NBER WORKING PAPER SERIES

RECONSIDERING THE COSTS OF BUSINESS CYCLES WITH INCOMPLETE MARKETS

Andrew Atkeson
Christopher Phelan

Working Paper No. 4719

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 1994

This paper is part of NBER's research program in Economic Fluctuations. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #4719
April 1994

RECONSIDERING THE COSTS OF BUSINESS CYCLES WITH INCOMPLETE MARKETS

ABSTRACT

In this paper, we measure the potential welfare gains from counter-cyclical 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 counter-cyclical 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 counter-cyclical policies? In taking up the first question, we see it as critical to distinguish whether the main effect of counter-cyclical policy is to directly reduce the income risk faced by each individual or is simply to reduce the correlation across individuals in the income risk that they face. 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 counter-cyclical policy aimed at reducing aggregate fluctuations may be simply to remove the correlation across individuals in the unemployment risk that they face. We then use asset price data to argue that in an incomplete markets framework, the potential welfare gains from counter-cyclical policy are close to zero.

Andrew Atkeson
Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
and NBER

Christopher Phelan
Department of Economics
University of Wisconsin-Madison
1180 Observatory Drive
Madison, WI 53706

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 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

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

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 the $\log(z_t)$ 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' 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 counter-cyclical 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 counter-cyclical 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 counter-cyclical 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 counter-cyclical policies? In taking up the first question, we see it as critical to distinguish whether the main effect of counter-cyclical policy is to directly 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 either situation, counter-cyclical policy will have a general equilibrium effect on welfare if it changes asset prices. However, in the second situation, the situation in which counter-cyclical 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 counter-cyclical 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 to 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 which focuses on unemployment risk as the individual risk that would be affected by counter-cyclical 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 which 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 counter-cyclical policy in an economy with incomplete markets. In the first of these sections, we discuss theoretically the different effects on welfare of counter-cyclical policies which reduce aggregate fluctuations by reducing individual income risk directly and counter-cyclical policies which 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 counter-cyclical 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 counter-cyclical 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 counter-cyclical policy are essentially zero.

2. Aggregate and Individual Risk in Incomplete Markets

Under the assumption that there are complete markets for insuring 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 counter-cyclical policy on the processes which generate individual income streams. All that matters is the effect of counter-cyclical 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 calculating the welfare gains from counter-cyclical 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 counter-cyclical 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 counter-cyclical 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 above, 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, counter-cyclical 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, is unchanged.

It is clear, then, that if the main effect of counter-cyclical policy is to remove correlations in individual risk, then 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' 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 which clear markets when aggregate consumption is constant as opposed to facing asset prices which 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' model are small. As a result, he finds a low cost of aggregate fluctuations.

Given this interpretation of Lucas' result, one 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 which 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 which smoothed consumption and that marginal utility over the cycle could improve the utility of the representative consumer substantially.

The equity premium in Lucas' 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, that 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 which fail to predict a large equity premium may also fail to measure accurately the volatility of marginal utility and thus understate the welfare costs of business cycles, we design our model for measuring the potential benefits of counter-cyclical policy to be consistent with the equity premium and other data on the volatility of asset returns.

This idea that counter-cyclical 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 which determines an individuals' 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 which 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 above. 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 that an alternative method for eliminating aggregate fluctuations in her model is to maintain her four-state Markov process describing individuals' labor market transitions but eliminate correlations across individuals in these transitions. That is, eliminate the correlation across individuals in the stability of their individual employment and unemployment. While both methods eliminate aggregate fluctuations, since prices are pinned down by the assumed storage technology, the welfare gain to simply eliminating correlations in employment stability across individuals is identically zero.

3. Counter-Cyclical Policy and Endogenous Unemployment Risk

We now consider the question of whether the main effect of counter-cyclical 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 counter-cyclical policy which 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 counter-cyclical 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 which 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 capitalists. 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 $E_0 \sum_t \beta^t c_t^k$ and workers have preferences $E_0 \sum_t \beta^t (c_t^l)^\gamma$, 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 capitalists. This rules out the possibility of long term contracts between a capitalist and a worker. Capitalists have no ability to commit to contracts; that is, 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

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

where $(1 - \frac{w}{b(z')})$ 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

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

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

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

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 $w = \frac{\gamma}{\gamma+1}\bar{b}$, the unconditional probability an agent is employed is again $\frac{1}{\gamma+1}$, and aggregate output is constant at $(\frac{1}{\gamma+1})^2\frac{\bar{b}}{2}$. 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 counter-cyclical 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]

[1] In this version of the model, it is difficult to consider the welfare implications of counter-cyclical 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 counter-cyclical policy on welfare, but impossible to derive closed form solutions for wages and unemployment risk due to wealth effects. If these wealth effects are

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 which predicts serial correlation in the unemployment rate and yet still has the feature that counter-cyclical 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 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 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

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

---

small, then the results on unemployment risk obtained above 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 counter-cyclical policy is implemented. Further, as above, wages rise and expected output and profits fall. The impact of counter-cyclical 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 counter-cyclical policy. If agents are not very risk averse, they lose.

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 which 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. One leaves unaffected his unconditional transition probabilities. Thus this policy smoothes the unemployment rate simply by removing the correlation in individuals' unemployment risk.

Our purpose in presenting these models is to demonstrate the possibility that the main effect of counter-cyclical 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

[2] For example, when z takes on two values, 1 and 2, with probabilities 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.

off higher wages for higher risks of being unemployed. Upon implementation of a counter-cyclical 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 counter-cyclical policy. Instead, we focus solely on the effects of counter-cyclical 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 counter-cyclical policy achieved through their general equilibrium effect on asset prices.

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

We have argued above 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 counter-cyclical policies which 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 which 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 counter-cyclical 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 the output of those workers employed by the government. In choosing parameters for this model, 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 which 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 $q(z)$ being the probability of z. At each date, government follows a policy of hiring $g_t(z)$ agents. Agents in government jobs at time t have high output y_t^h.

Agents not in government jobs at time t have probability $\pi(z)$ of having high output and probability $(1-\pi(z))$ of having low output. Thus at time t, agents have probability $g_t(z)+\pi(z)$ of being employed and producing y_t^h and probability $1-g_t(z)-\pi(z)$ of being unemployed and producing y_t^l. Agents have preferences

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

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 given

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

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

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 counter-cyclical 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 counter-cyclical 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

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

where $\tilde{c}_t$ represents the agent's consumption at date t under the counter-cyclical policv. 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 ten 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 counter-cyclical 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 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

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

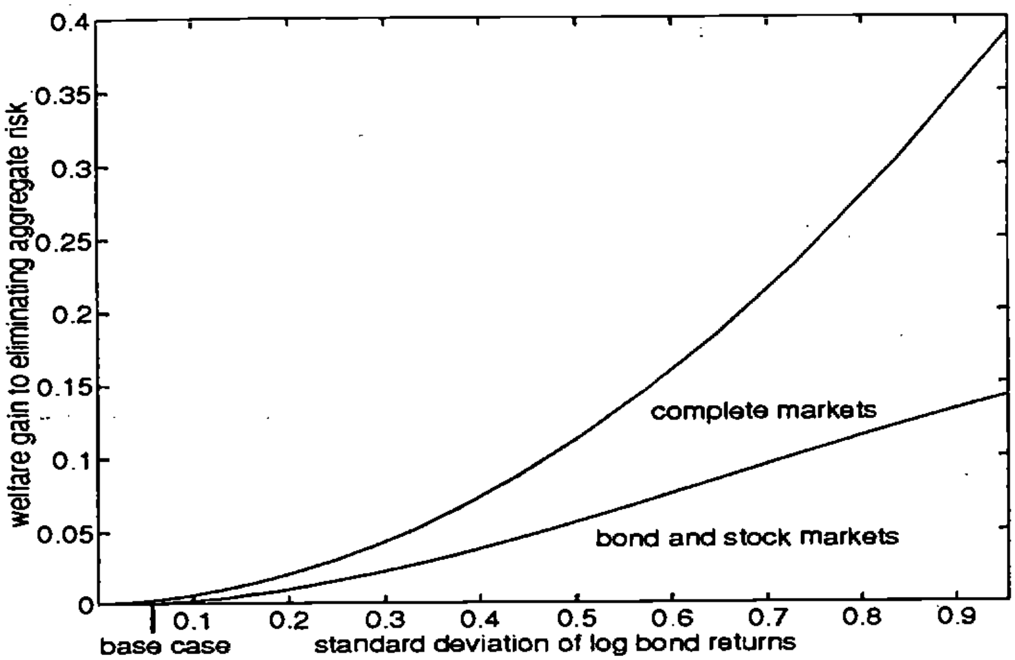

Figure 1. Gain to Eliminating Aggregate Risk

presented above are marked "base case" in Figure 1. Note that it is possible to construct examples which 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 John Heaton and Deborah 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 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 $E[\beta \frac{u'(c_{t+1}^i)}{u'(c_t^i)} \mid z_{t+1}]$, 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 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_t^i that satisfies

$$M_{t+1}(z_{t+1}) = E[\beta \frac{u'(y_{t+1}^i)}{u'(y_t^i)} \mid z_{t+1}]. \tag{14}$$

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 counter-cyclical 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 which 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 counter-cyclical policy in an economy with incomplete markets. In conducting this measurement, we see it as critical to distinguish whether the main effect of counter-cyclical 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, counter-cyclical 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 counter-cyclical policy simply reduces the correlation in individual risks, the indirect effect of counter-cyclical policy on asset prices is the only effect such a policy will have.

We present a model where the effect of counter-cyclical 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, counter-cyclical 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 counter-cyclical 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 which 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 counter-cyclical 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, counter-cyclical policy actually lowers the long-run average level of output. If policies which 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.

References

Attanasio, Orazio and Steven J. Davis (1993) "Relative Wage Movements and the Distribution of Consumption," manuscript, Stanford University.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay (1993) "The Econometrics of Financial Markets," manuscript, Princeton University.

Cochrane, John H. and Lars Peter Hansen (1992) "Asset Pricing Explorations for Macroeconomics," *NBER Macroeconomics Annual*, **7**, 115–165.

Constantinides, George M. and Darrell Duffie (1992) "Asset Pricing with Heterogeneous Consumers," manuscript, University of Chicago.

Deaton, Angus and Christina Paxson (1993) "Intertemporal Choice and Inequality," NBER working paper 4328.

De Long, J. Bradford, and Lawrence H. Summers (1988) "How Does Macroeconomic Policy Affect Output?," *Brookings Papers on Economic Activity*, **2**, 433–480.

Heaton, John and Deborah Lucas (1993) "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," NBER working paper 4249.

İmrohoroğlu, Ayşe (1989) "Costs of Business Cycles with Indivisibilities and Liquidity Constraints," *Journal of Political Economy*, **97(6)**, 1364–1383.

Lucas, Robert E. (1987) *Models of Business Cycles*. New York: Basil Blackwell.

7. 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,

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

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

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

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 (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

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

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

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

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