Review of Economic Dynamics **2,** 245–272 (1999)
Article ID redy.1998.0043, available online at http://www.idealibrary.com on IDEAL®

On the Welfare Effects of Eliminating Business Cycles*

Per Krusell†

Department of Economics, Harkness Hall, University of Rochester, Rochester, New York 14627

and

Anthony A. Smith, Jr.

Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213

Received May 1, 1998

We investigate the welfare effects of eliminating business cycles in a model with substantial consumer heterogeneity. The heterogeneity arises from uninsurable and idiosyncratic uncertainty in preferences and employment, where, regarding employment, we distinguish among employment and short- and long-term unemployment. We calibrate the model to match the distribution of wealth in U.S. data and features of transitions between employment and unemployment. Unlike previous studies, we study how business cycles affect different groups of consumers. We conclude that the cost of cycles is small for almost all groups and, indeed, is negative for some. *Journal of Economic Literature* Classification Numbers: C68, D31, D61, E32.

*This project originated in conversations with Rao Aiyagari. For valuable comments, we wish to thank Neil Wallace, Zvi Hercowitz, an anonymous referee, and seminar participants at Carnegie Mellon, the Institute for International Economic Studies, Penn State, Rochester, Toulouse, the 1998 Annual Meetings of the Society for Economic Dynamics, and the Aiyagari Memorial Conference on Dynamic Macroeconomics. Per Krusell thanks the Bank of Sweden Tercentenary Foundation and both authors thank the National Science Foundation for support.

†Per Krusell is also affiliated with the Institute for International Economic Studies, Stockholm University, Stockholm, Sweden, and the Center for Economic Policy Research, London, England.

1094-2025/99 $30.00

1. INTRODUCTION

In a provocative exercise, Lucas (1987) calculated an estimate of the welfare gain associated with the elimination of business cycles. Lucas's calculation was very simple. He translated the comparison between an economy with and without cycles into a comparison between an estimated time series representation of the actual postwar U.S. consumption path and the trend part of that representation. To obtain a welfare comparison, Lucas assumed an infinitely lived agent who maximizes expected utility and has constant relative risk aversion. The estimates implied welfare gains, translated into equivalent changes in average aggregate consumption, of no more than a very small fraction of 1%; for example, for logarithmic utility the welfare gain is 0.008%.

If one wants to claim that Lucas' estimate badly understates the possible gains, then there would seem to be three alternative routes to take. First, one can dispute his assumption that eliminating cycles leaves the trend level of output unchanged; perhaps instead it is possible to raise the average level of output by eliminating cycles. The second and third routes accept the notion that average output in some sense is unaffected. The second route stays within Lucas' general framework, but argues that other assumptions about preferences or about the stochastic process governing aggregate consumption under cycles are more realistic and lead to larger costs. The third route is to look at the effects of eliminating cycles in a more disaggregated fashion. In particular, one can study the effects of business cycles on different consumers in order to investigate whether cycles seem much more costly to some consumers than to others, a possibility Lucas mentioned. This paper is one, among a few others, that takes the third route.

The perspective we offer here is that it is quite plausible that the welfare costs of cycles are not so high on average, but may be very high for, say, the very poor or currently unemployed members of society. We therefore compare cyclical and noncyclical economies from the perspective of individual consumers as a function of their wealth and employment status. Our analysis is the first one that provides such comparisons: whereas some existing heterogeneous-agent analyses have asked whether the welfare costs of cycles may be higher on average across consumers in such models, these analyses have not studied the effects on subgroups of consumers.

We employ a dynamic equilibrium model where consumers differ in a number of respects: employment status and preferences (discount rates), which both are exogenous and stochastic and follow a process common to all consumers, and wealth, which is endogenous. Although consumers cannot insure themselves directly against idiosyncratic risks, consumers can save, and their savings can be used as a buffer to insure partially against adverse idiosyncratic outcomes. Our model is calibrated; in particular, we match U.S. employment and wealth data. Consequently, consumers differ widely in both their wealth holdings and their employment prospects. As a result, insurance possibilities and exposures to risk vary substantially in the population.

The economy with cycles is driven by exogenous stochastic movements in productivity and employment. We construct a corresponding no-cycle economy by replacing the aggregate shocks with their conditional expectations and by integrating the idiosyncratic shock processes with respect to the aggregate stochastic variables. We remove aggregate shocks as of a given point in time, solve for the equilibrium transition path toward the steady state without aggregate movements, and compare the welfare of each individual in this equilibrium to the welfare he would have obtained had the aggregate shocks remained.

Our main quantitative finding is that the costs of business cycles in our framework are extremely small for almost all consumers; they are even negative for some consumers. This finding hinges on two crucial features of our analysis. First, the elimination of stochastic aggregate movements in employment (in particular, in the economy with cycles, an individual is more likely to be unemployed in recessions than in booms) does not alter individual consumers' employment processes directly. Why? This follows from how we eliminate cycles. A simple 2-by-2 example explains this point: an agent's probability of employment is the same in a four-state process (employed/unemployed, good/bad aggregate state) as in the two-state process arrived at by summing the probabilities of employment in the two aggregate states. This example extends to serially correlated processes. Second, given that there are no direct "risk-reducing" effects on employment, only price effects—fluctuations in wage and rental rates—remain. In our economy, these fluctuations do not affect welfare much, and some consumers may even prefer them. One reason for the small cost of risk is that most consumers—dynasties in our case—save enough so that the remaining risk exposure is unimportant in utility terms.

The welfare effects of eliminating cycles do differ across consumers. Consumers who have no wealth, are up against a borrowing constraint, and are unemployed can suffer large losses from cycles—up to one or two percentage points in average consumption. However, there are vanishingly few consumers in this situation, since consumers take precautions ex ante and save enough to avoid it. We also find that rich consumers may gain considerably from eliminating cycles.

The first paper to introduce consumer heterogeneity for the purpose of studying the effects of eliminating business cycles was İmrohoroğlu (1989). Her framework is similar to ours, but has exogenous, nonfluctuating prices and does not allow a realistic calibration of the wealth distribution. Other papers that investigate the role of heterogeneity include Atkeson and Phelan (1994), whose work we discuss in detail in Section 2.3, Beaudry and Pages (1996), and Gomes, Greenwood, and Rebelo (1998). Atkeson and Phelan discussed the connection between aggregate and idiosyncratic risk, and they suggested as a serious possibility that the elimination of aggregate risk does not affect individual risk at all. They did not analyze a calibrated dynamic model, but focused on simple examples; one of these makes the point that an economy with a high market price of aggregate risk does not necessarily produce large welfare gains when this risk is eliminated.

Beaudry and Pages (1996) studied idiosyncratic wage risk that worsens in recessions: so-called reallocation shocks.[1] They assumed that when laid-off workers are reemployed, their new wages are much lower than their wages were before they were laid off, and that this wage difference disappears only slowly over time.[2] Since layoffs occur more frequently during recessions, they argued that cycles lead to an increase in both the variance and the persistence of idiosyncratic risk. Beaudry and Pages obtained higher costs of business cycles. However, their findings are based on the assumption that there is no idiosyncratic risk at all in the economy without cycles, and that workers in the economy with cycles cannot save to insure against wage risk.

Gomes, Greenwood, and Rebelo (1998) argued—as we do, but for a different reason—that the elimination of cycles may increase utility for many agents. The argument in that paper, whose main purpose is to study search unemployment in a context with incomplete markets against idiosyncratic risks, is based on the option value of search. Since low outcomes are not payoff-relevant, the search behavior results in payoffs which are convex in productivity (wage), so that more fluctuations in productivity may be preferred to less. As a final point on the literature, there is a view, expressed in Obstfeld (1994) and later in Tallarini (1997) and Dolmas (1998), that agents' preferences may be of the non-expected-utility type. Especially accompanied with low discount rates and aggregate time series that are well approximated by a random walk, such preferences can lead to significantly larger welfare costs of aggregate fluctuations.

Because the effects of eliminating cycles are quite complex in our general setup, we first describe, in Section 2, partial equilibrium effects in a two-period model. We then present the full model (specification in Section 3 and results in Section 4) along with its calibration and computational findings. We study two versions of the full model: one where the employment process is of the standard, two-state variety, and one where in addition there is a distinction between short- and long-term unemployment.

[1]The effects of reallocation shocks of the uninsurable kind are also considered in Attanasio and Davis (1996).

[2]Empirical support for this can be found in Bils (1985) and elsewhere.

2. PRELIMINARIES

We first briefly lay out in Section 2.1 a slightly simplified version of our general model. We then present in Section 2.2 a two-period model which is constructed to capture—in essence and notation—most of the ingredients in the multiperiod model. Quantitative issues and the effects of transition, along with other complications due to an infinite time horizon, are covered later in the paper. In Section 2.3 we discuss in detail how we eliminate cycles; Section 2.4 uses the two-period model to analyze the welfare consequences.

2.1. A Dynamic Model

We describe a dynamic model which is close to the one we study quantitatively in Section 3. For presentational purposes, the model in this section is slightly simpler: it has two employment states only—employed and unemployed—and no preference heterogeneity.

We use a Bewley-style model, similarly to Aiyagari (1994) and Huggett (1993), with aggregate uncertainty. In particular, the setup builds on the one studied in Krusell and Smith (1998). There is a large number (measure 1) of ex ante identical agents. Preferences are given by

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t),$$

where β^t is the discount rate between time 0 and time t and u is strictly increasing and strictly concave.

There is an exogenous aggregate shock, z: $z = z_g$ in good times and $z = z_b$ in bad times. This process follows a first-order Markov process, with $\pi_{z'|z}$ denoting the probability that next period's state is z' if the current state is z (primes are used for next period variables throughout).

Consumption in this economy derives from two sources: a constant returns to scale, Cobb–Douglas production function whose inputs are total capital and total labor input, and home production. The aggregate production function is

$$z\,\bar{k}^{\alpha}\bar{n}^{1-\alpha},$$

where k is capital (an overbar refers to a total) and n is labor. Home production, which accrues in the amount g to all unemployed agents, is a simple way of capturing a basic, exogenous level of insurance against employment shocks.[3] Aggregate output, including undepreciated capital, can be used to either consume or invest.

[3]This insurance could be thought of as unemployment insurance; incorporating a government budget constraint to reflect this is easy and would not change our analysis.

An agent's working status is described by ϵ: the agent either works, $\epsilon = 1$, or is unemployed, $\epsilon = 0$. When employed, each agent supplies one unit of labor input. Therefore, $\bar{n}$ equals $1 - u$, where u is the unemployment rate. We allow the unemployment rate to take on only two values: u_g in good times and u_b in bad times. That is, u and z move together perfectly (although in opposite directions). We employ a law of large numbers so that, conditional on the aggregate state, agents' employment statuses are uncorrelated. The individual employment status follows a first-order Markov chain. Notice that the restriction of the unemployment rate to two values forces the individual transition probabilities to depend on both today's and next period's aggregate states. We use $\pi_{\epsilon'|\epsilon z z'}$ to denote the probability of ϵ' conditionally on (ϵ, z, z'); $\pi_{\epsilon', z'|\epsilon z}$ refers to the joint (ϵ', z') outcome next period.

The markets in this economy are simple: labor and capital services are traded on competitive spot markets each period, at marginal product prices $w = w(\bar{k}, 1 - u, z)$ and $r = r(\bar{k}, 1 - u, z)$ in terms of current consumption goods, respectively.

We rule out insurance markets for idiosyncratic risk by assumption.[4] In addition, we assume that only one asset is traded. This asset is a claim to one unit of capital and, since capital is exchangeable for consumption one-to-one, it has the price 1. Its return next period is $1 - \delta + r$, and $r = z(\bar{k}/(1-u))^{\alpha-1}$, so the asset is risky. For simplicity, we do not include a riskless asset; in an earlier paper (Krussel and Smith, 1998), we studied a very similar model with a second, riskless asset, and we found that the allocations were very similar whether or not the second asset was present.

There is also a time- and state-independent lower bound, $\underline{k}$, on any agent's holdings of this asset. This lower bound—a borrowing constraint—precludes perfect insurance using the asset. Market clearing for assets means that agents' capital holdings sum up to the economy's total capital stock.

We defer a formal definition of equilibrium to Section 3. Suffice it to say here that agents choose consumption and savings each period subject to their budget constraint,

$$c + k' = r(\bar{k}, 1 - u_z, z)k + w(\bar{k}, 1 - u_z, z)\epsilon + g(1 - \epsilon) + (1 - \delta)k,$$

and their borrowing constraint so as to maximize their net present-value utility; they take all aggregate variables as given.

[4]Cole and Kocherlakota (1997) show that, under certain assumptions about the unobservability of shocks and behavioral variables, the one-asset allocation does implement a constrained optimum in the case of no aggregate shocks, provided that the borrowing constraint is the loosest possible ensuring that any debt is repaid.

2.2. A Two-Period Model

The model we now consider is simply a two-period version of the model just described. This means that the second period has no savings decisions; all income is consumed. For cleaner exposition, we will suppress any notation reflecting the current stochastic states, so that, for example, π_g is the probability that the second-period aggregate state is g, and we will lump together all current income of an agent into the variable ω.

The first-period budget thus is

$$c + k' = \omega,$$

whereas the second-period budgets satisfy, for each realization of the individual and aggregate stochastic states

$$c'_{0g} = k'R'_g + g',$$

$$c'_{0b} = k'R'_b + g',$$

$$c'_{1g} = w'_g + k'R'_g,$$

$$c'_{1b} = w'_b + k'R'_b.$$

Here, subscripts refer to second-period states. The agent's total, present-value utility can now be written as

$$u(c) + \beta\left\{\pi_g \pi_{0|g} u(c'_{0g}) + \pi_b \pi_{0|b} u(c'_{0b}) + \pi_g \pi_{1|g} u(c'_{1g}) + \pi_b \pi_{1|b} u(c'_{1b})\right\}.$$

2.3. The Elimination of Cycles

Our modeling of aggregate and idiosyncratic fluctuations—the two exogenous stochastic processes for z and ϵ—does not provide any guidance for answering what would happen if cycles were eliminated. First, as in Lucas' work, since the origins of fluctuations are exogenous, it is not clear how fluctuations could be eliminated at all. This is a clear weakness of the present approach. We will follow Lucas in not describing explicit stabilization policies in our experiments. Instead, we simply eliminate cycles directly by considering alternative shock processes—processes without stochastic aggregate movements.

Given that we eliminate cycles by directly altering the exogenous shock processes, is it clear what specific processes should replace the original ones? It is not: the exogeneity assumption leaves this issue unanswered entirely. Lucas replaces the shock in his model with its mean, assuming that there could be no average consumption gain—or loss—from eliminating cycles. We wish to follow this "neutrality" assumption. However, it is not obvious how to implement this idea here. One reason is that we have two sources of consumption movements, one aggregate and one idiosyncratic.

First, as regards the aggregate shocks, we replace z and u by their conditional means. In the long run (and in the two-period economy), this means that the economy without cycles has productivity $\pi_g z_g + \pi_b z_b$ and unemployment rate $\pi_g u_g + \pi_b u_b$, where π_g and π_b are the unconditional probabilities of good and bad aggregate states, respectively. Along the transition path, the productivity and unemployment variables are calculated the same way but with conditional probabilities, so that there is no direct gain or loss from eliminating cycles arising solely from the initial aggregate state.

Using the expected values of z and u in the economy without cycles, as we do, seems natural. As a result, however, average output (ignoring the endogeneity of capital) is not the same across the economies with and without cycles. This is because output is not linear in z and u: in particular, since production is convex in z and $1-u$ and since z and $1-u$ have a positive comovement, our procedure leads to output being slightly higher on average in the economy with cycles. Nonetheless, since it is easy to compute the size of this difference in percentage terms, we can identify this effect in the final welfare results.

Second, as regards the idiosyncratic shocks, there is again no guidance within the model for what the idiosyncratic shocks should look like in the economy without aggregate shocks. At one extreme, one might imagine, as did Beaudry and Pages (1996), that idiosyncratic shocks disappear entirely if aggregate shocks are eliminated. At the other extreme, one could imagine idiosyncratic risk being larger in the economy without aggregate shocks.

Atkeson and Phelan (1994) suggested that one useful principle here is to remove the correlation between the idiosyncratic shocks of different individuals, leaving each individual's shock process unchanged. They pointed out that one way to remove correlation is to continue to give each agent a z shock, but where the z shock is now idiosyncratic rather than common to all agents. This principle implies that any effect on welfare of eliminating cycles must come through changes in the price processes, and Atkeson and Phelan discussed a particular example of how substantial variability in bond prices under aggregate risk can have large effects on individuals' welfare.

Here, we adopt a different assumption, one which also removes any correlation across individuals. We assume that eliminating aggregate shocks amounts to *integration* over the aggregate shock. Suppose the individual variable of interest, y, is a function g of an idiosyncratic shock ϵ and an aggregate shock z, $y = g(\epsilon, z)$, and that the joint distribution of the shocks is given by a density $f(\epsilon, z)$. We then identify the idiosyncratic shock process in the absence of aggregate risk, $\tilde{y}$, with

$$\tilde{y} = \int_z g(\epsilon, z) f(z|\epsilon)\, dz$$

for each ϵ, with density $\int_z f(\epsilon, z)\, dz$.

A simple example which illustrates the difference between our assumption and that of Atkeson and Phelan is as follows: suppose y denotes an individual productivity (or wage) level, and that it is the sum of two independent N(0, 1) shocks, one idiosyncratic and one aggregate,

$$y = \epsilon + z,$$

implying that y is N(0, 2). Then we obtain

$$\tilde{y} = \epsilon,$$

which is N(0, 1). Here, $\tilde{y}$ is clearly less risky—it has a lower variance than y. Atkeson and Phelan's principle here would mean that individuals' shocks have the same variance (indeed are the same) whether or not there are cycles: to them, $\tilde{y}$ would still be equal to $z + \epsilon$ in the absence of cycles, with z now being an idiosyncratic shock which is uncorrelated across agents.

For a $\{0, 1\}$ employment process, such as the one that we consider in this paper, our assumption reduces to the Atkeson and Phelan assumption. In our case, we have $g(\epsilon, z) = \epsilon$, implying quite trivially that $\tilde{y} = 1$ with probability $\pi_{1g} + \pi_{1b}$ and $\tilde{y} = 0$ with probability $\pi_{0g} + \pi_{0b}$.

The integration principle we suggest extends to a setting with multiple time periods. It is straightforward to integrate over future zs in order to obtain probabilities *as of time zero* for employment at any future date. However, the new stochastic process also needs to be supplemented with an assumption about information sets; formally, the process needs to be assigned a σ algebra and an associated probability measure. This assumption is important because it affects the agent's ability to forecast. In particular, in the economy without cycles, what information does an agent use at a future date to forecast his subsequent employment? One possibility is to assume that the agent's information is the same in the economies with and without cycles. This assumption seems neutral and is the one we use. As in Atkeson and Phelan (1994), it amounts to letting the agent receive a z shock as well as an employment shock, but where the z shock is now idiosyncratic rather than common to all agents. In this case, the z shock has no effect on an agent's income in the current period: its only usefulness is to improve the agent's forecasts of future income. Notice that the multiperiod extension of the integration principle we propose here coincides with the Atkeson–Phelan principle in the case of a two-state process.

Other assumptions about information sets are possible. One would not allow the agent to see a z, so that the only information available to the agent is the employment outcome. As a consequence, the agent would like to "obtain information about the z" by using past employment realizations—the employment process, after all, is correlated with the unobservable z process upon which it was based. This assumption about information sets therefore leads to an employment process which is infinite-order Markov.

Even with an infinite history of past employment realizations, however, the agent still has less information with which to make forecasts in the no-cycle economy than in the cycle economy.

Yet another assumption, which is the one used in İmrohoroğlu (1989), is both to suppress z and to further restrict the employment process to be first-order Markov. This assumption cannot be made fully consistent with the integration principle (unless the z process is uncorrelated with the employment process). However, one can select future conditional probabilities of ϵ' given ϵ that coincide on average—across z realizations—with those of the original (ϵ, z) process. We investigated this alternative in our quantitative work and found that the employment duration process has fatter tails in the economy with cycles than in the economy without, which may explain in part why İmrohoroğlu obtains positive welfare costs of cycles.

2.4. Analysis of the Two-Period Model

Would the consumer in this economy like to have the aggregate uncertainty eliminated? To structure the analysis, we proceed in two steps. We first consider a decision-theoretic model for which we assume that prices without cycles are simply the conditional means of prices in the economy with cycles: $R' \equiv \pi_g R'_g + \pi_b R'_b$ and $w' \equiv \pi_g w'_g + \pi_b w'_b$. We then discuss briefly some general equilibrium issues.

In the economy without aggregate uncertainty, the utility of the agent is

$$u(c) + \beta\left\{\pi_0 u(c'_0) + \pi_1 u(c'_1)\right\},$$

where $c'_0 = k'R' + g'$ and $c'_1 = w' + k'R'$.

Consider first the case where the aggregate uncertainty economy actually has no randomness in prices: $R'_g = R'_b$ and $w'_g = w'_b$. Here, the consumer is indifferent as to which economy to live in, since the utility function and the constraints now are the same in the two economies. This means that the fact that there are four states instead of two in the aggregate uncertainty world is inconsequential—in particular, it does not amount to more risk for the individual consumer. This case is actually the one that İmrohoroğlu (1989) considered (although in a full dynamic setting): there are no price movements in her economy, and the only aggregate shock is the shock to the employment rate. Since, in addition, her prices are exogenous, the elimination of aggregate shocks cannot have any consequence: agents face the same employment process and prices as before, so their behavior and welfare are unaltered.[5]

[5] As discussed in Section 2.3 and as was also pointed out in Atkeson and Phelan (1994), the reason why her results are not zero identically is that, in the case without aggregate shocks, she actually uses a shock process for the individuals which is not the process one obtains by making the aggregate shock idiosyncratic.

Now consider the case when the aggregate uncertainty is directly payoff-relevant: $w'_g \neq w'_b$ and $R'_g \neq R'_b$. Here, it is useful to separate the utility function under aggregate uncertainty into two parts, each corresponding to one employment state. When the aggregate and the idiosyncratic shocks are uncorrelated, these two parts can be studied separately. Focusing on the unemployed state, let us compare

$$\pi_0 u(k'R' + g')$$

to

$$\pi_g \pi_{0|g} u(k'R'_g + g') + \pi_b \pi_{0|b} u(k'R'_b + g').$$

Due to the strict concavity of u, the former is strictly greater than the latter for all values of k' if R' is the convex combination of R'_g and R'_b with weights $\pi_{g|0}$ and $\pi_{b|0}$, respectively. If ϵ and z are independent, this is indeed the case; an analogous argument holds for the part of the utility function that conditions on employment in the next period. To summarize, if the unemployment rate were the same in good as in bad times and if general equilibrium effects on prices left no average increase or decrease in wages and rental rates as aggregate shocks were eliminated, all agents would strictly prefer to live in the economy without aggregate shocks.

However, because the individual and aggregate shocks are correlated in our formulation, the analysis is more complicated. We first consider the effects of fluctuating wages, keeping rental rates constant, and we thereafter make wages constant but let rental rates fluctuate. Finally, we comment on the general case.

Wage Fluctuations

Suppose that $R'_g = R'_b$ but that $w'_g > w'_b$. Then the first pieces of the second-period utility—those for the unemployed state—are the same for the two economies. Further, if $\pi_{g|1} < \pi_g$, then the part conditional on employment satisfies

$$\pi_1 u(w' + k'R') > \pi_1 u(w'_g \pi_{g|1} + w'_b \pi_{b|1} + k'R'),$$

which in turn is greater than

$$\pi_1 \pi_{g|1} u(w'_g + k'R') + \pi_1 \pi_{b|1} u(w'_b + k'R')$$

by strict concavity. Therefore, in this case, utility is strictly higher (for all savings levels, including the optimal one) without aggregate uncertainty. If, on the other hand, $\pi_{g|1} > \pi_g$, which happens if and only if $\pi_{1|g} > \pi_{1|b}$, then utility may be higher in the aggregate uncertainty world.[6] The intuition is

[6]The first inequality can be rewritten as $\pi_g \pi_{1|g} / (\pi_g \pi_{1|g} + \pi_b \pi_{1|b}) > \pi_g$, which simplifies to $\pi_{1|g} > \pi_{1|b}$.

clear: in the unemployed state, it does not matter what the aggregate state is—it does not affect consumption in this state. In the employed state, on the other hand, the agent may prefer aggregate uncertainty, provided that employment is more likely in the good state than in the bad state. In this case, the wage tends to be high when the agent is employed. Consequently, conditional on employment, the agent's expected wage is higher than the unconditional expected wage, implying that it is worse for the agent to receive the unconditional expected wage. Therefore, with low enough curvature in the utility function, the agent will prefer wage fluctuations.

Rental Rate Fluctuations

Suppose now that $R'_g > R'_b$ but that $w'_g = w'_b = w'$. Then, by an argument analogous to the one just given, utility next period when unemployed is higher without aggregate uncertainty if $\pi_{0|g} > \pi_{0|b}$ and $k' < 0$.[7] However, for the same argument to work when the agent is employed next period, it would have to be the case that $\pi_{1|g} > \pi_{1|b}$, which contradicts $\pi_{0|g} > \pi_{0|b}$.

To simplify the analysis, let us write $\pi_{g|1} = \pi_g + \nu$. This implies $\pi_{b|1} = \pi_b - \nu$, $\pi_{b|0} = \pi_b + \nu(\pi_1/\pi_0)$, and $\pi_{g|0} = \pi_b - \nu(\pi_1/\pi_0)$. Expressing total second-period utility as a function of ν, we have, after simplification,

$$\pi_1\pi_g u(w' + k'R'_g) + \pi_1\pi_b u(w' + k'R'_g) + \pi_0\pi_g u(k'R'_g) + \pi_0\pi_b u(k'R'_g)$$
$$+ \nu\pi_1[(u(w' + k'R'_g) - u(w' + k'R'_b)) - (u(k'R'_g) - u(k'R'_b))].$$

The sum of the first four terms in this expression is less than the utility without cycles. Moreover, since u is concave and $w' > 0$, the final term is negative provided $k' > 0$ and $\nu > 0$, or provided $k' < 0$ and $\nu < 0$. That is, one can show with either of these two provisions that agents prefer the economy without fluctuating prices. Is it possible, say, if $k' < 0$ and $\nu > 0$, that the utility is higher with price fluctuations? We have not been able to provide conditions under which this is true. Here, unlike in the example with wage fluctuations, the loss from eliminating risk relies on concavity of the utility function; this loss increases as concavity increases. A higher degree of concavity, however, also increases the gain from eliminating risk.

The preceding analysis shows that a consumer's views on aggregate risk depend on the nature of the risk. Although consumers, in general, prefer constant to fluctuating rental rates, they may, under certain conditions, prefer fluctuating to constant wages. In addition, a consumer's views on aggregate risk depend on his individual characteristics, such as his wealth and current employment status. Wealth is important because, in the absence of insurance markets and in the presence of a constraint on borrowing, it helps

[7]Similarly, utility without aggregate uncertainty is also higher if $\pi_{0|g} < \pi_{0|b}$ and $k' > 0$.

to insure the consumer against idiosyncratic risk. The preceding analysis of wage and rental rate fluctuations shows that a consumer's attitude toward risk, as captured by the degree of concavity in the consumer's utility function, plays a key role in determining the benefits of eliminating aggregate risk. More generally, the extent to which a consumer is well insured, as captured by the size of his wealth holdings, will play an important role in determining the benefits to eliminating risk. In other words, consumers who are very poor—especially those who are close to zero consumption—are likely to see large gains from eliminating aggregate risk. Wealthy consumers, on the other hand, do not appreciate this benefit as much.

Individual wealth is also important because it determines the composition of the consumer's income. Very wealthy agents mainly care about fluctuations in the rental rate, since wage income is a small part of their total income, whereas consumers with close to zero savings do not care about rental rates. For consumers with significant negative wealth, rental rate fluctuations again become important. These consumers worry about how much they have to pay in interest payments, and they are especially afraid of large interest rate realizations when they are unemployed and have no wage income.

We saw in the analysis above that if utility functions are rather flat or if most consumers are well insured (as turns out to be the case in our calibrated model), wage fluctuations are liked by all but the very poorest consumers, whereas rental rate fluctuations are disliked. It is therefore possible that consumers' views on the benefits of eliminating aggregate vary nonmonotonically with wealth: the very poorest consumers benefit from the elimination of cycles, because they are very concerned about aggregate risk (especially rental rate risk); the very richest consumers benefit too, because wage income is irrelevant for them; but consumers with modest wealth do not benefit, because wages are their main source of income.

The consumer's employment status also plays a role in determining his attitudes toward aggregate risk. When employment is positively serially correlated, as it is in the calibrated model, employed consumers are better insured, since they are more likely to receive wage income in the future. These consumers also care more about wage rate fluctuations than about rental rate fluctuations, since a larger part of their expected income is in the form of wages.

Turning now to some general equilibrium considerations, let us recall how prices—the returns to capital and the wage rate—are determined: they are given by the marginal products of an aggregate, Cobb–Douglas production function whose inputs are total capital and total employment. When we eliminate the aggregate exogenous shock by replacing the stochastic productivity and employment variables with their means, the functional form matters for the end result. In general, unlike in our above partial-equilibrium experiments, the economy without aggregate uncertainty will not have rental and wage rates that are the averages (taken across the aggregate state) of the corresponding rates in the economy with aggregate uncertainty. This occurs for two reasons: first, the capital stock is endogenous, and second, the pricing functions are not linear in z (and u).

When the amount of uncertainty changes for individuals, savings change. For the parameterizations we use in our quantitative model below, we find that savings are higher with more uncertainty—an effect of precautionary savings. This implies that one effect of eliminating aggregate uncertainty is to push wage rates downward and rental rates upward.

When z is replaced by its mean, will the rental rate be higher or lower than the average rental rate in the economy with stochastic productivity? The answer depends on specific parameter values. In terms of the rental rate function, it is convex in the capital/labor ratio: assuming that aggregate capital is the same across the two economies, if u_z is replaced by its mean, the average value of r is higher. However, the productivity variable fluctuates as well and is correlated with the input fluctuations. In our parameterizations, the unemployment rate fluctuates more than z does, and we find that average rental rates are higher in the economies without aggregate uncertainty, holding aggregate capital fixed. For parallel reasons, average wage rates are lower.

The differential effects on average wage rates and rental rates of eliminating cycles imply that individuals with low wealth have a reason to be against the elimination of cycles: the price of labor, which is what they care about, is lower on average. High-wealth individuals, on the other hand, see benefits from eliminating cycles.

In summary, the two-period model teaches us that (i) the elimination of stochastic movements in z and u does not necessarily lead to increases, and may even lead to decreases, in utility, and (ii) welfare effects differ across agents as a function of their employment and wealth statuses. The absolute and relative magnitudes of the effects we have discussed also depend on the aggregate state, on the serial correlation properties of the shocks, and on the size of the capital stock. Furthermore, not only does the parameterization of preferences matter in our economy, but the form of the aggregate production function matters as well.

3. THE QUANTITATIVE MODEL

We now turn to the model that we use in our computational experiments. This model is calibrated to observed data on employment, income, and wealth. Since this model has already been exposited in a simplified version in Section 2.1, we focus here on what is different in the general version, provide some formal aspects of the equilibrium definition, briefly discuss computation, and describe our calibration.

3.1. Setup

Compared to the model in Section 2.1, there are three changes: (i) we specialize to logarithmic utility; (ii) there are preference shocks; and (iii) we consider a distinction between long- and short-term unemployment.

The preferences are

$$E_0 \sum_{t=0}^{\infty} \beta_t \log c_t,$$

where β_t is a stochastic variable which is idiosyncratic—i.i.d. across agents—and describes the cumulative discounting between period 0 and period t. In particular, $\beta_{t+1} = \tilde{\beta}\beta_t$, where $\tilde{\beta}$ is a three-state, first-order Markov process.

Let $\epsilon \in \{1, 2, 3\}$, where 1 denotes long-term unemployed, 2 denotes short-term unemployed, and 3 denotes employed. The distinction between short- and long-term unemployment allows us to consider differences among the unemployed both in terms of their income when unemployed and their prospects for future employment. In particular, in the calibration we assume (i) that short-term unemployed receive higher unemployment insurance benefits, $g_2 > g_1 > g_3 = 0$, and (ii) that their probability of employment is higher, with the difference being more pronounced in recessions than in booms. As before, the individual employment status, jointly with the aggregate shock z, follows a first-order Markov chain.

Formally, a recursive competitive equilibrium for this economy is defined using the aggregate state variables. Let Γ denote the current measure of consumers over holdings of capital, employment, and preference status. Then, the state variable relevant to the individual includes (Γ, z) and the idiosyncratic vector $(k, \epsilon, \tilde{\beta})$. Let H denote the equilibrium transition function for Γ:

$$\Gamma' = H(\Gamma, z, z').$$

Consumers solve

$$v(k, \epsilon, \tilde{\beta}; \Gamma, z) = \max_{c,\, k'} \{u(c) + \tilde{\beta} E[v(k', \epsilon', \tilde{\beta}'; \Gamma', z') | z, \epsilon, \tilde{\beta}]\}$$

subject to

$$c + k' = r(\bar{k}, 1 - u_z, z)k + w(\bar{k}, 1 - u_z, z)\, I(\epsilon) + g_\epsilon + (1 - \delta)k',$$

$$\Gamma' = H(\Gamma, z, z'),$$

$$k' \geq \underline{k},$$

where $I(\epsilon) = 1$ if $\epsilon = 3$ and 0 otherwise. If

$$k' = f(k, \epsilon, \tilde{\beta}; \Gamma, z)$$

denotes the optimal saving decision for the agent, then an equilibrium can be defined as a law of motion H, individual functions (v, f), and pricing functions (r, w) such that (i) (v, f) solves the consumer's problem, (ii) (r, w) equal the marginal products of capital and labor, respectively, and (iii) H is generated by f and the law of motion for $(z, \epsilon, \tilde{\beta})$. The economy without cycles is defined in the same way, but using different processes for z (which is now deterministic) and ϵ.

3.2. Calibration

For the most part, our calibration is standard in that it is close to real-business-cycle practice. In particular, we interpret a period to be a quarter, and choose $\delta = 0.025$ and $\alpha = 0.36$.

We calibrate the discount factor process by assuming a symmetric distribution of $\tilde{\beta}$s—with 80% of the population on the middle value and 10% on each extreme point in any time period—and an expected duration of the extreme discount values of 50 years (approximating a lifetime). The specific numerical values of $\tilde{\beta}$ are selected so that some wealth distribution statistics are similar to what they are in the data; these are displayed at the end of this section. For the calibrations we consider, the difference between consecutive values of $\tilde{\beta}$ is roughly one-half of a percentage point.

We select aggregate shocks so that we approximate the movements in observed output fluctuations in the postwar United States; based on a u_b equal to 10% and a u_g of 4%, we therefore select $z_g = 1.01$ and $z_b = 0.99$, and we set the expected duration of each aggregate state to 2 years.

The borrowing constraint is set, roughly speaking, to be the loosest possible; in particular, we set it so that at a constant, high, interest rate, the agent is just able to pay back even with maximally bad individual employment luck. This means that the largest amount of borrowing is about 60–70% of average annual output per worker in the economy with cycles, not including home production.

We present results from two calibrations with different employment dynamics. In the first one, the short- and long-term unemployment states are collapsed into one state; this is the setup of Krusell and Smith (1998), which in turn follows the tradition of İmrohoroğlu's work. We refer to this as our baseline calibration. In this case, we select g so that the lower part of the wealth distribution looks like the data, implying a value corresponding to about 10% of the quarterly wage: $g = 0.0334$. The discount factors in this case are 0.9858, 0.9894, and 0.9930. The employment process here can be described by four 2-by-2 matrices, one for each (z, z'):

$$\begin{pmatrix} 0.33 & 0.67 \\ 0.03 & 0.97 \end{pmatrix}$$

for the transition $(z, z') = (z_g, z_g)$ (rows indicate the current state and columns indicate next period's state; row 1 is the state of unemployment and row 2 the state of employment),

$$\begin{pmatrix} 0.75 & 0.25 \\ 0.07 & 0.93 \end{pmatrix}$$

for (z_g, z_b),

$$\begin{pmatrix} 0.25 & 0.75 \\ 0.02 & 0.98 \end{pmatrix}$$

for (z_b, z_g), and

$$\begin{pmatrix} 0.60 & 0.40 \\ 0.04 & 0.96 \end{pmatrix}$$

for (z_b, z_b).[8]

As is explained in detail in Krusell and Smith (1998), we select parameter values to satisfy the requirements (i) that the aggregate unemployment can only take on two values and (ii) that the expected duration of unemployment is 1.5 quarters in the good aggregate state (that is, the duration given that the good aggregate state persists) and 2.5 quarters in the bad aggregate state. In this calibration, cycles have the property that individual risk is more severe in bad aggregate states: the unemployment rate is higher in this state, as is the expected duration of unemployment. As emphasized by Mankiw (1986), the countercyclicality of individual risk is crucial for generating increased risk premia.

In the second calibration, where we make the distinction between short- and long-term unemployment, we use the discount factors 0.9823, 0.9879, and 0.9935. We use $g_2 = 0.391$, which is about 50% of the quarterly wage, to roughly replicate the U.S. replacement ratio during the first quarter of unemployment,[9] and we select $g_1 = 0.038$ to match the left tail of the wealth distribution.

[8] The numbers in the matrices are rounded to two digits.

[9] Unemployment insurance at this rate can normally be collected for the first two quarters and sometimes longer. Here, we assume it can only be collected for one quarter for computational convenience.

The transition matrices between employment states here are

$$\begin{pmatrix} 0.50 & 0 & 0.50 \\ 0.25 & 0 & 0.75 \\ 0 & 0.03 & 0.97 \end{pmatrix}$$

for the (z_g, z_g) transition (rows indicate the current state and columns next period state; recall that 1 means long-term unemployed, 2 short-term unemployed, and 3 employed),

$$\begin{pmatrix} 0.17 & 0 & 0.83 \\ 0.03 & 0 & 0.97 \\ 0 & 0.03 & 0.97 \end{pmatrix}$$

for the (z_b, z_g) transition (that is, going from z_b to z_g),

$$\begin{pmatrix} 0.94 & 0 & 0.06 \\ 0.75 & 0 & 0.25 \\ 0.04 & 0.03 & 0.93 \end{pmatrix}$$

for the (z_g, z_b) transition, and

$$\begin{pmatrix} 0.99 & 0 & 0.01 \\ 0.03 & 0 & 0.97 \\ 0 & 0.03 & 0.97 \end{pmatrix}$$

for the (z_b, z_b) transition.[10]

The restrictions we impose on these matrices include (i) the restrictions on expected duration that we use, and nonnegativity of probabilities, which imposes nonlinear restrictions on parameters; (ii) the requirement that aggregate unemployment takes on only two values, which severely limits what can be assumed on the individual level; and (iii) the (definitional) restrictions that the long-term unemployed cannot transit to short-term unemployment and always go through short-term unemployment first.[11] We also

[10] The numbers in the matrices have been rounded to two significant digits; the exact number for $\pi_{1|1,z_b,z_b}$ is 0.9875.

[11] An exception to the first of these statements can be found in the transition from the good to the bad aggregate state. There it is possible to go directly from $\epsilon = 3$ to $\epsilon = 1$. We had to use this parameterization in order to avoid making the probability of employment next period higher for unemployed agents than for employed agents. However, we do make sure that agents who go from employment to long-term unemployment when the aggregate state goes from bad to good receive g_2, that is, are treated as short-term unemployed, during their first quarter of unemployment. This, in effect, formally forces the last period's value for z to be part of the current aggregate state. Thus, in our computation, we do need to make a distinc-

TABLE I
The Distribution of Wealth

	% of wealth held by top					Fraction with wealth < 0 (%)	Gini coefficient
	1%	5%	10%	20%	30%		
One kind of unemployed	24	54	72	87	91	11	0.81
Two kinds of unemployed	25	56	73	84	88	12	0.78
Data	30	51	64	79	88	11	0.79

impose the requirement that the probability of employment is always higher for currently employed than for currently unemployed.

The difference between average unemployment durations in the good and bad aggregate states is substantial in our calibration. The expected duration of unemployment in the bad aggregate state is 80 periods for long-term unemployed (e.g., a little less than half of a working lifetime), whereas it is only 2 periods in the good aggregate state.[12] Relatedly, the fraction of all unemployed agents consisting of long-term unemployed is much higher in the bad aggregate state than in the good aggregate state: 73% versus 33%. In fact, the total number of short-term unemployed almost does not change at all across the aggregate states in our calibration (it is 0.027 in the bad state and 0.0268 in the good state), so what a recession does is to add a number of long-term unemployed to the economy. Thus, there is a potential for more significant suffering from bad aggregate shocks among unlucky consumers in this calibration than in the one which does not distinguish short- from long-term unemployment.

With this calibration, we obtain the average long-run wealth distributions shown in Table I. The wealth distributions generated by the two model calibrations are both quite similar to U.S. data (we use data based in Wolff (1994) and Díaz-Giménez, Quadrini, and Ríos-Rull (1997)). The relatively parsimonious three-discount-factor setup allows us to roughly capture the broad features of observed wealth inequality: substantial skewness, with most of the capital held by the very richest agents, and a large mass of people with close to or below zero wealth.

tion between those bad aggregate states which have persisted and those which come directly following a good aggregate state, since they are associated with different amounts of total resources. This complication increases the possible values of the exogenous aggregate state from 2 to 3 but is the simplest way to ensure a sensible calibration of individual employment dynamics subject to maintaining the simplification that aggregate unemployment can only take on two values.

[12]But note that, since recessions last only eight periods on average, the average duration of an unemployment spell is relatively small: a little more than two periods, which is approximately the same average duration that obtains in the model with two idiosyncratic employment states.

3.3. Model Solution

We solve the model with aggregate fluctuations using the technique employed and described in Krusell and Smith (1998). In brief, this technique works as follows: agents act as if only a limited set of moments of Γ matter for the determination of prices, and the (aggregate) result of that behavior is shown to be almost perfectly consistent with their perceptions of how prices evolve. The technique can be applied because there is "approximate aggregation" in this class of models: the aggregates are determined mainly by those whose wealth is not near zero and these agents have almost identical savings propensities. For more details on the computation relying on approximate aggregation, see Krusell and Smith (1998).

Approximate aggregation does not imply that all the model properties are close to those of a standard representative-agent model. Aggregate capital accumulation is mainly determined by the very richest, since wealth is so unevenly distributed, and they behave like typical representative-agent, "permanent-income" consumers; hence the approximate aggregation result. However, the poorer consumers do not smooth consumption well at all: they can be referred to as "hand-to-mouth" consumers. Since their consumption is a much larger fraction of total consumption, this implies a much lower correlation between aggregate consumption and aggregate output than in representative-agent models. The fact that the consumption processes are quite different leaves open the possibility that the risk associated with cycles is substantial for many consumers.

In this paper, we also need to compute transition paths for economies without aggregate shocks. The central idea is to postulate a time path for aggregate capital, solve for agents' decisions given this path, and then verify that the time path for aggregate capital implied by agents' aggregated decisions matches the postulated time path. The Appendix describes the algorithm in detail.

The computed equilibrium laws of motion for aggregate capital describe the accuracy of our computations in the economies with cycles. They are

$$\log \bar{k}' = 0.100 + 0.960 \log \bar{k},$$

$$R^2 = 0.999991, \qquad \hat{\sigma} = 0.0056\%$$

in good times and

$$\log \bar{k}' = 0.095 + 0.960 \log \bar{k},$$

$$R^2 = 0.999987, \qquad \hat{\sigma} = 0.0075\%$$

in bad times for the baseline calibration. The R^2 figures indicate the extent of the deviation from rationality: when these laws of motion for capital are taken as given by agents, the agents' implied savings behavior aggregates up to a capital stock series which, when regressed on current capital, produces the stated coefficients and the reported R^2s and percentage standard errors $\hat{\sigma}$. Since aggregation does not hold strictly in this model due to the incompleteness of markets, any regression error could be avoided by using more information about the distribution of capital and a more general functional form than the log-linear one used here. However, since the fit is so impressive, only very tiny improvements in forecasts are possible for the agents. Moreover, these improvements in turn are even less important in utility terms—utility losses for consumers in this setup are extremely small even for significant departures from the optimal decision rules (this and similar points are elaborated on in Lucas (1987), Cochrane (1989), and Krusell and Smith (1996), among others).

For the calibration with short- and long-run unemployed, the corresponding equations are

$$\log \bar{k}' = 0.105 + 0.958 \log \bar{k},$$
$$R^2 = 0.99997, \qquad \hat{\sigma} = 0.0094\%$$

in good times,

$$\log \bar{k}' = 0.116 + 0.952 \log \bar{k},$$
$$R^2 = 0.9997, \qquad \hat{\sigma} = 0.028\%$$

in bad times when the last period was bad as well, and

$$\log \bar{k}' = 0.092 + 0.962 \log \bar{k},$$
$$R^2 = 0.99998, \qquad \hat{\sigma} = 0.0084\%$$

in bad times when the last period was good. The fit is a little worse here than in the baseline case, especially when two bad aggregate shocks hit in succession, but it is still impressive.

4. WELFARE EFFECTS OF ELIMINATING THE BUSINESS CYCLE

Since we want to record the welfare effects of eliminating cycles for different groups of agents, we need to solve for transition paths. This means that it is impossible to avoid movements in the capital stock, because agents adjust their savings in the new shockless aggregate environment toward the steady state. Movements in the exogenous aggregate variables can be separated into expected and unexpected parts. Our experiment is to eliminate only the unexpected part, that is, to replace the stochastic processes for z and u with their conditional expectations as of the initial date. This leaves a deterministic movement in z and u which disappears in the long run. Finally, recall that the idiosyncratic employment/preference shock process for the economy without aggregate fluctuations is the same one as for the economy with aggregate fluctuations.

4.1. The Baseline Case: Homogeneous Unemployed

With a two-state process for employment—the agent is either employed or unemployed—the long-run welfare gain from eliminating cycles turns out to be 0.138%. To compute the long-run welfare gain, we first compute the expected value of the steady-state distribution of lifetime utilities in the economy with aggregate shocks. We perform this calculation by averaging across both agents and time in a long simulation consisting of 10,000 time periods and 30,000 agents. We then perform a similar calculation in the economy without aggregate shocks. Finally, we convert the difference between the two expected values into a consumption equivalent in the same way that Lucas did. The long-run welfare gain in the economy with two employment states is nearly 20 times larger than the welfare gain (0.008%) in Lucas (1987), but remains very small nevertheless.

Turning to the differential effects across agents and, thus, to our transition experiments, Table II shows some of the main statistics we obtained. The most striking finding is that the welfare gains are very small, indeed negative, for almost all groups we cover in the table. On average, there is a welfare loss from eliminating cycles. The precise amount depends on the initial aggregate capital stock and aggregate exogenous state; the loss is higher starting from a low capital stock. Recall that the steady-state comparison, in contrast, gives an average welfare gain from eliminating cycles. That gain cannot be explained by a difference between the capital stock in the steady-state equilibrium without cycles and the average capital stock in the stationary stochastic equilibrium—the latter is higher. We have not

TABLE II
Welfare Effects for Different Agents from Eliminating Cycles:
The Two-State Employment Process

Initial state		Utility gain in percentage consumption						
			Employed agents (by wealth percentile)			Unemployed agents (by wealth percentile)		
z	$\bar{k}$	All agents	< 1	25–50	> 99	< 1	25–50	> 99
z_g	11.2	−0.008	−0.016	−0.007	−0.017	−0.016	−0.008	−0.017
z_g	12.3	−0.006	−0.010	−0.004	−0.021	−0.018	−0.004	−0.021
z_b	11.2	−0.010	−0.020	−0.012	0.014	−0.040	−0.013	0.014
z_b	12.3	−0.005	−0.008	−0.003	−0.031	−0.018	−0.003	−0.032

been able to isolate exactly what it is about the transition which produces welfare losses from eliminating cycles. It is apparent that the wage and interest rate paths, which slowly converge to their long-run constant levels, are not well liked by most agents compared to the stationary stochastic price processes of the economy with cycles. One possibility is that the transition experiment retains an aspect of the cycle—the slowly moving, deterministic price path—which was not present in the steady-state comparison, in which the no-cycle economy has constant prices and in which there is a modest welfare gain from eliminating cycles.[13]

Table II also shows substantial differential effects between employed and unemployed agents, and among agents with different wealth holdings. The employment status matters significantly only at low wealth levels; agents with substantial wealth do not particularly care whether they are employed or unemployed. How the employment status matters depends on the specific case. First, it is an important part of the individuals' present-value income. Second, it also determines wage income as a fraction of total income, and therefore how much the agent cares about wage fluctuations relative to rental rate fluctuations. As explained in the context of the two-period model in Section 2.2, fluctuations in the wage rate may be good for poor agents, since the probability of a good aggregate state, given employment, is higher than the unconditional probability of a good aggregate state in this calibration. The welfare effects across different wealth levels are nonmonotonic in some cases: eliminating cycles is more detrimental for the poor agents than for the middle class, and whether the rich are better or worse off in relative terms depends on the initial aggregate conditions.

It should also be recalled that the experiments in this section make output somewhat higher on average in the cyclical economy (see the discussion in Section 2.3). The difference, for both calibrations, is around 0.009% in the intermediate and long run, and a little less in the short run. Thus, one might want to add an amount of slightly less than 0.009 to all the numbers in the table, leaving most of the totals positive, although many group numbers remain negative.

In summary, whereas one might have expected a priori that poor, unemployed agents gain significantly from eliminating cycles, since the risk associated with price fluctuations is more costly for them, they do not, on average, even gain. The intuition based on risk is not totally wrong, however: the table does not show what happens to agents who are right at the borrowing constraint. In the cases where such agents are very close to zero

[13]Transitions sometimes involve nonmonotonic paths for capital; if the initial aggregate state is good (bad) and aggregate capital is high (low), the capital stock first climbs (falls), then falls (climbs).

consumption, there is a substantial gain from eliminating cycles: gains of close to 2% are possible. However, there are vanishingly few agents in this situation: since it causes significant pain to be close to zero consumption with a binding borrowing constraint, agents save enough over time that it is very unlikely for them to end up in such a situation.

The table does not show how the current degree of patience matters. The typical dependence on patience is as follows: higher discount factors lead to lower losses from eliminating cycles (or higher gains), presumably because the long-run welfare effects of eliminating cycles are positive, a fact which patient agents appreciate more than less patient agents.

Finally, in order to compare our results with those from models where the wealth distribution is much less skewed—as in İmrohoroğlu (1989)—we also calculated the (long-run) welfare effects of eliminating cycles in an economy without preference heterogeneity (yielding a wealth distribution with a Gini coefficient of around 0.3). Requiring a realistic wealth distribution from the model increases the welfare costs by about a factor of 7.

4.2. The Calibration with Short- and Long-Term Unemployment

The model with more severe unemployment shocks produces similar results for steady-state utility levels: the long-run gain from eliminating cycles is 0.068%. Table III shows the effects from the transition experiments. As in the case of the two-state calibration, the welfare gains from eliminating cycles are often negative. Here, in two of the four cases there is a loss on average across agents. As before, the differential effects across employment and wealth types are nontrivial and depend significantly on initial conditions. The main conclusion is that it is hard to generate significant welfare gains from eliminating cycles in this kind of model, even if the goal is simply to obtain large gains for some agents.

TABLE III
Welfare Effects for Different Agents from Eliminating Cycles:
The Three-State Employment Process

Initial state			Utility gain in percentage consumption								
			Employed agents (by wealth percentile)			Short-term unemployment (by wealth percentile)			Long-term unemployment (by wealth percentile)		
z	$\bar{k}$	All agents	< 1	25–50	> 99	< 1	25–50	> 99	< 1	25–50	> 99
z_g	11.2	−0.024	−0.059	−0.030	0.089	−0.150	−0.033	0.081	−0.340	−0.040	0.093
z_g	12.1	0.005	0.022	0.024	−0.241	−0.018	0.024	−0.225	−0.097	0.023	−0.252
z_b	11.2	−0.040	−0.038	−0.021	0.056	−0.075	−0.022	0.049	−1.019	−0.039	0.045
z_b	12.1	0.004	0.028	0.025	−0.261	0.012	0.025	−0.241	−0.148	0.018	−0.253

5. SUMMARY

We find that the steady-state average gains from eliminating cycles are higher than in Lucas' (1987) representative-agent calculations. However, these gains are still quite low in absolute terms: about 0.1% of average consumption. Moreover, if one takes the transition into account, which seems the only sensible thing to do, the gains are smaller than in Lucas' work, and often negative. Finally, although we do find significant differences among agents in how they are affected by eliminating cycles, only the vanishingly few agents who are very close to the zero consumption level can gain substantial amounts (up to 2%). Furthermore, apart from these very poorest agents, of whom there would be no more than a handful in a sample as large as that of the United States, those who gain the most among all the rest are often the wealthiest agents.

Our estimates of the welfare gains from eliminating business cycles suffer from a number of weaknesses. The model that we use to construct our estimates embodies an important weakness: it treats the aggregate and idiosyncratic shocks as exogenous. This means that the comparison between a situation with cycles and one without requires us to construct a different model, one without aggregate shocks. Assumptions in this different model regarding production possibilities and idiosyncratic shocks are difficult to defend. We choose assumptions in the spirit of Lucas; they are, in a certain sense, neutral. However, defending them further seems impossible without adopting a view about the dependence of aggregate cycles on government policy.

We also abstract from a number of features that may be quantitatively important. We assume that consumers have infinite lives (they are part of dynasties). Our results show that dynasties can protect their members very well in utility terms from idiosyncratic shocks using only one asset, even in the presence of a borrowing constraint. Models with finite lives and imperfect altruism may give different results in this respect. We also assume that there is no idiosyncratic wage heterogeneity. Idiosyncratic wage risk is likely much larger than aggregate wage risk, and a model with this feature might lead to larger estimates of the welfare costs of business cycles.

APPENDIX

This appendix describes the algorithm that we use to calculate the transition paths in the economies without cycles. The central idea is to postulate a time path for aggregate capital, solve for agents' decisions given this path, and then verify that the time path for aggregate capital implied by agents' aggregated decisions matches the postulated time path.

We calculate our initial guess for the time path of aggregate capital as follows. First, calculate a (deterministic) law of motion for aggregate capital by taking the average of the two laws of motion (one corresponding to good times and one corresponding to bad times) for aggregate capital in the economy with cycles. Second, starting from the initial value of aggregate capital, use the law of motion recursively to determine aggregate capital for the next 600 periods (by which point aggregate capital will have converged to a steady-state value).

Given a postulated time path for aggregate capital, the first step in the algorithm is to use the last 550 time periods in this time path to calculate a (deterministic) first-order autoregressive law of motion for aggregate capital. We calculate this law of motion using ordinary least squares. We use the last 550 time periods because both aggregate employment and the aggregate productivity shock have settled down to their steady-state values after 50 periods.[14] Aggregate capital, on the other hand, converges much more slowly.

The second step in this algorithm is to solve the consumer's recursive dynamic programming problem in the economy without cycles, taking as given the first-order autoregressive law of motion for aggregate capital. We perform this task by iterating on the value function as in Krusell and Smith (1998); see, in particular, the appendix to that paper. The solution to this problem is a decision rule and a value function, both of which depend on aggregate capital and individual state variables. (Neither aggregate employment nor the aggregate productivity shock is a state variable: instead, each is set equal to its steady-state value.)

The third step in the algorithm is to iterate backward starting in period 50 to calculate agents' decision rules during the first 50 periods of transition. We therefore calculate 50 different decision rules and value functions, one for each period. Specifically, in period 50, the agent takes as given the current values of aggregate capital, employment, and productivity and the value of aggregate capital in period 51. (Recall that the consumer takes as given a time path for aggregate capital. In addition, at each time period in the backward iteration, aggregate employment and productivity are set equal to their conditional expectations as of the initial time period in the economy with cycles.) In period 50, we use the value function computed in the second step of the algorithm to determine the future value of the agent's current savings decision. (Since one of the arguments of this value function is aggregate capital, the individual needs to know the value of aggregate capital in period 51.) The solution to the agent's problem in period 50 is a decision rule and a value function, each of whose arguments are the

[14]We use 550 time periods to calculate the law of motion for aggregate capital in the economy with two employment states. In the economy with three employment states, we use the last 525 time periods to calculate this law of motion because aggregate employment does not settle down to its steady-state value until period 75.

individual state variables. We then move back to period 49, in which period the agent takes as given the values of aggregate capital, employment, and productivity in period 49 as well as the period-50 value function (this function determines the future value of the agent's savings decision in period 49). We continue iterating backward in this fashion until the initial period, storing the decision rules for later use. Although the decision rules and value functions do not have aggregate variables as explicit arguments, the period-t decision rule and the period-t value function do depend implicitly (by means of the backward iteration) on the values of the aggregate variables in period t and subsequent periods.

The fourth step in the algorithm is to use the decision rules calculated above to simulate the behavior of the economy without cycles starting from the given initial distribution of agents. In particular, we simulate the behavior of 90,000 agents for 600 time periods, by which time the economy has converged to the new steady-state equilibrium (save for negligible "wiggles" due to simulation error). The initial distribution in the simulation is a typical distribution in the economy with cycles (given the specified values of aggregate capital and the aggregate productivity shock in the initial period). By adding up the savings decisions of the 90,000 agents, we obtain a time path for aggregate capital. For the same reasons that approximate aggregation holds in the economy with cycles, the transition path for aggregate capital varies only to a very small extent as the higher moments in the initial distribution vary over the range of values typically observed in the economy with cycles (holding fixed the values of aggregate capital and the aggregate productivity shock).

The fifth step in the algorithm is to compare the new path for aggregate capital to the one that agents took as given when computing their decision rules. If the two paths are close enough together, then we have computed an equilibrium path for aggregate capital. If they are not close enough together, then we calculate a new path for aggregate capital by taking a weighted average of the postulated path and the implied path. We then return to the first step of the algorithm and continue iterating until the postulated and implied time paths for aggregate capital are close enough together.

The accuracy of this algorithm depends on the accuracy with which one calculates decision rules and the extent to which the law of motion calculated in the first step of the algorithm fits the last 550 observations in the time path for aggregate capital. Along both of these dimensions, one can make numerical error arbitrarily small. First, one can use a finer grid when constructing cubic spline approximations to the value functions (see Krusell and Smith (1998, Appendix) for further details). Second, one can use more flexible functional forms to compute the law of motion in the first step of the algorithm. We find that increasing accuracy in either of these ways changes the numerical results only to a very small degree.

REFERENCES

Aiyagari, S. R. (1994). "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics* **109**, 659–684.

Atkeson, A., and Phelan, C. (1994). "Reconsidering the Costs of Business Cycles with Incomplete Markets," *in* "NBER Macroeconomics Annual 1994," (S. Fischer and J. Rotemberg, Eds.) pp. 187–207, National Bureau of Economic Research.

Attanasio, O., and Davis, S. (1996). "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy* **104**, 1227–1262.

Beaudry, P. and Pages, C. (1996). "The Cost of Business Cycles and the Stabilization Value of Unemployment Insurance," unpublished manuscript.

Bils, M. (1985). "Real Wages Over the Business Cycle: Evidence From Panel Data," *Journal of Political Economy* **93**, 668–689.

Cochrane, J. (1989). "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *American Economic Review* **79**, 319–337.

Cole, H., and Kocherlakota, N. (1997). "A Microfoundation for Incomplete Security Markets," working paper 577, Federal Reserve Bank of Minneapolis.

Díaz-Giménez, J., Quadrini, V., and Ríos-Rull, J.-V. (1997). "Dimensions of Inequality: Facts on the U.S. Distribution of Earnings, Income, and Wealth," *Federal Reserve Bank of Minneapolis Quarterly Review* 3–21.

Dolmas, J. (1998). "Risk Preferences and the Welfare Cost of Business Cycles," *Review of Economic Dynamics*, **1**, 646–676.

Gomes, Greenwood, and Rebelo (1998). Unpublished manuscript.

Huggett, M. (1993). "The Risk-Free Rate in Heterogeneous-Agents, Incomplete Markets Economies," *Journal of Economic Dynamics and Control* **17**, 953–969.

İmrohoroğlu, A. (1989). "The Cost of Business Cycles with Indivisibilities and Liquidity Constraints," *Journal of Political Economy* **97**, 1364–1383.

Krusell, P., and Smith, Jr., A. A. (1996). "Rules of Thumb in Macroeconomic Equilibrium: A Quantitative Analysis," *Journal of Economic Dynamics and Control* **20**, 527–558.

Krusell, P., and Smith, Jr., A. A. (1997). "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns," *Macroeconomic Dynamics* **1**, 387–422.

Krusell, P., and Smith, Jr., A. A. (1998). "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, forthcoming.

Lucas, Jr., R. E. (1987). "Models of Business Cycles," Basil Blackwell, New York.

Mankiw, N. G. (1986). "The Equity Premium and the Concentration of Aggregate Shocks," *Journal of Financial Economics* **17**, 211–219.

McCallum, B. T. (1986). "On "Real" and "Sticky-Price" Theories of the Business Cycle," *Journal of Money, Credit, and Banking* **18**, 397–414.

Mehra, R., and Prescott, E. C. (1985). "The Equity Premium: A Puzzle," *Journal of Monetary Economics* **15**, 145–162.

Obstfeld, M. (1994). "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability," *European Economic Review* **38**, 1471–1486.

Storesletten, K., Telmer, C., and Yaron, A. (1998). "Asset Pricing with Idiosyncratic Risk and Overlapping Generations," unpublished manuscript.

Tallarini, Jr., T. (1997). "Risk-Sensitive Real Business Cycles," unpublished manuscript.

Wolff, E. N. (1994). "Trends in Household Wealth in the United States, 1962–83 and 1983–89," *Review of Income and Wealth* **40**, 143–174.