or to preferences-that are cyclically important but not captured in either of the categories I have considered so far.¹⁰ Even so, on the basis of this evidence I find it hard to imagine that more than 30 percent of the cyclical variability observed in the postwar United States could or should be removed by changes in the way monetary and fiscal policy is conducted.

IV. Risk Aversion

The estimate of the potential gains from stabilization reviewed in Section II rests on assumed consumer preferences of the constant relative risk aversion (CRRA) family, using but two parameters—the subjective discount rate ρ and the risk-aversion coefficient γ —to characterize all households. This preference family is almost universally used in macroeconomic and public finance applications. The familiar formula for an economy's average return on capital under CRRA preferences,

(6)
$$r = \rho + \gamma g$$
,

where g is the growth rate of consumption, makes it clear why fairly low γ values must be used. Per capita consumption growth in the United States is about 0.02 and the after-tax return on capital is around 0.05, so the fact that ρ must be positive requires that γ in (6) be at most 2.5. Moreover, a value as high as 2.5 would imply much larger interest rate differen-

⁹ Questions of measurement errors are discussed in the paper and by Summers (1986) in the same volume. In Prescott (1986b), estimates of 0.5 to 0.75 for the contribution of technology shocks to output variance are proposed.

¹⁰ For example, Shapiro and Watson attribute a large share of output variance to a shock which they call "labor supply" [and which I would call "household technology," following Jess Benhabib et al. (1991) and Jeremy Greenwood and Zvi Hercowitz (1991)].

7

tials than those we see between fast-growing economies like Taiwan and mature economies like the United States. This is the kind of evidence that leads to the use of γ values at or near 1 in applications.

But the CRRA model has problems. Rajnish Mehra and Prescott (1985) showed that if one wants to use a stochastic growth model with CRRA preferences to account for the entire return differential between stocks and bonds—historically about 6 percent—as a premium for risk, the parameter γ must be enormous, perhaps 50 or 100.¹¹ Such values obviously cannot be squared with (6). This "equity premium puzzle" remains unsolved, and has given rise to a vast literature that is clearly closely related to the question of assessing the costs of instability.¹²

One response to the puzzle is to adopt a three- rather than two-parameter description of preferences. Larry G. Epstein and Stanley E. Zin (1989, 1991) and Philippe Weil (1990) proposed different forms of *recursive utility*, preference families in which there is one parameter to determine intertemporal substitutability and a second one to describe risk aversion. The first corresponds to the parameter γ in (6), and can be assigned a small value to fit estimated average returns to capital. Then the risk-aversion parameter can be chosen as large as necessary to account for the equity premium.

Thomas D. Tallarini, Jr. (2000) uses preferences of the Epstein-Zin type, with an intertemporal substitution elasticity of 1, to construct a real-business-cycle model of the U.S. economy. He finds an astonishing separation of quantity and asset price determination: The behavior of aggregate quantities depends hardly at all on attitudes toward risk, so the coefficient of risk aversion is left free to account for the equity premium *perfectly*.¹³ Tallarini estimates a welfare cost of aggregate consumption risk of 10 percent of consumption, comparable to some of the supply-side gains cited in Section I, and two orders of magnitude larger than the estimate I proposed in Section II.¹⁴ As Maurice Obstfeld (1994) shows, this result is basically the formula (4) with a coefficient of risk aversion two orders of magnitude larger than the one I used.

Fernando Alvarez and Urban J. Jermann (2000) take a nonparametric approach to the evaluation of the potential gains from stabilization policy, relating the marginal cost of business-cycle risk to observed market prices without ever committing to a utility function. Their estimation procedure is based on the observation that consumption streams with a wide variety of different risk characteristics-or something very nearly equivalent to them-are available for sale in securities markets. They use a mix of asset-pricing theory and statistical methods to infer the prices of a claim to the actual, average consumption path and alternative consumption paths with some of the uncertainty removed. They call the price differentials so estimated *marginal* welfare costs, and show that they will be upper bounds to the corresponding total cost: my compensation parameter λ . The basic underlying hypotheses are that asset markets are complete and that asset-price differences reflect risk and timing differences and nothing else.

The gain from the removal of *all* consumption variability about trend, estimated in this way, is large—around 30 percent of consumption.¹⁵ This is a reflection of the high risk aversion needed to match the 6-percent equity premium, and can be compared to Tallarini's estimate of 10 percent. But the gain from removing risk at what Alvarez and Jermann call business-cycle frequencies—cycles of eight

¹¹ See also Lars Peter Hansen and Kenneth J. Singleton (1983).

¹² Two especially informative surveys are John H. Cochrane and Hansen (1992) and Narayana R. Kocherla-kota (1996).

¹³ Similar results, obtained in a closely related context, were reported by Hansen et al. (1999).

¹⁴ James Dolmas (1998) uses still another preference family, obtaining much higher cost estimates than mine. Like Tallarini, Christopher Otrok (1999) develops and analyzes a complete real-business-cycle model. He uses a preference family proposed by John Heaton (1995). His cost estimates are close to mine. A recent paper by Anne Epaulard and Aude Pommeret (2001) contains further results along this line, and provides a very useful quantitative comparison to earlier findings.

¹⁵ Alvarez and Jermann offer many estimates in their Tables 2A–2D. My summary is based on Table 2D, which uses postwar (1954–1997) data and requires that consumption and dividends be cointegrated. From this table, I follow the authors and cite averages over the columns headed "8 years" and "inf."

years or less—is two orders of magnitude smaller, around 0.3 percent. Most of the high return on equity is estimated to be compensation for long-term risk only, risk that could not be much reduced by short-run policies that are neutral in the long run.

Accepting Shapiro and Watson's finding that less that 30 percent of output variance at business-cycle frequencies can be attributed to nominal shocks, the lower Alvarez and Jermann estimate of 0.3 should be reduced to 0.1 if it is to serve my purpose as an estimate of the value of potential improvements in stabilization policy. But it is important to keep in mind that this estimate is not smaller than Tallarini's because of a different estimate of risk aversion. Tallarini's estimate of $\gamma = 100$ is the parametric analogue of Alvarez and Jermann's "market price of risk," based on exactly the same resolution of the equity premium puzzle. The different cost estimate is entirely due to differences in the consumption paths being compared.

Resolving empirical difficulties by adding new parameters always works, but often only by raising more problems. The risk-aversion levels needed to match the equity premium, under the assumption that asset markets are complete, ought to show up somewhere besides securities prices, but they do not seem to do so. No one has found risk-aversion parameters of 50 or 100 in the diversification of individual portfolios, in the level of insurance deductibles, in the wage premiums associated with occupations with high earnings risk, or in the revenues raised by state-operated lotteries. It would be good to have the equity premium resolved, but I think we need to look beyond high estimates of risk aversion to do it. The great contribution of Alvarez and Jermann is to show that even using the highest available estimate of risk aversion, the gain from further reductions in businesscycle risk is below one-tenth of 1 percent of consumption. The evidence also leaves one free to believe-as I do-that the gain is in fact one or two orders of magnitude smaller.

V. Incomplete Markets and Distribution Effects

The calculations I have described so far treat households as identical and individual risks as diversifiable. But as Per Krusell and Anthony A. Smith, Jr. (1999) observe, "it is quite plausible that the welfare costs of cycles are not so high on average, but may be very high for, say, the very poor or currently unemployed members of society." Several recent studies have pursued this possibility.¹⁶ Doing so evidently requires models with incomplete risk sharing and differently situated agents.

Krusell and Smith (1999, 2002) study a model economy in which individual families are subject to three kinds of stochastic shocks. There is an aggregate productivity shock that affects everyone, and employment shocks that differ from person to person. Families are infinitely lived dynasties, but every 40 years or so a family draws a new head, whose subjective discount rate is drawn from a fixed distribution. Dynasties with patient heads will accumulate wealth while others will run their wealth down.¹⁷ The sizes of these shocks are chosen so that the model economy experiences realistic GDP fluctuations, unemployment spells have realistic properties, and the overall wealth distribution matches the U.S. distribution: In the model, the wealthiest 5 percent of households own 54 percent of total wealth; in reality, they hold 51 percent.

It is essential to the substantive question that motivates this study that neither the employment shocks nor the uncertainty about the character of the household head can be diversified away. Otherwise, the individual effects of the aggregate productivity shocks would be the same as in the representative agent models I have already discussed. One may argue over why it is that markets do not permit such diversification, but it seems clear enough that they do not: Where is the market where people can be insured against the risk of having irresponsible or incompetent parents or children?

These exogenous forces acting differentially across households induce different individual choices, which in turn lead to differences in individual capital holdings. The state space in this economy is very large, much larger than

¹⁶ For example, Ayse Imrohoroğlu (1989), Andrew Atkeson and Christopher Phelan (1994), Krusell and Smith (1999, 2002), Kjetil Storesletten et al. (2001), and Tom Krebs (2002).

¹⁷ This way of modeling wealth changes within a fixed distribution across families was introduced in John Laitner (1992).

anything people were working with numerically 15 years ago, and without the method developed in Krusell and Smith (1998) it would not have been possible to work out the predictions of this model. A key simplification comes from the fact that the impact on any one family of the shocks that hit others has to work through two prices, the real wage and the rental price of capital. These prices in turn depend only on the total stock of capital, regardless of the way it is distributed, and total employment, regardless of who has a job and who does not. By exploiting these features, solutions can be calculated using an iterative procedure that works like a dream: For determining the behavior of aggregates, they discovered, realistically modeled household heterogeneity just does not matter very much.

For *individual* behavior and welfare, of course, heterogeneity is everything. In the thought-experiments that Krusell and Smith run with their model, removal of the business cycle is defined to be equivalent to setting the aggregate productivity shock equal to a constant. It is important to be clear on what the effect of such a change would be on the behavior of the employment shocks to which individuals are subject, but the magical character of the experiment makes it hard to know how this question is best resolved. I will describe what Krusell and Smith did, and deal with some other possibilities later on.

Suppose that a shock $y = az + \varepsilon$ affects an individual's behavior, where z is the aggregate shock and ε is idiosyncratic. We project the individual shock on the aggregate, $\varepsilon = cz + \eta$, where the residual η is uncorrelated with z, and then think of an ideal stabilization policy as one that replaces

$$y = az + \varepsilon = (a + c)z + \eta$$

with

$$\hat{y} = (a+c)E(z) + \eta.$$

Not only is the direct effect of the productivity shock z removed but also the indirect effects of z on the individual employment shocks ε .¹⁸ In this particular application, removing the variance of the aggregate shock is estimated to reduce the standard deviation of the individual employment shocks by 16 percent.¹⁹

The first such thought-experiment Krusell and Smith describe involves a comparison between the expected utility drawn from the steady state of the economy with aggregate shocks and the expected utility from the steady state of the economy with aggregate shocks and their indirect effects removed in the way I have just described. The welfare gain from eliminating cycles in this sense turns out to be negative! In a model, like this one, in which markets for risk pooling are incomplete, people will engage in precautionary savings, overaccumulating capital in the effort to self-insure. This implies larger average consumption in the more risky economy. Of course, there are costs to accumulating the higher capital stock, but these costs are not fully counted in a steady-state comparison.

In any case, as Krusell and Smith emphasize, there is nothing really distributional about a steady-state comparison: Every infinitely lived dynasty is assigned a place in the wealth distribution at random, and no one of them can be identified as permanently rich or poor. The whole motivation of the paper is to focus on the situation of people described as "hand-to-mouth consumers," but a steady-state comparison misses them. This observation motivates a second thought-experiment-one with much more complicated dynamics than the first-in which an economy is permitted to reach its steadystate wealth distribution with realistic aggregate shocks, and then is relieved of aggregate risk. The full transition to a new steady state is then worked out and taken into account in the utility comparisons. In this experiment, we can identify individuals as "rich" or "poor" by their position in the initial wealth distribution, and discuss the effects of risk removal category by category.

The average welfare gain in this second experiment is about 0.1 of 1 percent of consumption, about twice the estimate in Section II of this paper. (Krusell and Smith also assume log utility.) But this figure masks a lot of diversity. Low wealth, unemployed people—people who

¹⁸ This is a linear illustration of the more generally defined procedure described in Krusell and Smith (1999).

¹⁹ Here and below, the numbers I cite are taken from Krusell and Smith (2002).

would borrow against future labor income if they could—enjoy a utility gain equivalent to a 4-percent perpetual increase in consumption. Oddly, the very wealthy can also gain, as much as 2 percent. Krusell and Smith conjecture that this is due to the higher interest rates implied by the overall decrease in precautionary savings and capital. Finally, there is a large group of middle wealth households that are made worse off by eliminating aggregate risk.

These calculations are sensitive—especially at the poor end of the distribution—to what is assumed about the incomes of unemployed people. Krusell and Smith calibrate this, roughly, to current U.S. unemployment insurance replacement rates. If one were estimating the costs of the depression of the 1930's, before the current welfare system was in place, lower rates would be used and the cost estimates would increase sharply.²⁰ It would also be interesting to use a model like this to examine the trade-offs between reductions in aggregate risk and an improved welfare system.

Storesletten et al. (2001) study distributional influences on welfare cost estimates with methods that are closely related to Krusell and Smith's, but they obtain larger estimates of the gains from removing all aggregate shocks. They use an overlapping generations setup with 43 working age generations, in which the youngest cohort is always credit constrained. In such a setting, the young are helpless in the face of shocks of all kinds and reductions in variance can yield large welfare gains. But if the age effects are averaged out to reflect the importance of intrafamily lending (as I think they should be) the gains estimated by Storesletten et al. under log utility are no larger than Krusell and Smith's.²¹ In contrast to earlier studies, however, the Storesletten et al. model implies that estimated welfare gains rise faster than proportionately as risk aversion is increased: From Exhibit 2, for example, the average gain increases from 0.6 of a percent to 2.5 as γ is increased from 2 to 4.

Two features of the theory interact to bring this about.²² First, and most crucial, is a differ-

ence in the way reductions in the variance of aggregate shocks affect risks faced at the individual level. In the Storesletten et al. simulations, a bad realization of the aggregate productivity shock increases the conditional *variance* of the idiosyncratic risk that people face, so aggregate and individual risks are compounded in a way that Krusell and Smith rule out. A second difference is that idiosyncratic shocks are assumed to have a random walk component, so their effects are long lasting. A bad aggregate shock increases the chances that a young worker will draw a bad individual shock, and if he does he will suffer its effects throughout his prime working years.

The effects of these two assumptions are clear: They convert small, transient shocks at the aggregate level into large, persistent shocks to the earnings of a small fraction of households. Whether they are realistic is question of fact. That individual earnings differences are highly persistent has been clear since Lee Lillard and Robert Willis's pioneering (1978) study. The fanning out over time of the earnings and consumption distributions within a cohort that Angus Deaton and Christina Paxson (1994) document is striking evidence of a sizeable, uninsurable random walk component in earnings. The relation of the variance of earnings shocks to the aggregate state of the economy, also emphasized by N. Gregory Mankiw (1986) in connection with the equity premium puzzle, has only recently been studied empirically. Storesletten et al. find a negative relation over time between cross-section earnings means and standard deviations in Panel Studies of Income Dynamics data. Costas Meghir and Luigi Pistaferri (2001) obtain smaller estimates, but also conclude that "the unemployment rate and the variance of permanent [earnings] shocks appear to be quite synchronized" in the 1970's and 1980's.

These issues are central to an accurate description of the risk situation that individual agents face, and hence to the assessment of welfare gains from policies that alter this situation. The development of tractable equilibrium models capable of bringing cross-section and panel evidence to bear on this and other macroeconomic questions is an enormous step forward. But Krusell and Smith find only modest effects of heterogeneity on the estimates of wel-

²⁰ See Satyajit Chatterjee and Dean Corbae (2000).

²¹ Based on Exhibits 2 and A.3.1.

²² Storesletten et al. do a good job of breaking the differences into intelligible pieces. I also found the example explicitly solved in Krebs (2002) very helpful in this regard.

fare gains from the elimination of aggregate risk, and even accepting the Storesletten et al. view entails an upward revision of a factor of only about 5.

The real promise of the Krusell-Smith model and related formulations, I think, will be in the study of the relation of policies that reduce the impact of risk by reducing the variance of shocks (like aggregate stabilization policies) to those that act by reallocating risks (like social insurance policies). Traditionally, these two kinds of policies have been studied by different economists, using unrelated models and different data sets. But both appear explicitly in the models I have reviewed here, and it is clear that it will soon be possible to provide a unified analysis of their costs and benefits.

VI. Other Directions

My plan was to go down a list of all the things that could have gone wrong with my 1987 calculations but, as I should have anticipated, possibilities were added to the list faster than I could eliminate them. I will just note some of the more interesting of these possibilities, and then conclude. The level of consumption risk in a society is, in part, subject to choice. When in an economy that is subject to larger shocks, people will live with more consumption variability and the associated loss in welfare, but they may also substitute into riskavoiding technologies, accepting reduced average levels of production. This possibility shows up in the precautionary savings-overaccumulation of capital-that Krusell and Smith (1999, 2002) found. As Garey Ramey and Valerie A. Ramey (1991) suggested, this kind of substitution surely shows up in other forms as well.

In an endogenous growth framework, substitution against risky technologies can affect rates of growth as well as output levels. Larry E. Jones et al. (1999) and Epaulard and Pommeret (2001) explore some of these possibilities, though neither study attributes large welfare gains to volatility-induced reductions in growth rates. Gadi Barlevy (2001) proposes a convex adjustment cost that makes an erratic path of investment in knowledge less effective than a smooth path at the same average level. In such a setting, reducing shock variability can lead to higher growth even without an effect on the average level of investment. He obtains welfare gains as large as 7 percent of consumption in models based on this idea, but everything hinges on a curvature parameter on which there is little evidence. This is a promising frontier on which there is much to be done. Surely there are others.

VII. Conclusions

If business cycles were simply efficient responses of quantities and prices to unpredictable shifts in technology and preferences, there would be no need for distinct stabilization or demand management policies and certainly no point to such legislation as the Employment Act of 1946. If, on the other hand, rigidities of some kind prevent the economy from reacting efficiently to nominal or real shocks, or both, there is a need to design suitable policies and to assess their performance. In my opinion, this is the case: I think the stability of monetary aggregates and nominal spending in the postwar United States is a major reason for the stability of aggregate production and consumption during these years, relative to the experience of the interwar period and the contemporary experience of other economies. If so, this stability must be seen in part as an achievement of the economists, Keynesian and monetarist, who guided economic policy over these years.

The question I have addressed in this lecture is whether stabilization policies that go beyond the general stabilization of spending that characterizes the last 50 years, whatever form they might take, promise important increases in welfare. The answer to this question is "No": The potential gains from improved stabilization policies are on the order of hundredths of a percent of consumption, perhaps two orders of magnitude smaller than the potential benefits of available "supply-side" fiscal reforms. This answer does depend, certainly, on the degree of risk aversion. It does not appear to be very sensitive to the way distribution effects are dealt with, though it does presuppose a system of unemployment insurance at postwar U.S. levels. I have been as explicit as I can be on the way theory and evidence bear on these conclusions.

When Don Patinkin gave his *Money, Interest, and Prices* the subtitle "An Integration of Monetary and Value Theory," value theory meant, to him, a purely static theory of general equilibrium. Fluctuations in production and employment, due to monetary disturbances or to shocks of any other kind, were viewed as inducing disequilibrium adjustments, unrelated to anyone's purposeful behavior, modeled with vast numbers of free parameters. For us, today, value theory refers to models of dynamic economies subject to unpredictable shocks, populated by agents who are good at processing information and making choices over time. The macroeconomic research I have discussed today makes essential use of value theory in this modern sense: formulating explicit models, computing solutions, comparing their behavior quantitatively to observed time series and other data sets. As a result, we are able to form a much sharper quantitative view of the potential of changes in policy to improve peoples' lives than was possible a generation ago.

REFERENCES

- Aiyagari, S. Rao. "On the Contribution of Technology Shocks to Business Cycles." *Federal Reserve Bank of Minneapolis Quarterly Review*, Summer 1994, *18*(3), pp. 22–34.
- Alvarez, Fernando and Jermann, Urban J. "Using Asset Prices to Measure the Cost of Business Cycles." Working paper, University of Chicago, 2000.
- Atkeson, Andrew and Phelan, Christopher. "Reconsidering the Costs of Business Cycles with Incomplete Markets." *NBER macroeconomics annual, 1994.* Cambridge, MA: MIT Press, 1994, pp. 187–207.
- Auerbach, Alan J. and Kotlikoff, Laurence J. Dynamic fiscal policy. Cambridge: Cambridge University Press, 1987.
- Bailey, Martin J. "The Welfare Cost of Inflationary Finance." *Journal of Political Economy*, April 1956, 64(2), pp. 93–110.
- **Barlevy, Gadi.** "The Cost of Business Cycles Under Endogenous Growth." Working paper, Northwestern University, 2001.
- Benhabib, Jess; Rogerson, Richard and Wright, Randall. "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." *Journal of Political Economy*, December 1991, *99*(6), pp. 1166–87.
- Blanchard, Olivier J. and Quah, Danny. "The Dynamic Effects of Aggregate Demand

and Supply Disturbances." American Economic Review, September 1989, 79(4), pp. 655–73.

- Brock, William A. and Turnovsky, Stephen J. "The Analysis of Macroeconomic Policies in Perfect Foresight Equilibrium." *International Economic Review*, February 1981, 22(1), pp. 179–209.
- Chamley, Christophe P. "The Welfare Cost of Capital Income Taxation in a Growing Economy." *Journal of Political Economy*, June 1981, 89(3), pp. 468–96.
- Chatterjee, Satyajit and Corbae, Dean. "On the Welfare Gains of Reducing the Likelihood of Economic Crises." Federal Reserve Bank of Philadelphia Research Working Paper No. 00/14, November 2000.
- Christiano, Lawrence J.; Eichenbaum, Martin and Evans, Charles. "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds." *Review of Economics and Statistics*, February 1996, 78(1), pp. 16–34.
- Cochrane, John H. and Hansen, Lars Peter. "Asset Pricing Explorations for Macroeconomics." *NBER macroeconomics annual*, 1992. Cambridge, MA: MIT Press, 1992, pp. 115–65.
- **Deaton, Angus and Paxson, Christina.** "Intertemporal Choice and Inequality." *Journal of Political Economy*, June 1994, *102*(3), pp. 437–67.
- **Dolmas, James.** "Risk Preferences and the Welfare Cost of Business Cycles." *Review of Economic Dynamics*, July 1998, *1*(3), pp. 646–76.
- Epaulard, Anne and Pommeret, Aude. "Recursive Utility, Endogenous Growth, and the Welfare Cost of Volatility." International Monetary Fund Working Paper No. WP/01/ 05, 2001.
- **Epstein, Larry G. and Zin, Stanley E.** "Substitution, Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns I: A Theoretical Framework." *Econometrica*, July 1989, *57*(4), pp. 937–69.
- . "Substitution, Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns II: An Empirical Analysis." *Journal of Political Economy*, April 1991, 99(2), pp. 263–86.
- Friedman, Milton. "The Role of Monetary Policy." American Economic Review, March 1968, 58(1), pp. 1–17.

- Greenwood, Jeremy and Hercowitz, Zvi. "The Allocation of Capital and Time over the Business Cycle." *Journal of Political Economy*, December 1991, *99*(6), pp. 1188–214.
- Hall, Robert E. and Rabushka, Alvin. *The flat tax*, 2nd Ed. Stanford, CA: Hoover Institution Press. 1995.
- Hansen, Gary D. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics*, November 1985, 16(3), pp. 309–27.
- Hansen, Lars Peter; Sargent, Thomas J. and Tallarini, Thomas D., Jr. "Robust Permanent Income and Pricing." *Review of Economic Studies*, October 1999, 66(4), pp. 873–908.
- Hansen, Lars Peter and Singleton, Kenneth J. "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *Journal of Political Economy*, April 1983, 91(2), pp. 249–65.
- Heaton, John. "An Empirical Investigation of Asset Pricing with Temporally Dependent Preference Specifications." *Econometrica*, May 1995, *63*(3), pp. 681–717.
- Imrohoroğlu, Ayse. "The Cost of Business Cycles with Indivisibilities and Liquidity Constraints." *Journal of Political Economy*, December 1989, 97(6), pp. 1364–83.
- Jones, Larry E.; Manuelli, Rodolfo E. and Stacchetti, Ennio. "Technology (and Policy) Shocks in Models of Endogenous Growth." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 7063, April 1999.
- Judd, Kenneth L. "The Welfare Cost of Factor Taxation in a Perfect Foresight Model." *Journal of Political Economy*, August 1987, 95(4), pp. 675–709.
- Kocherlakota, Narayana R. "The Equity Premium: It's Still a Puzzle." *Journal of Economic Literature*, March 1996, *34*(1), pp. 42–71.
- Krebs, Tom. "Growth and Welfare Effects of Business Cycles in Economies with Idiosyncratic Human Capital Risk." Working paper, Brown University, 2002.
- Krusell, Per and Smith, Anthony A., Jr. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy*, October 1998, 106(5), pp. 867–96.
 - . "On the Welfare Effects of Eliminating Business Cycles." *Review of Economic Dynamics*, January 1999, 2(2), pp. 245–72.
 - _____. "Revisiting the Welfare Effects of

Eliminating Business Cycles." Working paper, Carnegie Mellon University, 2002.

- **Kydland, Finn E. and Prescott, Edward C.** "Time to Build and Aggregate Fluctuations." *Econometrica*, November 1982, *50*(6), pp. 1345–70.
- Laitner, John. "Random Earnings Differences, Lifetime Liquidity Constraints, and Altruistic Intergenerational Transfers." *Journal of Economic Theory*, December 1992, *58*(2), pp. 135–70.
- Lillard, Lee and Willis, Robert. "Dynamic Aspects of Earnings Mobility." *Econometrica*, September 1978, *46*(5), pp. 985–1012.
- Lucas, Robert E., Jr. Models of business cycles. New York: Basil Blackwell, 1987.
- _____. "Inflation and Welfare." *Econometrica*, March 2000, 68(2), pp. 247–74.
- Mankiw, N. Gregory. "The Equity Premium and the Concentration of Aggregate Shocks." *Journal of Financial Economics*, September 1986, *17*(1), pp. 211–19.
- Meghir, Costas and Pistaferri, Luigi. "Income Variance Dynamics and Heterogeneity." Working paper, Institute for Fiscal Studies, 2001.
- Mehra, Rajnish and Prescott, Edward C. "The Equity Premium: A Puzzle." *Journal of Monetary Economics*, March 1985, *15*(2), pp. 145–61.
- Meltzer, Allan H. "The Demand for Money: The Evidence from Time Series." *Journal of Political Economy*, June 1963, *71*(3), pp. 219–46.
- **Obstfeld, Maurice.** "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability." *European Economic Review*, August 1994, *38*(7), pp. 1471–86.
- **Otrok, Christopher.** "On Measuring the Welfare Cost of Business Cycles." Working paper, University of Virginia, 1999.
- **Prescott, Edward C.** "Theory Ahead of Business Cycle Measurement." *Federal Reserve Bank* of Minneapolis Quarterly Review, Fall 1986a, 10(4), pp. 9–21.
- . "Response to a Skeptic." *Federal Reserve Bank of Minneapolis Quarterly Review*, Fall 1986b, *10*(4), pp. 28–33.
- . "Richard T. Ely Lecture: Prosperity and Depression." *American Economic Review*, May 2002 (*Papers and Proceedings*), 92(2), pp. 1–15.

- Ramey, Garey and Ramey, Valerie A. "Technology Commitment and the Cost of Economic Fluctuation." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 3755, June 1991.
- Rosen, Sherwin. "Public Employment and the Welfare State in Sweden." *Journal of Economic Literature*, June 1996, *34*(2), pp. 729–40.
- Sargent, Thomas J. and Sims, Christopher A. "Business Cycle Modeling without Pretending to Have Too Much A-Priori Economic Theory," in *New methods in business cycle research*. Proceedings from a Conference, Federal Reserve Bank of Minneapolis, 1977, pp. 45–109.
- Shapiro, Matthew D. and Watson, Mark W. "Sources of Business Cycle Fluctuations." *NBER macroeconomics annual, 1988.* Cambridge, MA: MIT Press, 1988, pp. 111–48.
- Solow, Robert M. "Technical Change and the Aggregate Production Function." *Review of*

Economics and Statistics, August 1957, 39(3), pp. 312–20.

- Storesletten, Kjetil; Telmer, Chris I. and Yaron, Amir. "The Welfare Cost of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk." *European Economic Review*, June 2001, 45(7), pp. 1311– 39.
- Summers, Lawrence H. "Capital Taxation and Accumulation in a Life Cycle Growth Model." *American Economic Review*, September 1981, 71(4), pp. 533–44.
- . "Some Skeptical Observations on Real Business Cycle Theory." *Federal Reserve Bank of Minneapolis Quarterly Review*, Fall 1986, *10*(4), pp. 22–27.
- Tallarini, Thomas D., Jr. "Risk-Sensitive Real Business Cycles." *Journal of Monetary Economics*, June 2000, 45(3), pp. 507–32.
- Weil, Philippe. "Non-Expected Utility in Macroeconomics." *Quarterly Journal of Economics*, February 1990, *105*(1), pp. 27–42.