Volume Title: NBER Macroeconomics Annual 1994, Volume 9

Volume Author/Editor: Stanley Fischer and Julio J. Rotemberg, eds.

Volume Publisher: MIT Press

Volume ISBN: 0-262-05172-4

Volume URL: http://www.nber.org/books/fisc94-1

Conference Date: March 11-12, 1994

Publication Date: January 1994

Chapter Title: Reconsidering the Costs of Business Cycles with Incomplete Markets

Chapter Author: Andrew Atkeson, Christopher Phelan

Chapter URL: http://www.nber.org/chapters/c11010

Chapter pages in book: (p. 187 - 218)

Andrew Atkeson and Christopher Phelan
UNIVERSITY OF PENNSYLVANIA AND NBER, AND NORTHWESTERN UNIVERSITY AND UNIVERSITY OF WISCONSIN-MADISON

Reconsidering the Costs of Business Cycles with Incomplete Markets

1. Introduction

In *Models of Business Cycles*, Robert Lucas (1987) puts forward a disarmingly simple argument that the potential welfare gains from eliminating the fluctuations in aggregate consumption associated with business cycles are, at best, extremely small. His argument is as follows: Assume that aggregate consumption is described by the statistical model $c_t = (1 + g)^t z_t$, where z_t is a lognormally distributed random shock and $(1 + g)$ is a deterministic trend. Assume that there are complete markets or perfect insurance against individual income risk, and that consumers have identical constant relative risk aversion (CRRA) preferences, so that individual consumption is simply a fraction of aggregate consumption. Then ask the question, How much would consumers pay to move to a world in which aggregate (and individual) consumption has no fluctuations? In this alternative world, aggregate consumption is described by the model $c_t = (1 + g)^t$. In particular, solve for the fraction λ of consumers' current consumption stream that satisfies

$$E_0 \sum_{t=0}^{\infty} \beta^t \big((1 + \lambda)(1 + g)^t z_t\big)^{\gamma} = \sum_{t=0}^{\infty} \beta^t (1 + g)^{t\gamma}. \tag{1}$$

Interpret λ as the constant fraction of aggregate consumption at each date, and state that consumers would need to be paid to give them the same utility they obtain from the consumption stream with no aggregate fluctuations. With the assumption that log (z_t) is distributed $N(-\sigma_z^2/2, \sigma_z^2)$ we can calculate $(1 + \lambda) = \exp((1 - \gamma)\sigma_z^2/2)$. Using numbers like $\gamma = 0$ (log utility) or $\gamma = -4$ for the curvature of the utility function and .013 for σ_z, the standard deviation of aggregate consumption around trend, we get welfare costs of $\lambda = .00008$ or .00042. That is, we get the answer that the welfare costs of aggregate fluctuations are virtually zero.

One concern about Lucas's calculation of the welfare costs of aggregate fluctuations centers on the assumption in his model that there are complete markets for insuring individual income risk. In particular, in a setting with substantial idiosyncratic income risk and incomplete markets for sharing that risk, the marginal utility of consumption for each individual in the economy can be considerably more variable than would be the case if there were complete markets. Given this possibility, it would seem that large welfare gains might be obtained from a countercyclical policy if that policy, directly or indirectly, allowed consumers to obtain smoother consumption streams in equilibrium.

In this paper, we measure the potential welfare gains from countercyclical policy in an economy with incomplete markets. In the course of conducting this measurement, we focus on two questions as central to the determination of those potential gains: (1) What is the likely effect of countercyclical policy on the nature of the income risk faced by individuals in the economy? and (2) What are the likely general equilibrium effects brought about as asset prices change due to the implementation of countercyclical policies? In taking up the first question, we see it as critical to distinguish whether the main effect of countercyclical policy is to reduce directly the income risk faced by each individual or is simply to reduce the correlation across individuals in the income risk that they face. In either situation, countercyclical policy will have a general equilibrium effect on welfare if it changes asset prices. However, in the second situation, the situation in which countercyclical policy simply reduces the correlation across individuals in the risks that they face, this is the only effect such a policy will have.

In considering the likely effect of countercyclical policy on individual risk, we begin with the observation that one of the salient features of the business cycle is that fluctuations in aggregate hours worked and aggregate wages paid are not shared evenly across the population. That is, the income of workers employed continuously over the cycle does not fluctuate very much in comparison with the income fluctuations experienced by those who transit from employment to unemployment or from unemployment to employment over the cycle. This observation provides some justification for focusing on the unemployment risk an individual faces as the principal individual income risk connected to the business cycle.

In related work, Ayşe İmrohoroğlu (1989) presents a calculation of the costs of business cycles in an environment with incomplete markets that focuses on unemployment risk as the individual risk that would be affected by countercyclical policy. Her model does not allow any asset markets through which agents might share their unemployment risk. Instead, she assumes that each agent has an individual storage technology that he can use to smooth his consumption in response to his income fluctuations. To calculate the potential welfare gains from eliminating aggregate fluctuations, she compares agents' steady-state utility when they are faced with two different exogenous patterns of unemployment risk—one that represents the risks that individuals face when there are business cycles and one that represents the risks that they face when there are no business cycles. She finds costs of aggregate fluctuations that tend to be small, but whose size is quite sensitive to the exact specification of the individual's storage technology. We discuss her paper and its relation to our work in some detail.

In the next three sections, we reconsider the calculation of the potential welfare gains from countercyclical policy in an economy with incomplete markets. In the first of these sections, we discuss theoretically the different effects on welfare of countercyclical policies that reduce aggregate fluctuations by reducing individual income risk directly and countercyclical policies that reduce aggregate fluctuations by reducing the correlation across individuals in their income risk. In the second of these sections, we present a model of the wage and employment risk faced by individuals over the cycle in which the levels of those risks are chosen endogenously. On the basis of that model, we argue that the main effect of countercyclical policy aimed at reducing aggregate fluctuations may be simply to remove the correlation across individuals in the unemployment risk that they face. In this case, the main impact of countercyclical policy on individual welfare is through its general equilibrium impact on asset prices. In the third of these sections, we use asset price data in a model with incomplete markets to assess the potential gains from removing the correlation in individuals' unemployment risk. As a theoretical point, we show that the potential welfare gains from eliminating the correlation in individuals' income risk in a given environment is *smaller* when there are incomplete markets than when there are complete markets. On the basis of our interpretation of asset price data in an incomplete markets framework, we argue that the potential welfare gains from countercyclical policy are essentially zero.

2. *Aggregate and Individual Risk in Incomplete Markets*

Under the assumption that there are complete markets for ensuring individual income risk, aggregate risk is the only risk that affects individual consumption. If aggregate risk is eliminated, all individual consumption risk is also eliminated. For this reason, in calculating the welfare costs of aggregate risk with complete markets, it is not necessary to consider the direct effect of countercyclical policy on the processes that generate individual income streams. All that matters is the effect of countercyclical policy on aggregate income. With incomplete markets, on the other hand, this is not the case. With incomplete markets, the fluctuations in an individual's consumption are determined by the fluctuations in that individual's income and the extent to which that agent can trade his variable income for smooth consumption through storage technologies or asset markets. As a result, when one is calculating the welfare gains from countercyclical policy in an economy with incomplete markets, it is necessary to describe more precisely how such policies affect individuals' income risk and market opportunities. Obviously, this can be done in a number of different ways. In the following example, we highlight how one's calculation of the potential gains from countercyclical policy depends on various assumptions regarding individuals' income risk and market opportunities.

Consider a world with a continuum of agents in which, each period, each agent faces a probability $\pi(z)$ of being employed and, as a result, receiving high income y^h and probability $(1 - \pi(z))$ of being unemployed and receiving low income y^l, where z is an aggregate state of nature. Here z indexes the aggregate state of the economy in that it determines not only an individual's probability of receiving high income but also the proportion of individuals receiving high income. This proportion changes over time as the aggregate state z changes, and, thus, this economy experiences aggregate fluctuations.

Consider two ways that countercyclical policy might reduce aggregate fluctuations in this economy. One method would involve reducing the variance of individual income: In the extreme such a policy could eliminate aggregate fluctuations by eliminating entirely the variance in each individual's income. A second method would involve reducing the correlation across individuals in the income risk that they face: In the extreme, such a policy could eliminate aggregate fluctuations by eliminating the correlation across individuals in the income risk that they face without changing the characteristics of the income risk faced by any single individual. For the example earlier, the first type of policy might set individual income constant each period and equal to its unconditional mean $E_z\{\pi(z)y^h + (1 - \pi(z))y^l\}$. If z were i.i.d., the second type of policy might set the probability $\pi(z)$ constant and equal to its unconditional mean $E_z\pi(z)$ but leave the realizations y^h, y^l unchanged. This policy would eliminate aggregate fluctuations in income but would leave the unconditional distribution of individual income unchanged.

With complete markets, these two policies have the same impact on aggregate income and, thus, have the same impact on consumer welfare. With incomplete markets, these policies have different effects. To begin, assume that there are no asset markets and that agents have no ability to store consumption, so that individual consumption is always equal to individual income. In this case, under the first type of policy, each agent's utility is improved to the extent that the volatility of his individual income streams is reduced. Under the second type of policy, each agent's utility is unchanged since this type of policy does not change the unconditional distribution of agents' income and, thus, consumption, streams. Going further, assume that asset markets or technologies for storing consumption do exist but that, in equilibrium, countercyclical policy leaves asset prices or these technological returns unchanged. Then it is easy to show that, under the second type of policy, agents' utility is also unchanged. Agents use asset markets or storage technologies to transform income streams into consumption streams. If asset prices or technological returns are unchanged and the distribution of individual income streams is unchanged, then this transformation and, thus, individual utility are unchanged.

It is clear that if the main effect of countercyclical policy is to remove correlations in individual risk, the benefits of such a policy will be realized through the general equilibrium impact of such a policy on asset prices. If asset prices do change when the correlations in individual risk are reduced, then agents may gain if they are able to trade their variable income for smoother consumption at the new prices.

For example, one interpretation of Lucas's calculation of the welfare costs of aggregate risk under complete markets is that it is the answer to the question, How much would the representative agent pay to face asset prices that clear markets when aggregate consumption is constant as opposed to facing asset prices that clear markets when aggregate consumption fluctuates? When aggregate consumption fluctuates, asset prices also fluctuate to induce agents to choose to have fluctuating consumption. When aggregate consumption is constant, asset prices are constant, and agents choose to have constant consumption. The asset price fluctuations implied by Lucas's model are small. As a result, he finds a low cost of aggregate fluctuations.

Given this interpretation of Lucas's result, another concern about his calculation is that his model is not even remotely consistent with commonly observed features of asset price data. If the purpose of the model is to measure the amount that the representative agent is willing to pay to move from a world with asset prices like those currently observed to a world with asset prices that allow insurance against all risk, it would seem important that the model with aggregate fluctuations be consistent with currently observed asset prices. In particular, when interpreted in the context of a complete-markets, frictionless-trading, consumption-based model of asset prices, the equity premium is evidence that the representative consumer's marginal utility of consumption is dramatically different at different stages of the business cycle and, thus, that policies that smoothed consumption and marginal utility over the cycle could improve the utility of the representative consumer substantially.

The equity premium in Lucas's model is essentially zero. To understand the implications of the equity premium in standard versions of such models, recall that asset prices in these models are described by a stochastic process M_t, known as a pricing kernel, which satisfies $M_t = E_t[R_{i,t+1}M_{t+1}]$, where $R_{i,t+1}$ is the gross return on asset i if held from date t to date $t + 1$ and M_t is identified with the marginal utility of consumption at t for the representative consumer. As surveyed in Cochrane and Hansen (1992), the conclusion of empirical work on asset prices using frictionless, complete-markets, consumption-based models is that the pricing kernel and, thus, the marginal utility of consumption of the representative consumer must be extremely volatile if these models are to have hopes of matching the equity premium. In light of the concern that models that fail to predict a large equity premium may also fail to measure accurately the volatility of marginal utility and, thus, understate the welfare cost of business cycles, we design our model for measuring the potential benefits of countercyclical policy to be consistent with the equity premium and other data on the volatility of asset returns.

This idea that countercyclical policy may improve welfare by changing asset prices raises the question of how İmrohoroğlu (1989) finds positive gains to removing business cycles, since, in her model, she assumed that there were no asset markets but that agents had a linear storage technology for smoothing their income. The answer lies in the way she removes aggregate risk. She let the probability $\pi(z_{t+1})$ that an agent is employed at date $t + 1$ depend on whether the agent was employed at date t. She further allowed the aggregate shock z to follow a Markov process. In particular, she uses data on the conditional duration and level of unemployment in booms to define a transition matrix that determines an individual's movements between employment and unemployment in booms. Likewise, she defines the corresponding individual transition matrix for recessions. She also uses data on the duration of booms and recessions themselves to define a transition matrix that determines the evolution of aggregate state. She then specifies a world with no aggregate fluctuations by using data on the unconditional duration and level of unemployment to calibrate a single transition matrix determining individuals' movements between employment and unemployment.

One interpretation of her model is that individual employment and unemployment have two different levels of stability. One can think of an individuals' labor market status as having four states: stable employment, unstable employment, stable unemployment, unstable unemployment, with a transition matrix between these four states defined implicitly by the transition matrices described earlier. Aggregate fluctuations arise due to an assumed perfect correlation across individuals in the stability of their employment and unemployment. In booms, everyone has either stable employment or unstable unemployment, and in recessions, everyone has unstable employment or stable unemployment. This interpretation suggests an alternative method for eliminating aggregate fluctuations in her model. Maintain her four-state Markov process describing individuals' labor market transitions but eliminate correlations across individuals in these transitions, i.e., eliminate the correlation across individuals in the stability of their individual employment and unemployment. While both methods eliminate aggregate fluctuations, with prices pinned down by the assumed storage technology, the welfare gain to simply eliminating correlations in employment stability across individuals is identically zero.

3. *Countercyclical Policy and Endogenous Unemployment Risk*

We now consider the question of whether the main effect of countercyclical policy is to reduce individual income risk or to eliminate correlations across individuals in the income risk that they face. To address this question, we build a model of the income risk individuals face. In our model, labor market frictions prevent firms and workers from using the production technology to share the risk implied by random productivity, say, by proportionally reducing the wages received and hours worked by all agents in a recession. In our equilibrium, some workers become unemployed and receive no income, while other workers keep their jobs and receive income. The extent of wage and unemployment risk is determined endogenously: Workers trade off higher wages against a higher risk of unemployment. Also, in this model, equilibrium wages are rigid in the sense that anticipated, or long-term, changes in worker productivity are reflected in wage changes, while unanticipated, or short-term, changes in worker productivity are reflected in changes in the number of workers employed and unemployed. We demonstrate in this model that, while a countercyclical policy that smoothes out aggregate fluctuations in worker productivity does eliminate the wage risk faced by employed workers and the profit risk faced by owners of firms, it has no effect on the unemployment risk faced by an individual worker. Instead, this policy simply eliminates the correlations across workers in the unemployment risk that they face. Given that the wage risk faced by workers employed continuously over the cycle that we observe in the data is relatively small, we use this result to argue that the main effect of countercyclical policy on workers is to eliminate the correlations across individuals in the unemployment risk that they face.

Consider an economy with two types of agents: capitalists and workers. Time is discrete and denoted by $t = 0, 1, 2, 3, \ldots$. Each capitalist is endowed with a production technology that transforms the labor of one worker into θ units of consumption, where θ is random and cannot be verified by the workers. At each date, the distribution of the productivity term θ_t is the same for each capitalist. Specifically, let it be uniformly distributed over $[0, b(z_t)]$, where z_t is an aggregate state variable. The realization of θ_t is independent across capitalist. The aggregate state z_t follows a Markov process with transition matrix $\pi(z', z)$. Each worker is endowed with one unit of labor and derives no utility from leisure. Capitalists have preferences $E_0 \Sigma_t \beta^t c_t^k$ and workers have preferences $E_0 \Sigma_t \beta^t (c_t^l)^\gamma$, where $\gamma \in (0, 1]$.

At each date, each worker is matched with a capitalist prior to the realization of the aggregate state z and the individual productivity term θ for the capitalist. These matches last only one period, and each period each worker matches with a different capitalist. This rules out the possibility of a long-term contract between a capitalist and a worker. Capitalists have no ability to commit to contracts, i.e., if a capitalist and a worker match on the basis of an agreement to a wage w, the capitalist will fire the worker rather than pay w if $\theta < w$ is realized. In each period, the aggregate state z is revealed only after that period's spot market trades between capitalists and workers have been completed. This assumption, together with the assumption that θ_t cannot be observed by the worker, rules out the possibility of workers and capitalists contracting on a wage w_t that depends on the realization of θ_t or of z_t. Further, this implies that the capitalist will not pay the worker and the worker will not work if $\theta_t < w_t$, giving the worker wage income of zero at t.

For simplicity, assume that there are no asset markets or storage technologies so that each worker simply eats his wages each period, and each capitalist simply eats his profits. This assumption rules out wealth effects that complicate employment contracts. Assume that there are more capitalists than workers, so that the wage offered to workers each period is chosen to maximize the worker's ex ante welfare, or solves

$$\max_{w} \sum_{z'} w^{\gamma}\left(1 - \frac{w}{b(z')}\right)\pi(z', z), \tag{2}$$

where $(1 - w/b(z'))$ is the worker's probability of being hired if the wage is w and z' is the aggregate state realized at t.

The solution to this problem is to set the wage at

$$w(z) = \frac{\gamma}{\gamma + 1}\left(E_z \frac{1}{b(z')}\right)^{-1}, \tag{3}$$

where $E_z[1/b(z')] = \Sigma_{z'}[1/b(z')]\pi(z', z)$. The worker's unconditional probability of being employed is constant each period at $1/(\gamma + 1)$. Nevertheless, the number of workers employed in aggregate state z' is $(1 - w(z)/b(z'))$. The value of aggregate output conditional on transition (z', z) is

$$y(z', z) = \left(1 - \frac{w(z)}{b(z')}\right)\frac{(b(z') + w(z))}{2}. \tag{4}$$

Thus, the rate of unemployment and the level of output fluctuate over time as the aggregate state changes. Note that the wage in the next period rises when the aggregate state today indicates higher expected productivity in the next period and falls when the aggregate state today indicates lower expected productivity in the next period. That is, anticipated changes in worker productivity are reflected in wages. On the other hand, unemployment in the next period rises when the aggregate state tomorrow is lower than its expected value today and likewise falls when the aggregate state tomorrow is higher than its expected value today. That is, unanticipated changes in worker productivity are reflected in changes in the level of unemployment.

Now consider the effects of a government policy that somehow sets $b(z')$ equal to a constant $\bar{b}$, and in particular, let $\bar{b}$ be the expectation of $b(z')$ under its stationary distribution. Under this policy, the wage is set to $w = (\gamma/\gamma + 1)\bar{b}$, the unconditional probability an agent is employed is again $1/(\gamma + 1)$, and aggregate output is constant at $(1/(\gamma + 1))^2(\bar{b}/2)$. Clearly, this policy eliminates aggregate fluctuations. Wages rise on average and the long run expected value of aggregate output and profits fall. Thus, while this countercyclical policy does eliminate the wage risk faced by workers who are employed, the unemployment risk that workers face is unchanged. That is, this policy simply removes the correlation across workers in this risk.[1]

The preceding model is obviously stylized. To what extent does it generalize? One clear deficiency of this simple model is that the unemployment rate is i.i.d. over time. Here we present an extension of the preceding model that predicts serial correlation in the unemployment rate and yet still has the feature that countercyclical policy simply removes the correlation in individual's unemployment risk. Assume now that matches between capitalists and workers can last more than one period, but the productivity of a match θ remains constant as long as that match lasts. Let $\mu(z)$ represent the probability in aggregate state z that an ongoing match continues for at least one more period. Assume that workers who are unemployed enter into new matches requesting wage w. They are employed if the productivity of that new match θ is

1. In this version of the model, it is difficult to consider the welfare implications of countercyclical policy, since the implementation of such a policy changes the division of output between the two types of agents. If one alters the model to assume that workers have an equal ownership share in all firms and, thus, divide aggregate profits between them, it is possible to consider the impact of countercyclical policy on welfare but impossible to derive closed form solutions for wages and unemployment risk due to wealth effects. If these wealth effects are small, then the results on unemployment risk obtained earlier are approximately correct in this altered version of the model. In fact, in every numerical example that we have tried, workers' unemployment risk actually increases when countercyclical policy is implemented. Further, as earlier, wages rise, and expected output and profits fall. The impact of countercyclical policy on welfare is unclear. If agents are very risk averse, the benefits of smoothing profits outweighs the loss in expected output, and, thus, agents gain from countercyclical policy. If agents are not very risk averse, they lose.

greater than w. They remain employed at that wage by the capitalist until that match dies. If a match dies at the end of period t, the worker enters the search pool at the beginning of period $t + 1$ and draws a new match with another capitalist. When z is i.i.d., the solution to this model can be described by the following equations: Let V_s be the beginning of period value of being in the search pool and $V(w)$ be the value of being employed in a match at wage w. Then,

$$V(w) = u(w) + \beta E_z[\mu(z)V(w) + (1 - \mu(z))V_s], \tag{5}$$

and

$$V_s = \max_w E_z\left[\frac{w}{b(z)}\beta V_s + \left(1 - \frac{w}{b(z)}\right)V(w)\right]. \tag{6}$$

The reservation wage w^* is the argmax of the right-hand side of Equation (6) and is constant over time. Note that this is a standard search model except for the fact that the wage a worker receives is his reservation wage rather than his productivity.

In this model, at each date t, individuals who are employed have a higher chance of being employed at $t + 1$ than do individuals who are unemployed at t. Likewise, individuals who are unemployed at t have a higher chance of being unemployed at $t + 1$ than do individuals who are employed at t. This serial correlation in the individuals' employment prospects introduces serial correlation in the aggregate unemployment rate. Thus, even though z is i.i.d., the unemployment rate in this model is serially correlated.[2]

In this model it remains the case that government policies that stabilize $b(z)$ and $\mu(z)$ simply remove the correlation in individuals' risk in transiting from unemployment to employment or vice versa. As before, the probability that an agent in the search pool becomes employed equals $1 - wE_z(1/b(z))$. With quite a bit of algebra, one can show that, again as before, the optimal w for an agent adjusts so that this unconditional probability is constant. Thus, this policy smoothes the unemployment rate simply by removing the correlation in individuals' unemployment risk.

2. For example, when z takes on two values, 1 and 2, with probabilities $\frac{1}{2}$ of each realization, $b(1) = 8$, $b(2) = 12$, $\mu(1) = .95$, $\mu(2) = .99$, $\beta = .98$, and $\gamma = .3$, then the steady-state unemployment rate is .06 and the serial correlation in the unemployment rate is .8.

Our purpose in presenting these models is to demonstrate the possibility that the main effect of countercyclical policy may simply be to eliminate the correlation across individuals in unemployment risk when that risk is determined endogenously as a result of workers' and firms' strategies in the labor market. In these models, workers choose strategies for accepting or rejecting wage offers trading off higher wages for higher risks of being unemployed. Upon implementation of a countercyclical policy, workers' search strategies adjust in such a way to hold constant each individual's unconditional unemployment risk.

From here on, we proceed with the presumption that the business cycle component of the wage risk of the employed and the profits risk of the owners of firms is unimportant in determining the potential welfare gains from countercyclical policy. Instead, we focus solely on the effects of countercyclical policy on individual unemployment risk. We maintain that the effect of such policies is to eliminate the correlation across individuals in the unemployment risk that they face. In the next section, we consider the potential gains from countercyclical policy achieved through their general equilibrium effect on asset prices.

4. *Asset Price Data and the Gains to Eliminating Aggregate Risk*

We have argued earlier that calculation of the welfare costs of aggregate risk requires calculation of the changes in asset prices that result from eliminating aggregate risk. In this section, we use a simple model of asset prices in incomplete markets to argue that what little welfare gains exist from eliminating aggregate risk are due to the elimination of variability in bond returns. With aggregate fluctuations, agents tend to want to borrow when bond returns are high and want to lend when bond returns are low. Without aggregate fluctuations, this correlation between individuals' demands for borrowing and interest rates disappears. For computational reasons, we abandon the model of the previous section and instead study asset prices and aggregate risk in an extremely simple production economy. Nevertheless, this model illustrates the point that the potential welfare gains from countercyclical policies that simply eliminate the correlation in individual risks are small since the observed variation in bond returns is small.

Consider now a two-period model of unemployment risk assuming two asset markets: an uncontingent bond market and a stock market. Here, a share of stock is an asset that pays an aggregate state-contingent dividend with a mean normalized to unity. The dividend is not necessarily equal to aggregate consumption. In this model, we introduce an explicit countercyclical policy based on government hiring whose effect, like the policy in the previous section, is to eliminate the correlation across individuals in the unemployment risk that they face. Government hiring is financed from the sale of the output of those workers employed by the government. We do not attempt to choose parameters governing the pattern of individual unemployment risk to match direct observation of that risk. Instead, we choose the extent of that risk so that the model generates means and standard deviations of log stock and bond returns to match the data. We then calculate the welfare gains from eliminating the correlation across individuals in that risk. Our intention here is to ensure that we begin with a model that is potentially consistent with existing asset price data. In the following section, we remark on some of the more elaborate models of asset prices in incomplete markets currently in the literature. Some of these models take up the question we avoid of whether it is possible to match asset price data in a model with incomplete markets in which the extent of individual income risk is set to match direct observations of that risk.

Our model is an incomplete-markets model of asset prices, individual, and aggregate risk. Let there be a continuum of agents in the model. Let time consist of two dates $t = 1, 2$. Assume that agents produce output and consume at dates $t = 1, 2$ and trade assets at date $t = 1$. At each date, agents either produce high output (are employed) or low output (are unemployed). Their output is denoted by y_t^j, $j \in \{h, l\}$. At each date, a random aggregate state $z_t \in \{B, G\}$ is drawn, with $q(z_t)$ being the probability of z_t. At each date, government follows a policy of hiring $g_t(z_t)$ agents. Agents in government jobs at time t have high output y_t^h. The private sector hires $\pi(z_t)$ agents at time t. Thus, at time t, agents have probability $g_t(z_t) + \pi(z)$ of being employed and producing y_t^h and probability $1 - g_t(z_t) - \pi(z)$ of being unemployed and producing y_t^l. Agents have preferences

$$U = E_0\left\{\frac{c_1^\gamma}{\gamma} + \beta\frac{c_2^\gamma}{\gamma}\right\}. \tag{7}$$

Consider the following asset market structure. After agents learn of their employment status and the aggregate shock z_1 at date $t = 1$, they trade a risk-free bond and stock. The risk-free bond is a sure claim to one unit of consumption at date $t = 2$, and a share of stock is a claim to $d(z_2)$ units of consumption at $t = 2$, where z_2 is the aggregate state at date $t = 2$. Agents are initially endowed with zero bonds and zero shares. Both assets are in zero net supply. Let $p_b(z_1)$ be the price of the bond and $p_s(z_1)$ be the price of the stock given the value of the aggregate shock z_1 realized at date $t = 1$. Let $b^{j_1}(z_1)$ denote the bond holdings at date $t = 1$ of agents of type j_1 given aggregate shock z_1 and $s^{j_1}(z_1)$ denote their corresponding stock holdings. Agents' budget constraints are

$$c_1(z_1, j_1) = y_1^{j_1} - p_b(z_1)b^{j_1}(z_1) - p_s(z_1)s^{j_1}(z_1) \tag{8}$$

$$c_2(z_1, z_2, j_1, j_2) = y_1^{j_2} + b^{j_1}(z_1) + d(z_2)s^{j_1}(z_1). \tag{9}$$

The market clearing condition in the bond market is

$$(\pi(z_1) + g_1(z_1))b^h(z_1) + (1 - \pi(z_1) - g_1(z_1))b^l(z_1) = 0, \tag{10}$$

and for the stock market is

$$(\pi(z_1) + g_1(z_1))s^h(z_1) + (1 - \pi(z_1) - g_1(z_1))s^l(z_1) = 0. \tag{11}$$

To measure the welfare cost of aggregate fluctuations, we calculate equilibrium and consumer welfare in this model first given a constant level of government hiring $g_t(z) = \bar{g}$. We then calculate equilibrium and consumer welfare under a countercyclical government hiring policy $g_t(z)$ that attains the same unconditional mean level of output as in the first policy but that also eliminates all aggregate fluctuations. We assess the welfare cost of aggregate risk as the constant fraction that, if added to agents' consumption stream under the first policy, would give them the same ex ante utility as attained under the second policy.

Certainly, the welfare cost of aggregate risk obtained from this model depends upon the parameters chosen. We have argued that this cost depends upon the extent to which asset prices change when aggregate risk is eliminated. We choose the parameters of this model under the first policy to match data on the mean and standard deviation of asset returns and aggregate consumption growth. Clearly, when aggregate risk is eliminated, both the mean and the standard deviation of asset returns and aggregate consumption growth will change. We take figures of .018 and .033 for the mean and standard deviation of the log of aggregate consumption growth, .06 and .169 for the mean and standard deviation of log stock returns, and .018 and .055 for the mean and standard deviation of log bond returns from Campbell, Lo, and

MacKinlay (1993). As our base case, we choose parameters, $\beta = .615$, $\gamma = -.35$, $y_1^h = 1.039$, $y_1^l = .166$, $y_2^h = 1.060$, $y_2^l = .170$, $\pi(B) = .8075$, $\pi(G) = .9325$, $g_t(z) = \bar{g} = .0625$, $d(B) = .84$, $d(G) = 1.16$, $q(B) = q(G) = .5$. We obtain from the model .020 and .079 as the mean and standard deviation of the log of aggregate consumption growth, .068 and .169 for the mean and standard deviation of log stock returns, and .018 and .051 for the mean and standard deviation of log bond returns. That is, we essentially match the target moments for bond and stock returns as well as mean log consumption growth, while overstating the standard deviation of log consumption growth.

Under the alternative countercyclical policy, we set $g_t(B) = .125$, $g_t(G) = 0$, which sets the rate of unemployment to a constant .0675. Under this policy, the bond and the stock trade at the same constant price.[3] The log of aggregate consumption growth is now constant at .020, and the bond return is .018. The welfare cost of aggregate fluctuations is calculated by finding the number λ such that

$$E_0\left\{\frac{((1+\lambda)c_1)^\gamma}{\gamma} + \beta\frac{((1+\lambda)c_2)^\gamma}{\gamma}\right\} = E_0\left\{\frac{\tilde{c}_1^\gamma}{\gamma} + \beta\frac{\tilde{c}_2^\gamma}{\gamma}\right\}, \tag{12}$$

where $\tilde{c}_t$ represents the agent's consumption at date t under the countercyclical policy. For our base case parameters, we find $\lambda = .0002$, or a welfare gain to eliminating aggregate risk of two one-hundredths of 1% of aggregate consumption. For comparison, the welfare gain to eliminating aggregate risk in this economy under complete markets is 10 times greater. This measure of welfare gains to removing aggregate risk is not sensitive to the choice of β, but of course mean stock and bond returns are.

At the beginning of this section, we argued that the low gain from countercyclical policy in the model was due to the low variability of bond returns found in the data. To see the effect of increasing the standard deviation of bond returns on the welfare costs of business cycles, we perform the following experiment.

To increase the variability of expected asset returns in the economy with aggregate fluctuations, we vary $\pi(B)$, the percentage of agents who are employed given the bad realization of the aggregate shock. This increases the variability of aggregate consumption growth and thus the variability of bond returns. In Figure 1 we plot the welfare gain

3. The stock and bond trade for the same price because without aggregate uncertainty, stock dividends are uncorrelated with individual consumption.

to eliminating aggregate risk against the standard deviation of log bond returns. For comparison, we also include the welfare gain to eliminating aggregate risk given complete markets for the same range of parameter values. The bond variability and corresponding welfare gains for the first example presented earlier are marked "base case" in Figure 1. Note that it is possible to construct examples that generate high costs to business cycles but only by assuming large variability in bond returns.

5. *Asset Pricing in Incomplete Markets*

The two-period economy of the previous section is obviously quite simple. There are several more elaborate models of asset prices in environments with incomplete markets. In this section we discuss two of them, Constantinides and Duffie (1992) and Heaton and Lucas (1992), and their relationship of our simple model to their work. In short, in the previous section, we followed the method outlined in Constantinides and Duffie to build a model in which agents are not exceptionally risk

Figure 1 GAIN TO ELIMINATING AGGREGATE RISK

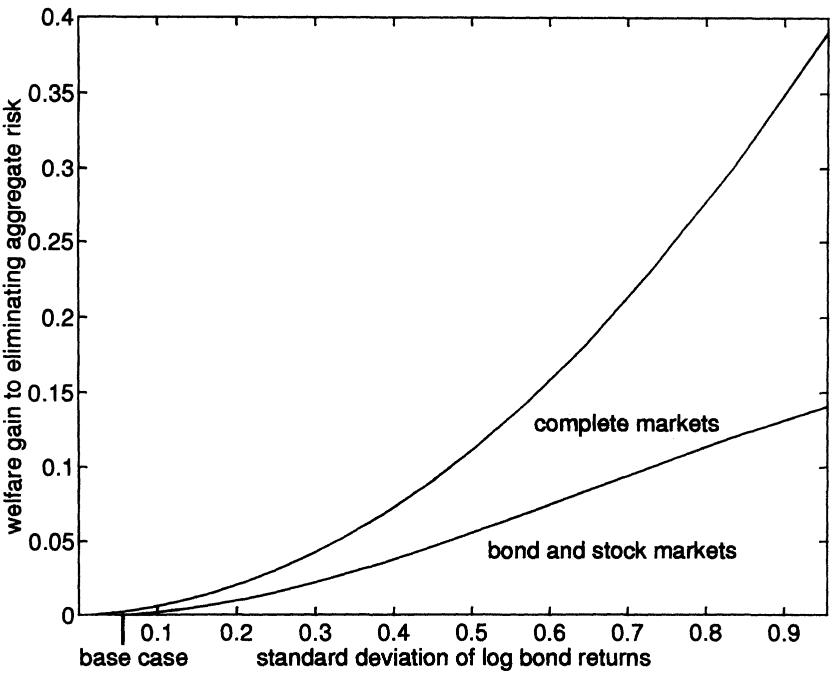

averse that has a large equity premium and relatively smooth aggregate consumption. In constructing that model, we made no effort to use data on individuals' income risk. Heaton and Lucas use data on such risk in constructing their model of asset prices. We discuss the implications of their findings in this section.

As mentioned before, in standard frictionless-trading, consumption-based asset pricing models, asset prices are described by a stochastic process $M_t(z_t)$ known as a pricing kernel that satisfies

$$1 = E_t[R_{i,t+1}M_{t+1}(z_{t+1})], \tag{13}$$

where $R_{i,t+1}$ is the gross return on asset i if held from date t to date $t + 1$, and z_t is an aggregate state of nature. Here, $M_{t+1}(z_{t+1})$ no longer identified with the marginal utility of aggregate consumption for the representative consumer. Instead, the term $M_{t+1}(z_{t+1})$ is identified with $E[\beta(u'(c^i_{t+1})/u'(c^i_t))|z_{t+1}]$, the conditional expectation of individual consumers' marginal utility of consumption for all consumers whose portfolio choices are interior. In such models, the highly variable pricing kernels implied by observed asset price data can be obtained if there is sufficient variability in the conditional variance of individuals' consumption, and, thus, in the conditional expectation of agents' marginal utility of consumption, across aggregate states of nature.

Constantinides and Duffie (1992) demonstrate a procedure for constructing individual consumption sequences to match a wide variety of pricing kernels under the assumption that agents have identical constant absolute risk aversion (CARA) or CRRA preferences. Moreover, this procedure can be used to construct an equilibrium model of a given pricing kernel by endowing each individual consumer with an idiosyncratic income process y^i_t that satisfies

$$M_{t+1}(z_{t+1}) = E\left[\beta \frac{u'(y^i_{t+1})}{u'(y^i_t)} \middle| z_{t+1}\right]. \tag{14}$$

Agents consume their income directly in equilibrium with no asset trade. Of course, it is not necessary in this model that individuals be endowed with their final consumption stream. It is simply the case that computation of the model is much more complicated if asset trades must also be calculated.

In the context of the model economy of Constantinides and Duffie, with frictionless trade of assets, the equity premium implies that the variance of the marginal utility of individual consumption conditional on aggregate state z_{t+1} be both highly variable and highly (negatively) correlated with stock returns. Thus, in this model, the equity premium could be taken as evidence that individual consumers face substantially more risk in recessions than in booms. In fact, it is precisely this principle that we use in constructing our numerical example in the previous section. As we saw in the previous section, though, evidence that individual consumers face substantially more risk in recessions than in booms is not evidence of a high cost of business cycles. If countercyclical policy eliminates aggregate fluctuations by eliminating correlations in individual risk, then the welfare gains from such a policy in our model are likely to be quite small.

As noted earlier, we made no attempt in our two-period model to use data on the income risk faced by individuals to choose parameters. Heaton and Lucas (1992) begin their paper with a review of several papers attempting to use data on individual income variability to calibrate incomplete markets models of asset prices. Finding that these earlier attempts were not successful in generating sizable equity premia, they turn to a study of the role of trading frictions in determining asset prices. We suspect that models that attempt to match asset prices by restricting the trading opportunities of agents with market frictions will deliver lower gains to eliminating correlations in individual risk than reported here. Note again that the gain to eliminating correlations in individual risk under autarky is zero. Assuming trading frictions should move agents closer to autarky.

6. *Conclusion*

In this paper, we measure the potential welfare gains from countercyclical policy in an economy with incomplete markets. In conducting this measurement, we see it as critical to distinguish whether the main effect of countercyclical policy is to reduce the income risk faced by each individual or is simply to reduce the correlation across individuals in the income risk that they face. In the first case, countercyclical policy can have a direct effect on welfare by reducing the risks individuals in the economy face. It also can have an indirect effect on welfare if it changes market-clearing asset prices. In the second situation, in the case in which countercyclical policy simply reduces the correlation in individual risks, the indirect effect of countercyclical policy on asset prices is the only effect such a policy will have.

We present a model in which the effect of countercyclical policy is simply to eliminate the correlation across individuals in the unemployment risk that they face. This model is based on the idea that the unemployment risk that an individual faces is determined in equilibrium by his choice of search strategy in the labor market. In particular, agents trade off a higher reservation wage against a higher probability of remaining unemployed. Once agents' search strategies adjust, countercyclical policy does not reduce any individual's chance of becoming unemployed, it simply ensures that a large number of agents are not unemployed at the same time.

Given this result, we use asset price data to calculate the general equilibrium effects of countercyclical policy on welfare in an incomplete markets environment. We find this effect on welfare to be near zero, since, with incomplete markets, eliminating the correlation in individual income risk does not open up many new opportunities for agents to smooth their consumption. On the other hand, with complete markets, eliminating the correlation in individual income risk allows agents to smooth consumption completely. Thus, the costs of aggregate fluctuations in the incomplete markets economy is smaller than in the complete markets economy for any given technical specification of an economy. Incomplete markets imply a lower gain to eliminating correlations in individual risk in another sense as well. With complete markets, it is difficult to generate an equity premium without assuming a technical specification that itself might imply a high gain to eliminating aggregate fluctuations. With incomplete markets this is not the case.

In this paper, we have not considered the extent to which government might be able to use countercyclical policy to raise the long-run average level of output in the economy and, thus, agents' long-run level of consumption and utility. In our model of the endogenous determination of unemployment risk, countercyclical policy actually lowers the long-run average level of output. If policies that raise long-term average levels of output do exist such policies certainly might produce large welfare gains. De Long and Summers (1988) examine these possibilities. Nevertheless, such gains would have little to do with "smoothing" the business cycle.

In thinking about formulating policies to reduce individual risk, it may be useful to investigate more specific steps government might take to enhance agents' market opportunities for trading risky income for smooth consumption. Deaton and Paxson (1993) and Attanasio and Davis (1993) present micro data that suggests that individuals may face considerable idiosyncratic uncertainty over the long term in their consumption streams. One possible reason that agents might have difficulty insuring idiosyncratic risk is that agents may have limited commitment possibilities. In fact, in our model of endogenous unemployment risk, the assumption of limited commitment plays a key role in preventing risk sharing. The extent to which contracts requiring commitment are enforceable is in many ways determined by government policy.

Appendix: Mathematical Derivation of Endogenous Wage Model

$V(w)$: Value of match paying w.
V_s: Value of being in search pool.
By definition,

$$V(w) = w^{\gamma} + \beta E_z[\mu(z)V(w) + (1 - \mu(z))V_s]. \tag{15}$$

Solving for $V(w)$ delivers

$$V(w) = \frac{1}{1 - \beta\bar{\mu}}(w^{\gamma} + \beta(1 - \bar{\mu})V_s), \tag{16}$$

where $\bar{\mu} = E_z\mu(z)$. Again, by definition,

$$V_s = \max_{w} E_z\left[\frac{w}{b(z)}\beta V_s + \left(1 - \frac{w}{b(z)}\right)V(w)\right], \tag{17}$$

or replacing in for $V(w)$ and collecting terms

$$V_s = \max_{w}\left\{\left[wB\beta + (1 - wB)\frac{\beta(1 - \bar{\mu})}{1 - \beta\bar{\mu}}\right]V_s + \frac{1 - wB}{1 - \beta\bar{\mu}}w^{\gamma}\right\}, \tag{18}$$

where $B = E_z(1/b(z))$. Taking the first-order condition of this and simplifying delivers

$$\beta B\bar{\mu}(1 - \beta)V_s - Bw^{\gamma} + \gamma w^{\gamma-1} - \gamma w^{\gamma}B = 0. \tag{19}$$

If one solves Equation (17) for V_s given an optimal choice of w (with again much simplification), one derives

$$V_s = \frac{1 - wB}{(1 - \beta)(1 - w\beta B\bar{\mu})}w^{\gamma}. \tag{20}$$

Replacing for V_s in the first-order condition (19) and simplifying delivers

$$w^2\left[B^2\beta\bar{\mu}\gamma\right] + w[B(\beta\bar{\mu}(1 - \gamma) - (1 + \gamma))] + \gamma = 0. \tag{21}$$

Note here that if $\bar{\mu} = 0$, Equation (21) implies Equation (3)—the formula for w for the simple one-period-match version of this model if one assumes an i.i.d. distribution for z. For general $\bar{\mu}$, solving for w delivers

$$w = \frac{-(\beta\bar{\mu}(1-\gamma) - (1+\gamma)) + \sqrt{(\beta\bar{\mu}(1-\gamma) - (1+\gamma))^2 - 4\beta\bar{\mu}\gamma^2}}{2\beta\bar{\mu}\gamma B} \tag{22}$$

The unconditional probability that an unemployed agent becomes employed equals wB. From Equation (22), this probability is independent of B; thus, setting $b(z)$ and $\mu(z)$ to their mean values leaves individual unemployment risk unaffected.

REFERENCES

Attanasio, O., and S. J. Davis. (1993). Relative wage movements and the distribution of consumption. Manuscript, Stanford University.

Campbell, J. Y., A. W. Lo, and A. C. MacKinlay. (1993). The econometrics of financial markets. Manuscript, Princeton University.

Cochrane, J. H., and L. P. Hansen. (1992). Asset pricing explorations for macroeconomics. *NBER Macroeconomics Annual* 7:115–165.

Constantinides, G. M., and D. Duffie. (1992). Asset pricing with heterogeneous consumers. Manuscript, University of Chicago.

Deaton, A., and C. Paxson. (1993). Intertemporal choice and inequality. NBER working paper 4328.

De Long, J. B., and L. H. Summers. (1988). How does macroeconomic policy affect output? *Brookings Papers on Economic Activity* 2:433–480.

Heaton, J., and D. Lucas. (1993). Evaluating the effects of incomplete markets on risk sharing and asset pricing. NBER working paper 4249.

İmrohoroğlu, A. (1989). Costs of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy* 97(6):1364–1383.

Lucas, R. E. (1987). *Models of business cycles*. New York: Basil Blackwell.

Comment

JOHN HEATON
Hoover Institution, Massachusetts Institute of Technology, and NBER

Atkeson and Phelan ask us to reconsider the question asked by Lucas (1987): How costly are business cycles? To answer this question we must consider how important are the aggregate or common shocks that individuals face. Further, the extent of available asset markets is likely to affect the answer. In a world of complete markets, individuals are able to trade away all of the idiosyncratic risk that they face, and each individual faces the same aggregate risk. Under some preference assumptions (such as time-additive constant relative risk aversion), each individual consumes a constant fraction of aggregate consumption. In this case the welfare costs of business cycles can be calculated using a representative consumer. Using a representative agent, Lucas (1987) found that the cost of business cycles is small.

There are several possible objections that could be made about Lucas's calculation. First, he assumed that consumption is stationary about a trend. If shocks to aggregate consumption are permanent, then the costs of cycles will be larger. Second, the preferences of the representative agent may be different from those assumed by Lucas. If the representative agent is more risk averse than Lucas assumes, then the costs of cycles will be larger. A related issue is the observation that the equity premium may provide a direct measurement of the market price of business cycles. Since the average equity premium is quite large, the cost of business cycles may in fact be large. A fourth objection is that Lucas's calculations may understate the costs of business cycles because a representative agent model is not appropriate due to market incompleteness. İmrohoroğlu (1989) examined this last objection and calculated the costs of business cycles to be much larger than that obtained by Lucas (1987).

Atkeson and Phelan reexamine the issue of market incompleteness and ask whether the costs of business cycles really are larger in an incomplete markets setting. They consider the effect that a change in aggregate uncertainty may have on individual uncertainty. In a simple model of employment risk, Atkeson and Phelan show that in reaction to a reduction in aggregate risk, individual agents may choose to face more idiosyncratic employment risk. As a result, the welfare gains to reducing aggregate employment risk may not be large. One possible way that the change in aggregate risk may affect welfare is through a general equilibrium effect on prices. To assess this affect, Atkeson and Phelan consider a general equilibrium model of asset markets and show that the gain from reducing aggregate fluctuations in the model is *smaller* in a world of incomplete markets than in a world of complete markets.

I first consider a model similar to Atkeson and Phelan's in which agents have access to limited securities markets and the equilibrium effects of a reduction in aggregate risk can be easily analyzed. In particular, the fact that a reduction in aggregate risk may have no effect on equilibrium welfare can be easily seen. Further, the model can resolve the equity premium puzzle of Mehra and Prescott (1985) so that it captures the appearance of large costs of business cycles. The model is a special case of the model presented in Constantinides and Duffie (1993) and is related to Mankiw's (1986) model.

1. *General Equilibrium Model*

In the model there is a continuum of agents and agent i's income is given by:

$$y_t^i = c_t^a \exp\left(z_t^i + \theta \log c_0^a - \theta \log c_t^a - \alpha t\right) - d_t, \tag{1}$$

where c_t^a is aggregate consumption at time t, and d_t is the dividend due to a holder of a share of stock at time t. z_t^i is given by

$$z_t^i = z_{t-1}^i + \sigma\epsilon_t^i + \left[2\beta + 2\theta \log(c_t^a/c_{t-1}^a)\right]^{1/2}\eta_t^i. \tag{2}$$

In Equation (1) $\alpha = \beta + \sigma^2/2$. The utility of each agent is given by

$$U = \frac{1}{1-\gamma} E\left(\sum_{t=0}^{\infty} c_t^{1-\gamma}\right), \tag{3}$$

where $\gamma > 0$ is the coefficient of relative risk aversion. The shocks $\{\epsilon_t^i\}$ and $\{\eta_t^i\}$ are assumed to be independent of each other and to be independently distributed over time and across agents. Further, for each t these shocks are normally distributed with mean zero and unit variance.

Each agent is endowed with one unit of the stock, and the number (mass) of agents is normalized to be one. Under the assumption of complete markets, there is a representative agent with preferences given by Equation (3). However, suppose that it is not possible for individuals to write contracts directly on their future labor income, and the only assets that can be traded are the stock and a risk-free bond. A no-trade equilibrium exists where the price of the stock and the bond are calculated as if there is a representative agent with CRRA preferences, but the preferences of the representative agent are different from those in Equation (3). Instead, the discount factor of the representative agent is $\beta^* \equiv \beta \exp(\gamma\alpha - \gamma^2\alpha)$, and the coefficient of relative risk aversion is given by $\gamma^* \equiv (1 - \theta - \theta\gamma)\gamma$. Notice that if $\theta < 0$, the equity premium predicted by the model is larger than in the complete markets case. In other words, if the variance of the shocks to individual income is countercyclical, the equity premium is larger.

Suppose now that aggregate consumption growth, $\{\log(c_t^a/c_{t-1}^a)\}$, is i.i.d. over time and takes on the values 1.051 and 0.985 with equal probability. If $\beta^* = 1.09$ and $\gamma^* = 29$, then the model predicts an equity premium of 2% and a bond return of 3%. If the true value of $\gamma = 4$ and $\beta = 0.98$, then under complete markets the equity premium is 0.5%, and the bond return is 8.4%. Notice that for the model with incomplete markets to predict an equity premium of 2% when $\gamma = 4$, θ must be -1.3, which implies substantial countercyclical movement in the variance of shocks to labor income. However, the model does predict a much larger equity premium than the complete markets model and should be a better vehicle for examining the cost of business cycles.

Under the assumption of complete markets, the costs of business cycles is assessed by finding λ (an increase in the growth rate of consumption) such that

$$(1+\lambda)^{1-\gamma}E\{(c_{t+1}^a/c_t^a)^{1-\gamma}\} = \mu^{1-\gamma}, \tag{4}$$

where μ is mean (gross) consumption growth. If $\gamma = 4$, $\lambda = 0.0017$ so that the costs of business cycles are relatively small. This is consistent with the small equity premium predicted by the complete markets model. However, under incomplete markets, we use the *induced preferences* of individual j over aggregate consumption growth:

$$E\{(c_{t+1}^i/c_t^i)^{1-\gamma}\} = E\{(c_{t+1}^a/c_t^a)^{1-\hat{\gamma}}\}. \tag{5}$$

In this case $\hat{\gamma} = 19$ and, analogous to Equation (4), $\hat{\lambda} = 0.0137$. The costs of business cycles are much larger in a world of incomplete markets as in İmrohoroğlu (1989).

However, Atkeson and Phelan find that the cost of business cycles is smaller in an incomplete markets setting. They ask us to consider the effects on individual income risk of a change in aggregate income risk if individuals take on more idiosyncratic risk when aggregate risk is reduced. In the earlier calculation, individual income risk is substantially reduced. Suppose instead that when aggregate shocks are eliminated, the distribution of each individual's labor income *stays the same*, and income is independent across agents. In this case aggregate income growth is constant so that there is a welfare gain under complete markets. However under incomplete markets, an autarkic equilibrium still exists; there is no equity premium, but *individuals are no better off*. As a result, the welfare gains of reducing business cycles are actually larger in a world of complete markets. This means that in an incomplete markets setting, we must be very careful when assessing the effects on welfare of changes in government policy that affect aggregate income risk.

There are several issues that are raised by this type of analysis and this model of asset markets. First, what if the shocks to idiosyncratic income are calibrated using data on individual income? Heaton and Lucas (1994) examine an asset pricing model and calibrate the income shocks using the PSID. The model generates trade in asset markets, but to generate an equity premium, market frictions in the form of transactions costs are needed. The model does generate a "net-of-transaction-costs" equity premium, but it probably is not large enough to generate important gains to eliminating aggregate income variability. However, an important caveat about this model and the model of Constantinides and Duffie (1993) is that they do not exactly match the wealth distribution and the distribution of asset holdings. As a result, welfare calculations based on them must be interpreted with some caution.

A second issue raised by the chapter is, How should we think about research that tries to disentangle aggregate versus idiosyncratic or sectoral shocks? This research may be difficult to interpret given that economic agents may react to changes in aggregate shocks so that the risk they face stays approximately constant. It would have been nice to see some discussion of this literature in the chapter, since determining the exact structure of aggregate shocks versus idiosyncratic shocks is central to the issue of the cost of business cycles.

The essential issue is to determine the likely effect on individual income variability of policy changes that affect aggregate income variability. Atkeson and Phelan provide one example in which a reduction in aggregate variability causes individual agents to choose strategies that result in more idiosyncratic income risk. This result will not generalize to all settings; however, it is an interesting case that needs further consideration. To the extent that the analysis generalizes, costs of business cycles calculated using a representative agent may in fact put an upper bound on the costs.

REFERENCES

Constantinides, G. M., and D. Duffie. (1993). Asset pricing with heterogeneous consumers. Manuscript, University of Chicago.

Heaton, J., and D. Lucas. (1994). Evaluating the effects of incomplete markets on risk sharing and asset pricing. Sloan School of Management working paper.

İmrohoroğlu, A. (1989). Costs of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy* 97:1364–1383.

Lucas, R. E. (1987). *Models of business cycles*. New York: Basil Blackwell.
Mankiw, N. G. (1986). The equity premium and the concentration of aggregate shocks. *Journal of Financial Economics* 17:211–219.
Mehra, R., and E. Prescott (1985). The equity premium: A puzzle. *Journal of Monetary Economics* 15:145–161.

Comment

JAMES TOBIN
Yale University

The Atkeson–Phelan paper is an interesting and ingenious entry in the popular recent literature minimizing the welfare gains of business cycle stabilization. But I hope that the Federal Reserve and the Council of Economic Advisers do not shut up shop just yet.

I begin with a few general remarks.

If nature imposes on Robinson Crusoe a fluctuating income, beyond his capacity to smooth by storage, maybe it is hard to construct a model of his utility that would induce him to pay more than a pittance for a bargain with nature to smooth his harvests. This is not true, of course, if Crusoe faced finite probabilities of famine and starvation, somehow excluded in calculations like those of Atkeson and Phelan. The same is true, with the same exception, of a society of representative consumers facing identical aggregative risks with no idiosyncratic risks. Uncorrelated idiosyncratic risks, without aggregative risk, offer in principle opportunities to smooth everyone's consumption simply by transfers each period from the lucky to the unlucky, given that today's lucky are tomorrow's unlucky. The accomplishment of such transfers is, I guess, what the authors mean by "complete markets." I guess that implies the opportunity to sell one's income stream for a smooth consumption stream, a bargain ruled out in practice by moral hazard and ruled out in law by prohibitions of slavery. It is noteworthy that most practical devices for redistributing incomes to maintain the consumption of those whose wage incomes are interrupted involve government programs, and sometimes entail compulsory membership in order to forestall adverse selection: workmen's compensation, unemployment insurance, social security. The authors' main point is consistent with this observation. They find that the smoothing afforded "in incomplete markets" by

I owe a great deal to discussions of the issues raised by the paper with my colleague William Brainard. He is an old hand at this sort of topic (see Brainard and Dolbear, 1971). He's not responsible for anything I say.

buying and selling bonds and equities are minimal, even with stabilization measures when those measures do not ameliorate individual employment and income prospects.

I approach this paper as an amateur. I am not an aficionado of this literature or of the "real business cycle theory" from which it is a spinoff. Worse yet, I retain the prejudices of my Keynesian generation, which grew up to believe that cyclical recessions and depressions are socially expensive market failures and that diminution of their magnitude is possible and can significantly enhance societal welfare. I have reason to believe that the organizers of this conference were well aware of my disqualifications for this assignment and willfully ignored them. A paper like the one before us has the merit of challenging the endangered species of which I am a member to justify propositions they had regarded as axiomatic.

I shall make two sets of comments. The first and more important concerns what lessons for practical government stabilization policies are suggested by the paper. The paper is, after all, a metaphor. It examines a model economy that is designed to mimic certain features of a real-world economy that bear on the question being investigated but otherwise bear almost no resemblance to actuality. How seriously should the metaphor be taken? The second set concerns some questions and problems of detail I encountered in the specific model of the paper.

1. *On Interpreting and Applying the Model*

This paper, like other papers in the literature, estimates welfare gains from eliminating or diminishing aggregate fluctuations of income and consumption without altering their expected levels and growth rates. The opposing view, my own, is that business cycles are for the most part asymmetrical departures from full-employment equilibrium, as represented for example by "potential" GDP and consumption. In this view the essence of "stabilization" is to raise the mean performance of the economy while also lowering its variance. Consider estimates of the "Okun gap" between potential and actual per capita GNP due to unemployment rates in excess of the lowest inflation-safe unemployment rate, if you like, the "NAIRU" or natural rate. The average gap from 1946 to 1992 in percent of potential per capita GNP is 2.1 if negative gaps are counted or 2.7 if they are counted as zeros. In nonmacro contexts no one would regard losses of this magnitude, \$136 to \$175 billion currently, as peanuts. Gaps of this kind were much larger before World War II, and it is quite possible they would have been larger these past five decades in the absence of countercyclical policies, both by "built-in stabilizers" in fiscal and financial institutions and by discretionary fiscal and monetary demand management (DeLong and Summers, 1988). Furthermore, it may well be that the growth of potential GNP itself is weakened by large and prolonged departures from full-employment potential, because they adversely and irreversibly affect aggregate saving and investment and the human capital embodied in the work force.

I realize that these Keynesian possibilities are ruled out in the paper by the equilibrium setup of the model economy. In the land of Atkeson and Phelan, cyclical fluctuations result from exogenous shocks to labor productivity, certainly not from aggregate demand shocks. Markets never fail to clear, even momentarily. If there is money in their economy and if prices are quoted in it, it plays no role at all in the fluctuations of real variables. Here "stabilization policy" does not mean what it usually connotes, demand management by fiscal and monetary policies. It would be a misuse of the model and its numerical conclusions to abandon conventional demand management unless you subscribe to the moving-equilibrium view of real business cycle theory.

It is true that the authors' model economy does concern fluctuations in employment and unemployment, and it does envisage compensatory public employment as the instrument of aggregative stabilization. Somehow employers and workers agree on (real) wages each period before they know the random draw of productivity growth. After the draw employers fire workers whose productivity is less than the agreed wage. Compensatory public employment is an inverse function of the productivity shock. Competition among employers sets the wage at the level maximizing workers' expected utility, as workers trade off wages against the probabilities of unemployment. Thus, unemployment is endogenous and voluntary. Moreover, Atkeson and Phelan contrive to make each individual's unemployment prospects independent of the existence and strength of the public jobs program.

Thus, in the model the only consequence of countercyclical policy is aggregate stabilization, i.e., in the authors' words, diminishing or eliminating correlations among individuals in their income fluctuations. This accomplishes very little, because it limits the source of welfare gain to the possibility that the policy makes it easier for individuals to smooth their consumption streams by trading assets among themselves.

At the end of their paper, Atkeson and Phelan do recognize the logical possibility that stabilization might raise mean incomes and consumptions, although their model excludes it. That would require complementary interaction between aggregate stabilization and improvement in individual prospects, a more likely possibility in the real world than in their world. Observation suggests that the impacts of aggregate fluctuations are concentrated on a minority of the population and labor force, indeed on persons with the least access to asset markets that permit consumption smoothing and hedging against unemployment risks. Tenure professors, professionals, and salaried white-collar workers in general are much less troubled by personal cycles, much less vulnerable to aggregate cycles, and much better situated to use asset markets. (True, in our most recent business cycle, these groups were suddenly more vulnerable than before.) The consequence is that aggregate stabilization is of particular benefit to the vulnerable minority. Their personal expectations of intermittent unemployment are greatly reduced, not just rendered heterogeneous in timing. Atkeson and Phelan do not consider this possibility.

Work sharing to spread the impacts of reductions in demand for person-hours of labor would limit the damage of aggregate unemployment and diminish the welfare gains from stabilization. But unless job sharing included the higher echelons, the basic point remains that aggregate stabilization diminishes the personal prospects of unemployment.

2. *Questions about the Model*

My main complaint is the absence of any connection between the labor market described in Section 3 and the two-period economy with two asset markets and two states of nature in Section 4. For one thing, the utility function in Section 3, which governs wages, unemployment, and profits is not the same as the utility function of Section 4, which determines saving and dissaving in bonds and stocks. The purpose of Section 4 is to give numerical content to a $2 \times 2 \times 2$ model by assigning "realistic" parameter values—to incomes of employed and unemployed and stock dividends in high- and low-productivity states, to time preference, to relative risk aversion. The criterion of realism is that these numbers imply values of endogenous variables—means and standard deviations of aggregate consumption growth, stock returns, and bond (actually bill) returns. The calibration is not altogether successful, partly because a two-year economy cannot mimic an ongoing infinite economy.

Specifically, the time preference discount is outlandishly high if the two periods are meant to be years. If they are meant to be decades or quarter centuries, then calibration of other parameter values by real-world annual values is not proper.

In the two-period model, the "stock" pays just one simple dividend, higher in the good state than in the bad state. Thus, there is no counterpart of the capital-gain component of total return taken as the stylized fact.

The most serious complaint is that the "stock" in the model is a completely arbitrary and artificial construct. The authors make no effort to identify it with claims on the profits generated in the wage-employment model of Section 3. (Positive profits are generated because risk-averse workers settle on wages systematically lower than productivity in order to lower unemployment risk.) Indeed, the Section 4 model appears to involve no identity that tells us that wages plus dividends equal production on each date in each state of nature. Likewise, the government budget is not modeled. How are the jobs programs paid for? Evidently, neither by debt issues nor by taxes.

Although the authors stress the possible role of asset prices in contributing to consumption smoothing, they do not tell us the asset prices generated by their model, with and without government countercyclical job creation, and in the two possible initial-period states. One reason for interest in these numbers is to see whether interest rates are procyclical or countercyclical. They might be higher in a low-productivity initial state, reflecting the higher marginal utility of current consumption. This would be a mark against the credibility of real business cycle models that interpret fluctuations as intertemporal substitutions, as against demand-side disequilibrium models, which imply low interest rates in recessions because investments are constrained by current and expected demand rather than by saving.

REFERENCES

Brainard, W. C., and F. T. Dolbear. (1971). Social risk and financial markets. *American Economic Review* LXI: 360–370.
DeLong, J. B., and L. H. Summers. (1988). How does macroeconomic policy affect output? *Brookings Papers on Economic Activity* 2:433–480.

Discussion

Michael Woodford pointed out that the crucial assumption in the model seemed to be a natural rate hypothesis and not the assumption that productivity shocks are driving aggregate fluctuations. Under a natural rate assumption, the model's welfare implications are relevant, regardless of the source of the shocks driving the aggregate fluctuations. Atkeson agreed with this interpretation and said that it was primarily out of convenience that only productivity shocks were included.

Peter Diamond noted that knowing the costs associated with risk aversion and individual unemployment experiences was important not just for stabilization policy but also for public finance issues such as choosing parameters for unemployment insurance. He cited work by Jonathan Gruber, who has used the PSID to look at the effect of unemployment insurance on people's consumption, and suggested that this work could be used to calibrate the individual fluctuations in Atkeson and Phelan's model.

Diamond also asked why employment and unemployment were modeled as single-period, independent events instead of a more common hazard approach, where the probability of employment conditional on being employed would be much higher than conditional on being unemployed. If individuals' unemployment risks were allowed to compound from period to period, then the fluctuations in consumption would tend to be higher. He also noted that if long-term unemployment spells lowered the probability of becoming employed by more than short spells, then there would be room for stabilization policy to improve welfare by reducing the probability of long-term unemployment, even if there was no effect on the mean level of unemployment. Diamond added that more attention needed to be focused on characterizing the distribution of the shocks since individual decisions will be very sensitive to this distribution.

In response to Diamond, Atkeson said that the optimization problem faced by workers would be much more complicated if unemployment tended to persist. However, he suspected that even in more complicated settings, the endogeneity of the wages and, thus, of the unemployment probabilities would yield similar results.

Olivier Blanchard suggested that an interesting empirical exercise to test the implications of the model would be to compute the unconditional hazard rates in and out of unemployment and employment in the absence of any aggregate fluctuations. Blanchard said that a simple way to do this would be to take the work Steve Davis and others have done in identifying reallocative versus aggregate shocks and to estimate hazard functions conditional on setting the aggregate shocks to zero. He guessed that the results would be consistent with the model, that there would still be a substantial amount of reallocation in the economy, and that the unconditional probabilities might not change very much.

Herschel Grossman asked why agents in the model were always induced to choose exactly the same probability of being unemployed. Intuitively, if the fluctuations in people's marginal products were reduced, then workers would be likely to choose some combination of a higher wage and a lower probability of being unemployed. Atkeson responded that the model was cooked to deliver this result, but that the basic idea was that when workers choose the wage, they don't know what the aggregate realization is, and so they end up picking an unconditional probability of unemployment.

Miles Kimball suggested allowing for a separation between risk aversion and intertemporal substitution. In such a specification, the costs of fluctuations would be more sensitive to low intertemporal substitution than they are to high-risk aversion. Kimball also suggested that one way to look at the effect of aggregate fluctuations on the average level of employment would be to see if the unemployment rate is convex in various business cycle indicators. If it is convex, then the fluctuations will tend to increase the average unemployment rate.

Daron Acemoglu observed that even if there are no obvious market failures preventing agents from trading with each other, there may be important negative search externalities in periods of high unemployment that would leave room for stabilization policy to raise the aggregate level of activity.

David Romer said that focusing on consumption-based utility as Lucas did was misleading. He suggested that a more realistic specification would include hours worked, with perhaps some asymmetry to capture the intuition that the utility loss from working less in a recession was not equal to the utility gain from working less during recessions.

Randall Wright asked why the focus was on business cycles rather than other fluctuations such as seasonal or weekly fluctuations. Atkeson answered that the paper was not meant to address business cycles per se, but rather aggregate risk.