REVISITING THE WELFARE EFFECTS OF ELIMINATING BUSINESS CYCLES

Preliminary and Incomplete

Per Krusell and Anthony A. Smith, Jr.[1]

November 2002

Abstract

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 status, where we distinguish between 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. Our analysis takes into account the transition from a steady state with cycles to one without cycles. Unlike previous studies, therefore, we can study how business cycles affect different groups of consumers. We conclude that the cost of cycles is positive but quite small on average. There are, however, large differences across groups: very poor consumers can gain a lot when cycles are removed, as can very rich consumers, whereas the majority of consumers—the "middle class"—sees very small negative effects from removing cycles. Inequality rises substantially upon removing cycles: the Gini coefficient on wealth moves from around 0.8 to 0.9.

[1]Per Krusell's affiliations are the University of Rochester, the Institute for International Economic Studies, and CEPR; Anthony A. Smith, Jr.'s affiliation is Carnegie Mellon University. This paper is a corrected and extended version of a paper published in the *Review of Economic Dynamics* 2:1, 245–272. The present paper contains drastically revised computations and analysis but also shares large parts of the published paper. For valuable comments, we wish to thank Burhan Kuruşçu, Fatih Guvenen, Zvi Hercowitz, Neil Wallace, an anonymous referee, and seminar participants at Carnegie Mellon, Georgetown University, Institute for International Economic Studies, M.I.T., Ohio State, Penn State, Rochester, Stanford, Toulouse, University of Montreal, University of Pennsylvania, Yale, the 1998 Annual Meetings of the Society for Economic Dynamics, and the Aiyagari Memorial Conference on Dynamic Macroeconomics.

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 the actual postwar U.S. consumption path and a path given by the trend of the actual path. 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 one percent; for example, for logarithmic utility the welfare gain is 0.008%.

If one wants to claim that Lucas's 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 only eliminate recessions. The second and third routes accept the notion that average output in some sense is unaffected. The second route stays within Lucas's general framework but argues that other assumptions about preferences or about the stochastic process governing the aggregate data 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 but 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, de facto 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. The no-cycle economy converges to a steady-state equilibrium in which aggregate variables do not move over time. Our goal is to evaluate the welfare effects of eliminating cycles on individual consumers. For this purpose, we cannot simply compare the steady-state equilibrium in the no-cycle economy with the stationary stochastic equilibrium in the economy with cycles, since individuals lose their identities in such a comparison. Instead, we remove aggregate shocks at a given point in time, solve for the equilibrium transition path toward steady state, and compare the welfare of each individual along this transition path to the welfare he would have obtained had the aggregate shocks remained.

We find that the costs of business cycles in our framework are small on average. However, behind this average we discover quite a large variation in how different groups are affected: some gain but some lose from eliminating aggregate cycles. The most important source of this variation is the change in the nature of the idiosyncratic employment process that we hypothesize. In the presence of aggregate as well as idiosyncratic risk, one needs to take a stand on how the removal of aggregate risk influences the risk facing an individual. We propose a procedure which follows that used by Lucas in his removing aggregate risk: we propose to "integrate out" the aggregate risk from the individual's employment process. By integrating the aggregate risk out of an income process we mean averaging over the aggregate states conditional on each idiosyncratic state. However, since individual employment is correlated with the aggregate state—when the aggregate state is good, the employment rate is high, and so are each individual's chances to find a job—the individual (employment) variable is not fully idiosyncratic, which makes integration a nontrivial task. With correlation, one thus first needs to construct a pure process for idiosyncratic luck that, by definition, is uncorrelated across individuals. Conditional on every realization of this new idiosyncratic variable, one can then integrate out the influence of aggregate risk on individual employment. Following this procedure, we find that the removal of aggregate risk lowers individual employment risk by about 16% in the long run.

With less idiosyncratic risk, two quantitatively significant implications for welfare follow. First, the poorest consumers can gain up to several percentage points in consumption equivalents from eliminating cycles. This contrasts Lucas's numbers, which are several orders of magnitude lower. Second, due to the lower income risk, the amount of precautionary saving in the economy falls. In the closed economy that we study, this raises the interest rate. This effect is small but nevertheless significantly raises the welfare of the very richest, who own very large amounts of wealth; the wealth distribution in the initial state reproduces the observed Gini coefficient for wealth and thus has (a small number of) very wealthy individuals. This effect on the welfare of the very richest also amounts to several percentage points of consumption equivalents. The middle class, in contrast, sees an improvement because of the lowering of risk, but it is very well insured in utility terms, so this effect is almost nil. Moreover, the middle class, of which a typical agent is employed, sees a fall in the wage, and as a result the middle class—in total around 65% of the agents—experience a welfare loss from eliminating aggregate cycles.

In addition to calculating welfare costs across individuals who differ in their economic status at the moment of eliminating cycles, we obtain striking implications for long-run inequality. The Gini coefficient for wealth distribution in the steady state without aggregate cycles is over 0.9, which can be compared to an initial average Gini coefficient of about 0.8 in the economy with cycles. That is, although the Gini coefficient for the earnings distribution is now lower, the wealth Gini coefficient goes up significantly. Behind this result is the assumption that discount rates differ across consumers: as there is less risk, consumers with different discount rates tend to corners, thus making wealth more dispersed. In particular the poorest can afford to become even poorer, given that their income risk is less severe and their discount rates tend to be significantly above the interest rate: they "want to" become poorer. For example, the number of households with negative assets goes from 11% to 31%. On the other hand, the very richest become even richer, due to both the increase in interest rates, which propagates their wealth, and a tendency to save even more: these agents are the most patient in the population. These effects on inequality are long-run effects but about half of the increase in the Gini coefficient appears within 10 years.

Not only does the new steady state, i.e., the long-run state of the economy without aggregate uncertainty, have sharply higher inequality but it also delivers significantly lower average welfare than the economy with cycles. This result has two origins. First, since there is less need for insurance and capital is lower, consumption will be lower on average. Second, there is much more inequality, which pushes average welfare downward. Of course, agents choose to have very different wealth levels and the benefits of these choices were realized during the transition—when the agents who end up poor in the steady state had a high level of consumption. In sum, the welfare gains from eliminating cycles all take place along the transition path, and they are in part associated with consuming some of the initial capital stock.

The first paper to introduce consumer heterogeneity for the purpose of studying the effects of eliminating business cycles is 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), Beaudry and Pages (1997), Gomes, Greenwood, and Rebelo (1998), Storesletten, Telmer, and Yaron (2001), and Krebs (2002). Atkeson and Phelan discuss the connection between aggregate and idiosyncratic risk, and they suggest as a serious possibility that the elimination of aggregate risk does not affect individual risk at all. They do not analyze a calibrated dynamic model, but focus 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 (1997) study idiosyncratic wage risk that worsens in recessions: so-called reallocation shocks.[2] They assume 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 only disappears slowly over time.[3] Since layoffs occur more frequently during recessions, they argue that cycles lead to an increase in both the variance and the persistence of idiosyncratic risk. Beaudry and Pages obtain higher costs of business cycles. However, their findings are based on the assumption that there is no idiosyncratic risk at all in the

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

[3]Empirical support for this can be found in Bils (1995) and elsewhere.

economy without cycles, and that workers in the economy with cycles cannot save to insure against wage risk. Along similar lines, Storesletten, Telmer, and Yaron (2001) as well as Krebs (2002) arrive at different results by considering a life-cycle, so that temporary shocks can have a larger impact (in the former paper), more permanent shocks, so that self-insurance is less effective (in both papers), and growth effects, using an "Ak" setup (in the latter paper). These authors consequently find larger effects of eliminating business cycles.

Gomes, Greenwood, and Rebelo (1998) argue—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.

In the representative-agent 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. Finally, Barlevy (2001) studies how fluctuations may have effects on growth by modeling the R&D process. He finds that the welfare gains from eliminating cycles can be up to two orders of magnitude larger than those Lucas found.

Because the effects of eliminating cycles are quite complex in our general setup, in Section 2.3 we first provide a comprehensive discussion of both partial and general 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.

2 Preliminaries

In Section 2.1, we first briefly discuss different routes one might take in answering our main question. We then discuss our theoretical model framework. In Section 2.2 we first lay out a slightly simplified version of our general model. The following section, Section 2.3, then describes a two-period model which is constructed to capture—in essence and notation—most of the ingredients in the multi-period 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 2.4 we discuss in detail how we eliminate cycles; Section 2.5 uses the two-period model to analyze the welfare consequences. Finally in Section 2.6 we comment on the relation between the question we study here and another issue of interest: the sharing of country-specific risk across countries.

2.1 Methodology

Lucas's (1987) model economy is very simple: consumption is exogenous and there is only one shock—one to the aggregate consumption process. In this economy, Lucas views the elimination of cycles as simply setting the shock to zero (its unconditional mean). To allow for heterogeneity, one way to extend what Lucas did is to use data on individual consumption. In particular, postulate and estimate the dependence of individual consumption on a purely idiosyncratic component and on aggregate variables. Next, for any value of the idiosyncratic component, take the average across the aggregate variables: this delivers a new consumption process. Finally, evaluate individual utility given this new process. We did not follow this procedure for two reasons. First, the procedure requires a long enough panel of individual consumption data that one can reliably estimate a process for individual consumption which identifies the aggregate from the idiosyncratic component as well as delivers an accurate assessment of the serial correlation properties of the shocks to individual consumption. Existing data do not grant this possibility. Second, this kind of calculation tends to underestimate the costs of cycles: if one instead models the randomness the agents are subject to, it may be better for them to change consumption in some other way than just averaging it across the aggregate states. That is, the utility value of not having aggregate cycles is underestimated. Although this is a problem also in Lucas's analysis, it is likely more quantitatively important here, where individual consumption volatility is much higher than in a representative-agent setup calibrated to aggregate data.

An alternative is to use individual income—for which data is arguably more reliable—and employ a model to infer consumption by assuming rational behavior given a certain set of asset markets. A first step would be to estimate a wage process and a process for asset returns and then to compare the utility outcome for a rational agent facing these processes with one where the same agent faces the same processes with their aggregate components removed. A second step would be to add an equilibrium component to the analysis, i.e., to also model where wages and rental rates come from. Since savings likely change as a result of eliminating risk, this seems a potentially important channel not to forget, at least if one believes that the economy is closed. Moreover, labor supply could change, leading to changes in the wage rate.

In this paper we follow a simple version of the latter procedure: we only model idiosyncratic differences in employment (and not in wage per hour worked), we assume that labor supply is inelastic, and we assume that all agents face the same return on saving. The asset structure is simple: there is only one asset—aggregate capital—and an exogenous borrowing constraint. The aggregate shock is modelled as exogenous changes in aggregate productivity and labor demand, and we study the general equilibrium effects for all different consumers of replacing the latter shocks with their conditional means.

2.2 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 (a bar 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.[4] Aggregate output, including undepreciated capital, can be used to either consume or invest.

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 both on 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.[5] There is, however, one asset market. 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, Krusell and Smith (1998), we study a very similar model with a second,

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

[5] 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.

riskless asset, and we find that the allocations are very similar whether or not the second asset is present. It should be kept in mind that the present calculations of the welfare gains from eliminating cycles are in all likelihood overestimates, as a second asset would allow for more risk sharing in the economy with cycles.

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,

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

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

2.3 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

$$\begin{aligned} c'_{0g} &= k'R'_g + g' \\ c'_{0b} &= k'R'_b + g' \\ c'_{1g} &= w'_g + k'R'_g \\ c'_{1b} &= w'_b + k'R'_b. \end{aligned}$$

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

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

2.4 The elimination of cycles

The heterogeneous-agent model we use here allows an indirect way—given the assumptions underlying the model—of deducing how individuals' consumption processes depend on aggregates. However, our modelling 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's 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. Perhaps the resulting welfare calculations represent upper bounds on the welfare gains, then, since presumably any policy measures designed to stabilize carry welfare-reducing distortions with them.

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.

2.4.1 The impact on aggregate variables

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 adjust for this difference in the final welfare figures.

2.4.2 The impact on individual employment

Second, as regards the individual-specific variables, 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 do Beaudry and Pages (1997), 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) suggest that one useful principle here is to remove the correlation between the idiosyncratic shocks of different individuals, leaving each individual's shock process unchanged. Atkeson and Phelan (1994) point out that the 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 discuss 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 two independent random variables, $y = g(i, z)$, where i is an idiosyncratic shock and z is an aggregate shock. The assumption that the two shocks are independent amounts to a definition of "idiosyncratic"; the densities are denoted $f_i(i)$ and $f_z(z)$, respectively. We then identify the idiosyncratic shock process in the absence of aggregate risk, $y^{w/o}(i)$, with the following:

$$y^{w/o}(i) = \int_z g(i, z) f_z(z) dz$$

for each i, with density $f_i(i)$.

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 jointly normal shocks, one individual-specific, but not necessarily fully idiosyncratic, shock ϵ and the aggregate shock z:

$$y = \epsilon + z.$$

We assume that the marginal distribution of each of these shocks is N(0,1) and that the covariance between the two shocks is ρ. If $\rho = 0$, so that ϵ and z are independent, we deduce that y is N(0,2). Then we obtain, using our integration principle, that

$$y^{w/o} = \epsilon,$$

which is N(0,1). Here, $y^{w/o}$ 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, $y^{w/o}$ would still be equal to $z + \epsilon$ in the absence of cycles, with z now being an idiosyncratic shock which is uncorrelated across agents.

If ϵ and z are correlated, integration requires first projecting ϵ onto z. This delivers $i + \rho z$, where i and z are now independent by construction and i has variance $\sigma_i^2 = 1 - \rho^2$, since we assumed that both ϵ and z were N(0,1). Now integration implies that

$$y^{w/o} = i,$$

which is N(0,1-ρ^2): this process has lower variance than in the case where ϵ and z were correlated.

In our model framework, the individual-specific income process y depends crucially on the employment process ϵ, which is not, in general, independent of z, like in the last example. In order to find the $\epsilon^{w/o}$—the employment process when there are no aggregate shocks—we therefore need to do the equivalent of the linear projection that was appropriate in that last example: we need to design a purely idiosyncratic variable i such that the employment/income outcome ϵ is a function of i and z.[6] It turns out that this can be done as follows: let i be uniform on [0,1], and define $\epsilon(i, z_g)$ to be 1 if $i \leq \pi_{1|g}$ and 0 otherwise and $\epsilon(i, z_b)$ to be 1 if $i \leq \pi_{1|b}$ and 0 otherwise. We will assume in this discussion that $\pi_{1|g} > \pi_{1|b}$, i.e., that ϵ and z are positively correlated.

Integration is now straightforward. If $i \leq \pi_{1|b}$, the individual is employed no matter what happens to the aggregate shock, so the integration is trivial: $\epsilon^{w/o} = 1$ for such values of i. Similarly, if $i > \pi_{1|g}$, the individual is unemployed no matter what: $\epsilon^{w/o} = 0$. Finally, if $\pi_{1|b} < i < \pi_{1|g}$, the individual is employed only if the aggregate state is good, which occurs with probability π_g; thus, integration for such values of i implies that $\epsilon^{w/o} = \pi_g \cdot 1 + (1 - \pi_g) \cdot 0 = \pi_g$. Thus, our new employment variable $\epsilon^{w/o}$ has the following 3-state distribution: 1 with probability $\pi_{1|b}$, π_g with probability $\pi_{1|g} - \pi_{1|b}$, and 0 with probability $1 - \pi_{1|g}$. Note that the new income variable thus has a different support—one more state—and that it is less risky: some probability mass has been moved from the extremes 0 and 1 into a middle state. In a dynamic economy, where individual employment is correlated over time, one can follow the same principles but it is quite a nontrivial affair to find the process for $\epsilon^{w/o}$. Suffice it to say here that this new process (i) will change nature—it will increase its support—as time evolves; (ii) will not be first-order Markov, but rather will be a function of two state variables which in turn are a function of all present and past values of i and evolve recursively; and (iii) will settle down to a stationary process with full support on [0,1]; Section 4 and Appendix III outline all the details.

Atkeson and Phelan, as mentioned, did not change the individual processes upon removing the aggregate shocks. In contrast, İmrohoroğlu (1989) did but she restricted the new employment process to be first-order Markov, something which is inconsistent with our integration principle. Similarly, the procedure used in Storesletten, Telmer, and Yaron (2001) is also inconsistent with the integration principle, although both these authors and İmrohoroğlu propose individual processes in their economies without aggregate cycles that have some intuitive appeal. Finally, Krebs (2002) does adhere to the integration principle in his recent paper.

2.5 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 several small steps. The uncertainty of aggregate origin faced by the agent has two parts: employment risk and price fluctuations. In Section 2.5.1 we first analyze the former alone in the case when prices (counterfactually) do not fluctuate; here, prices are constant and identical in the worlds with and without aggregate uncertainty. In Section 2.5.2 we then shut down the correlation between individual and aggregate employment, i.e., we study the effects of price fluctuations alone when the employment process is identical with and without aggregate uncertainty. We then move to general-equilibrium

[6] In the published version of this paper, we stated the integration principle as we did here (although with somewhat less detail) but we failed to apply it correctly to our economy. Our mistake amounted to treating ϵ as independent of z and as a result we obtained the exact same employment process with and without aggregate shocks. Thus, our results coincided with those that would followed from Atkeson and Phelan's procedure.

effects in Section 2.5.3, emphasizing how prices are influenced by the savings behavior of agents.

2.5.1 Partial equilibrium I: effects on the employment process, constant prices

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 $w' \equiv \pi_g w'_g + \pi_b w'_b$. In this subsection, we focus solely on how the elimination of cycles influence individual employment uncertainty, and we thus assume that the aggregate uncertainty economy actually has no randomness at all in prices: $R'_g = R'_b$ and $w'_g = w'_b$. We then analyze fluctuating prices in the next subsection, where we assume that the individual's employment process is not affected by eliminating cycles (i.e., as assumed by Atkeson and Phelan). These two sections thus provide partial insights from a partial-equilibrium perspective.

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

$$u(c) + \beta\left\{(1-\pi_{1|g})u(c'_0) + \pi_{1|b}u(c'_1) + (\pi_{1|g} - \pi_{1|b})u(c'_{0/1})\right\},$$

where $c'_0 = k'R' + g'$, $c'_1 = w' + k'R'$, and $c'_{0/1} = k'R' + \pi_g w' + \pi_b g'$. Here, note the third possible consumption outcome, $c'_{0/1}$, which occurs when $\pi_{1|g} - \pi_{1|b} > 0$; we shall assume that $\pi_{1|g} - \pi_{1|b} \geq 0$ for convenience.

Given our breakdown of individual employment in terms of i and z, it is useful to rewrite second-period utility under cycles so as to emphasize the independence of i and z:

$$(1-\pi_{1|g})[\pi_g u(c'_{0g}) + \pi_b u(c'_{0b})] + \pi_{1|b}[\pi_g u(c'_{1g}) + \pi_b u(c'_{1b})] + (\pi_{1|g} - \pi_{1|b})[\pi_g u(c'_{1g}) + \pi_b u(c'_{0b})],$$

where the consumption levels are defined as before.

With Atkeson and Phelan's way of eliminating cycles, the consumer would be indifferent as to which economy to live in. In contrast, with our assumption about the effects of eliminating cycles, whenever $\pi_{1|g} - \pi_{1|b} > 0$, the elimination of cycles strictly reduces uncertainty and the individual is better off. To see this, first note that $c'_{\epsilon g} = c'_{\epsilon b} = c'_\epsilon$ for $\epsilon \in \{0, 1\}$ because prices do not fluctuate. This makes the first two terms in the second-period utility identical across the two worlds, independently of the chosen k'. Second, for any k' chosen by the agent we see that the third term is higher without cycles, since a strictly concave u implies

$$u(k'R' + \pi_g w' + \pi_b g') > \pi_g u(k'R' + w') + \pi_b u(k'R' + g').$$

2.5.2 Partial equilibrium II: no effects on the employment process, fluctuating prices

Now consider the case when the aggregate uncertainty is directly payoff-relevant: $w'_g \neq w'_b$ and $R'_g \neq R'_b$. Here, let us ignore any effects on reducing individual employment risk and instead adopt Atkeson and Phelan's assumption so as to isolate the effects of prices.

It is useful now 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, since $\pi_g \neq \pi_{g|0}$ in our baseline calibration, so that the individual and aggregate states are correlated, the analysis becomes more nontrivial. We first consider the effects of fluctuating wages, keeping rental rates constant, and we thereafter make wages constant but let rental rates fluctuate.

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.[7] The intuition is 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 payoff from the aggregate state 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 $k' < 0$.[8] 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 $\pi_{0|g} > \pi_{0|b}$.

[7] 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 $\pi_{1|g} > \pi_{1|b}$.
[8] Similarly, utility without aggregate uncertainty is also higher if $\pi_{0|g} < \pi_{0|b}$ and $k' > 0$.

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

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

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. Unlike in the example with wage fluctuations, here the loss from eliminating risk results from concavity in the utility function, which is also the force underlying the gains from eliminating risk.

The role of ex-ante heterogeneity

The preceding analysis shows that a consumer's views on aggregate risk depend on the nature of the risk. These views, moreover, depend on individual characteristics: current wealth, employment status, and time preference.

Wealth is important because, in the absence of insurance markets and in the presence of a constraint on borrowing, it helps to insure the consumer against idiosyncratic risk. The preceding analysis of wage and rental rate fluctuations shows that a consumer's attitude towards 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 fluctuation 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 cycles vary nonmonotonically with wealth: the very poorest consumers benefit from the elimination of cycles, as they are very concerned about aggregate risk (especially rental rate risk); the very richest consumers benefit too, as wage income is irrelevant for them; but consumers with modest wealth do not benefit, as 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 employment, 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.

Finally, consider discount rate heterogeneity. In our two-period model, patient consumers will tend to be wealthy in the second period: they choose a large k'. Moreover, if one sees this economy as a snapshot from an infinite-horizon world, initial asset wealth is also strongly positively correlated with patience. As we saw above, wealth/large positive values for k' makes consumers care less about risk and more about rental rates. To the extent the latter fluctuate, we showed that rich consumers prefer the economy without fluctuations, although this effect is likely to be weak since it is zero with linear utility and these agents are well insured.

2.5.3 General equilibrium considerations

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).

How are prices affected?

Assuming that total savings do not change, 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 $\bar{k}'$ is the same across the two economies, if i_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 slightly higher in the economies without aggregate uncertainty even if we do not change the individual process as we eliminate cycles. For parallel reasons, average wage rates are slightly lower.

More importantly, when the amount of uncertainty changes for individuals, savings change, and this has implications for prices in our closed economy. For the parameterizations we use in our quantitative model below, we find that savings are significantly 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. As we shall see below, this effect will turn out to have quantitatively important implications.

In contrast, this effect is very weak if one does not change the employment process as cycles are eliminated, i.e., if one proceeds with the Atkeson and Phelan methodology.

The role of ex-ante heterogeneity

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

Employed workers will tend to lose from the downward effect on wages. For unemployed workers, this effect is also present, but it is weaker, especially for those agents with a high discount rate.

Discount rate heterogeneity is an important determinant of the response to changes in the average levels of wages and rental rates. For the most patient agents, whose discount rates will be the closest to the rental rates, an increase in the rental rate can have a large effect on savings. In general equilibrium, moreover, the smaller is the amount of uninsured individual risk, the larger are the effects of discount-rate differences on savings behavior of the different agents: the impatient borrow even more from the patient. As cycles are eliminated, therefore, one should see a larger equilibrium dispersion in wealth. In particular, this will make second-period *average* utility low, given that u is strictly concave. However, this effect on average utility comes from conscious choices of impatient consumers who choose to become poor in the second period.

In summary, the two-period model teaches us that (i) for a given agent, the elimination of stochastic movements in z and u do 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.

2.6 International risk-sharing

Suppose we consider several economies of the sort just studied and contemplate the possible sharing of country-specific aggregate productivity and employment risk across countries. As one possibility, one could study a large set of countries with no world uncertainty. In particular, one could imagine that there is such a world and that a given economy wants to examine whether or not to open up and join the international risk-sharing arrangement. Does the analysis in this paper have any implications for this question? The answer is that it depends on the specifics of how the international agreements would work.

Suppose that the international risk-sharing agreement does not allow capital to flow on net, so that there are only insurance payments flowing across borders: a country realizing a high z shock would pay and a country realizing a low z shock would receive a transfer. With a large number of countries and no world risk, these transfers would have to be actuarially fair in world equilibrium. Moreover, every consumer would be able to use the international insurance market to insure against the part of his employment risk which has aggregate origins. Would this deliver the case we look at where cycles are eliminated? It depends on the timing of the insurance contracting. If agents can insure against z_{t+1} knowing z_t and i_{t+1}, which is uncorrelated with z_{t+1} and therefore has no information about it, the answer is a qualified yes. The qualification is that the insurance against price risk may or may not be exactly the same in the two economies. Ignoring price risk, however, the agent would then ensure that his income at $t+1$ is precisely the probability-weighted average between the wage and the unemployment benefit at $t+1$, assuming that his idiosyncratic shock i_{t+1} is in the intermediate range where his employment fate would depend on the aggregate shock. If the agent's idiosyncratic shock were outside this range, he would not be interested in insurance, since the aggregate shock would not influence his employment outcome.

If, however, agents can only insure against z_{t+1} without knowing i_{t+1}, our experiment here would not correspond to the international risk sharing case, because now the agent would obtain 4 outcomes for income—he would occasionally ex post have to make a payment when unemployed (because the aggregate shock was good) and receive a payment when employed (because the aggregate shock was bad), thus obtaining less effective income insurance than the one that results from eliminating cycles.

Sensible international agreements would evidently also allow capital flows from high-productivity to low-productivity countries, if capital can flow quickly enough to allow this. Castro (2001) considers one such economy, where the international agreements do not allow direct insurance but rather international self-insurance: countries can borrow and lend subject to an exogenous borrowing constraint. Castro, however, assumes that individual risk-sharing within countries is perfect, and it is an open question how international capital flows would affect different groups within the economy through the risk channel.

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 2.2, 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 2.2, 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 thus

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

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:

$$c + k' = r(\bar{k}, 1 - u_z, z)k + w(\bar{k}, 1 - u_z, z)\, I_{\epsilon=3} + g_\epsilon + (1 - \delta)k'$$
$$\Gamma' = H(\Gamma, z, z')$$
$$k' \geq \underline{k},$$

where $I_{\epsilon=3} = 1$ 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 the resulting wealth distribution is similar to the data. For the calibrations we consider, the difference between consecutive values is roughly one-half of a percentage point.

We select aggregate shocks so that we approximate the movements in observed output fluctuations in 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 allowed borrowing is about 60-70% of average annual income.

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 the Krusell and Smith (1998) paper, 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 next period's state; row one 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).[9]

As is explained in detail in Krusell and Smith (1998), we selected 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.

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

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[10], and we select $g_1 = 0.038$ to match the left tail of the wealth distribution.

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's 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.[11]

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.[12] We also impose the requirement that

[10] 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.

[11] 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.

[12] 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 last period's value for z to be part of the current aggregate state. Thus, in our computation, we do need to make a distinction between those bad aggregate states which have persisted and those which come directly following a good aggregate state,

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.[13] 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 of 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. Our calibration intentionally exaggerates the dire consequences for employment dynamics of bad aggregate shocks.

With this calibration, we obtain the following average long-run wealth distributions:

The distribution of wealth							
	% of wealth held by top					Fraction with	Gini
	1%	5%	10%	20%	30%	wealth < 0	coefficient
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

The wealth distributions coming out of the two model calibrations are both quite similar to U.S. data (we use data based on Wolff (1994) and Díaz-Giménez, Quadrini, and Ríos-Rull (1996)). 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.

3.3 Model solution: approximate aggregation

We solve the model numerically 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

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.

[13]But note that, since recessions last only 8 periods on average, the average duration of an unemployment spell is relatively small: a little more than 2 periods, which is approximately the same average duration that obtains in the model with 2 idiosyncratic employment states.

class of models: most agents with nontrivial wealth holdings have almost identical savings propensities, making higher moments of the distribution of asset holdings play a very minor role in determining aggregates. 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. Appendix I describes the algorithm in detail.

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

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

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

in good times and

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

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

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, reproduces the stated coefficients and the reported R^2's 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:

$$\log \bar{k}' = 0.105 + 0.958 \log \bar{k}$$

$$R^2 = 0.99997 \quad \hat{\sigma} = 0.0094\%$$

in good times,

$$\log \bar{k}' = 0.116 + 0.952 \log \bar{k}$$

$$R^2 = 0.9997 \quad \hat{\sigma} = 0.028\%$$

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

$$\log \bar{k}' = 0.092 + 0.962 \log \bar{k}$$

$$R^2 = 0.99998 \quad \hat{\sigma} = 0.0084\%$$

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, as agents adjust their savings in the new shock-less aggregate environment toward the steady state. The movements in the exogenous aggregate variables can be separated into an expected and an unexpected part. Our experiment is to eliminate only the unexpected part, that is, to replace the stochastic z process with its conditional expectations as of the initial date. This leaves a deterministic movement in z and u which disappears in the long run.

The wage income process without cycles is significantly more complex than the one with aggregate cycles. Using the baseline case as an illustration, recall from the two-period economy that the employment process in the economy without cycles has a 3-point support: 0, 1, and a number strictly between 0 and 1, namely, the conditional probability of a good aggregate state. Intuitively, agents with a good enough individual shock i_2 would always be employed, no matter what z_2 was, and agents with a bad enough i_2 would always be unemployed, whereas agents with an intermediate value for i_2, for whom the aggregate state would determine the employment fate, will receive the average outcome in the economy without cycles.

In a three-period economy, idiosyncratic shocks i_t need to be constructed for every time period, and we assume that these are iid and uniform on [0,1]. Serial correlation in employment outcomes simply means that if the agent became employed in period 2 (either because of a low i_2 draw or a combination of an intermediate i_2 draw and a good aggregate shock), the cutoff for i_3 below which the agent would be employed in period 3 would be lower. In the aggregate-uncertainty economy, due to the Markov process we consider there would be 8 possible cutoffs. These would depend on whether the agent was employed in period 2 and on the outcomes of the aggregate state in periods 2 and 3 (recall that the probabilities of the two possible aggregate employment rates at $t+1$ depend both on z_t and on z_{t+1}). In the economy without cycles, conditional on a sequence of idiosyncratic shocks (i_2, i_3), the employment outcome in period 3 would be the average employment outcome—across all realizations of the aggregate sequences (z_2, z_3)—in the aggregate uncertainty world. Because there are 8 different cutoffs in period 3, defining 9 subintervals of [0,1], there are 9 possible values for employment, including 0, 1, and seven intermediate values. The latter correspond to the cases when the aggregate shock would have determined the agent's employment outcome in at least one of the periods. The number of possible period-3 employment outcomes for a given individual from the perspective of period 2 depends on i_2: it can be either 5 (if the agent had an extreme value for i_2, i.e., either employed for sure or unemployed for sure in that period, leaving only 4 cutoffs) or 9 (if the agent's employment outcome in period 2 would have depended on the aggregate state in that period).

By the logic above, there are for any individual and at any point in time t at most 9 possible realizations for employment next period. What these 9 values are, however, depends on the time period and on the agent's history, which in full includes the initial employment status and the sequence $\{i_2, i_3, \ldots, i_t\}$. Because of our Markov structure, this history can be summarized by a vector of dimension 2. Denote the joint probability that the agent would have been employed and the aggregate state would have been good in any period t by P_{gt}, and denote the joint probability that the agent would have been employed and the aggregate state would have been bad in the same period by P_{bt}. It is then possible to use i_{t+1} and the knowledge of t—which tells us the conditional probabilities that the aggregate states are good and bad at $t+1$—to update any pair (P_{gt}, P_{bt}) into a new pair $(P_{g,t+1}, P_{b,t+1})$. The exact form of this mapping is contained in Appendix III. Moreover, the employment outcome in period $t+1$ will simply be $1 \cdot P_{g,t+1} + 0 \cdot (1 - P_{g,t+1}) + 1 \cdot P_{b,t+1} + 0 \cdot (1 - P_{b,t+1})$, since these outcomes are disjoint, which equals $P_{g,t+1} + P_{b,t+1}$. As we move forward in time and there is a larger possible set of aggregate and individual histories, each corresponding to a probability vector, the number of possible values a wage can take increases and, as time approaches infinity, for any subinterval of [0,1], there is a positive probability of a wage outcome lying in that interval.

4.1 The baseline case: homogeneous unemployed

The stationary income process obtained for the agent in the economy without aggregate uncertainty generates an unconditional density as portrayed in the following two figures:

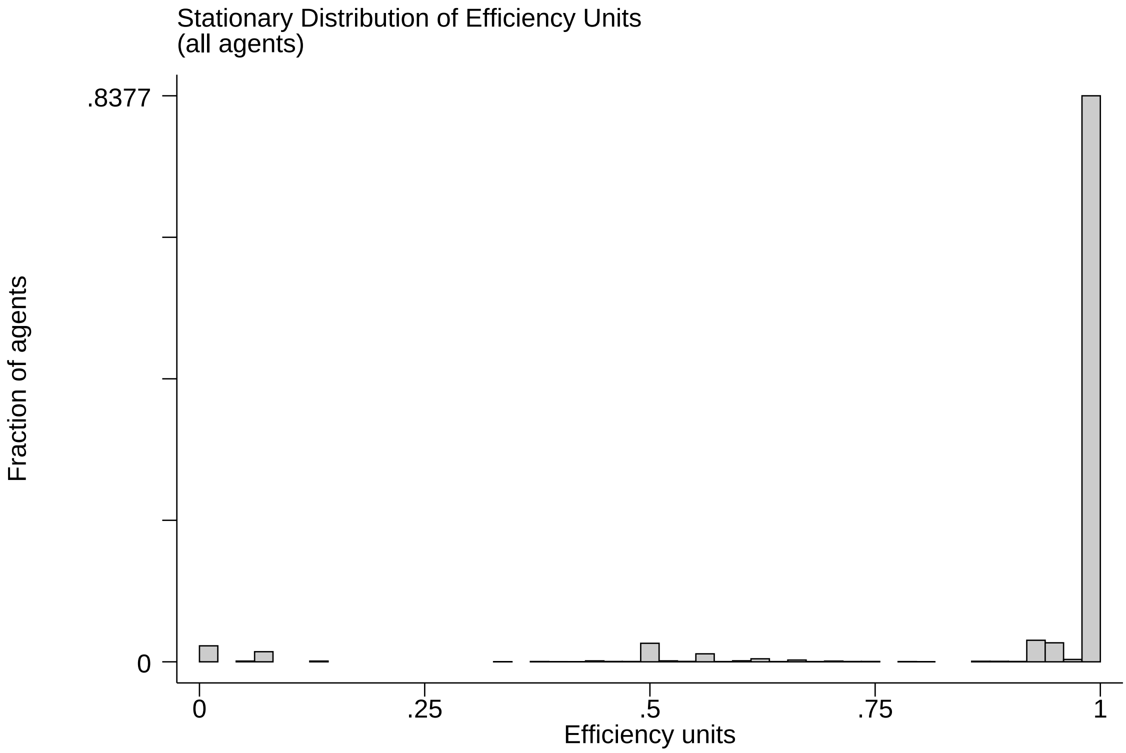

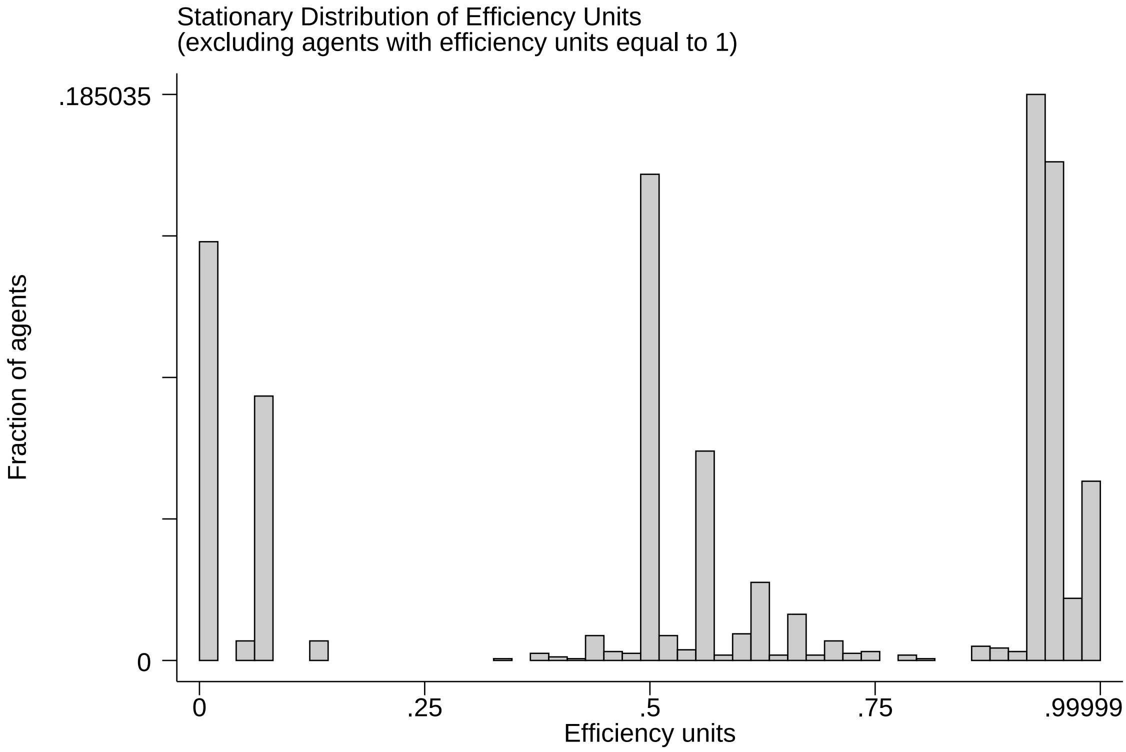

The second of these is just a more detailed picture that obtains when the values of "1" are eliminated. As can be seen from the graphs, the new process is quite concentrated on the values 1 and 0, but it also has significant probability mass on other numbers. The standard deviation of this process is about 16% lower than that of the process under aggregate uncertainty. Of course, this process only obtains asymptotically, and the wage process changes quite slowly.

With a 2-state process for employment—the agent is either employed or unemployed—the long-run welfare gain from eliminating cycles turns out to be negative: -0.66% of consumption across dates and states. To compute this 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 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 turns out to be negative for two reasons. First, as we have discussed earlier, the reduction in the amount of risk that individuals face reduces precautionary savings and hence reduces the amount of capital in the economy. This means that average consumption in the economy without cycles is lower than average consumption in the economy with cycles. Second, as we discuss in detail below, the reduction in individual risk leads to large increases in economic inequality. The cross-sectional variance of the logarithm of consumption, for example, increases by almost 10%, from an average of 0.064 in the economy with cycles to 0.07 in the economy with cycles (the Gini coefficient for consumption also increases by 10%, from 0.14 to 0.154). Since consumers have concave value functions, this spreading out of consumption reduces average welfare in the economy without cycles.

As we have noted earlier, the long-run welfare comparison is misleading since the identities of the consumers are lost. Indeed, turning to the differential effects across consumers, and thus, to our transition experiments, a very different picture of the welfare effects of eliminating cycles emerges. To begin, the following table contains summary measures of welfare changes across different groups.

Initial state		Fraction	Average utility gains in percentage consumption					
$\bar{k}$	z	gaining	All	$\epsilon = 1$	$\epsilon = 0$	$\beta =$ low	$\beta =$ middle	$\beta =$ high
11.2	z_g	0.363	0.089	0.083	0.222	0.218	0.005	0.631
11.2	z_b	0.400	0.110	0.089	0.297	0.266	0.025	0.632
12.3	z_g	0.330	0.087	0.085	0.140	0.226	0.002	0.624
12.3	z_b	0.364	0.096	0.087	0.175	0.260	0.010	0.619

The table contains 4 transition experiments: starting from two different initial capital stocks (each one is associated with a randomly drawn distribution of assets from the stationary stochastic process for this distribution under aggregate uncertainty) and from two different values of the aggregate shock. The results in the table reveal, first, that the *average* welfare gain from eliminating cycles is a little more than one magnitude larger than that computed by Lucas for the same period utility function: up to 0.09-0.11% of consumption from Lucas's 0.008%. Second, the distribution of welfare gains reveals substantial heterogeneity. For example, only a little over a third of the population even realize a gain; the rest *lose* from eliminating cycles. Summary statistics from the point of view of employment and patience type, as of the point in time when cycles are removed, are also displayed in the table. They reveal that unemployed agents lose on average 2–3 times more from cycles and that, across agents with different time preference rates, the most and least patient gain the most. In contrast, the middle group has a welfare gain that on average is similar to Lucas's number. The largest gainers are actually the most patient group, with numbers over 0.6%, i.e., close to two orders of magnitude larger than Lucas's representative-agent number.

The above table does not reveal who the losers from eliminating cycles are. We need a breakdown across wealth groups to find the losers, and the following table provides this kind of information for one of the four transition experiments. The tables corresponding to each of the other three transition experiments differ only slightly from this one; they are contained in the Appendix II.

Average Utility Gains by Wealth Group ($\bar{k} = 11.2$, $z = z_g$)

	Utility gain in percentage consumption							
	< 1	1–5	5–25	25–50	50–75	75–95	95–99	> 99
All	0.365	0.177	0.051	−0.055	−0.094	0.227	1.083	1.685
$\epsilon = 1$	0.240	0.153	0.045	−0.057	−0.095	0.227	1.083	1.686
$\epsilon = 0$	0.838	0.419	0.165	−0.002	−0.062	0.218	1.080	1.676

Here, a sharp U-shape across wealth levels appears. We see, in particular, that the 25th–75th percentiles in the wealth distribution are significant losers. For example, those in the 50th–75th percentile lose one order of magnitude more in consumption equivalents than Lucas's representative agent gains. Moreover, we see that the losses are larger for the employed than for the unemployed. These agents are not very vulnerable to risk, and they lose particularly from the lower wages that are due to lower aggregate savings.

The biggest gainers are richest group; the top percentile in wealth gain more than 1.5% in consumption equivalents. Clearly, the gains here derive from the increased interest rate. The poorest, represented as the bottom percentile here, gain between around 0.2% (the employed) and 0.8% (the unemployed); here, the diminished risk seems to be the reason why the gains are high.

Yet another cut of the heterogeneity in gains is compiled in the next table, where there is a focus on the groups at the extreme ends in the wealth distribution across different patience and employment statuses.

Utility Gains for Different Types of Agents ($\bar{k} = 11.2$, $z = z_g$)

	Wealth percentile						
Type of agent	constr.	0.005	0.05	0.5	0.95	0.995	0.999
$\epsilon = 1$, $\beta =$ low	0.533	0.446	0.289	0.083	0.428	1.204	1.518
$\epsilon = 1$, $\beta =$ middle	0.238	0.165	0.039	−0.092	0.670	1.491	1.808
$\epsilon = 1$, $\beta =$ high	0.065	0.021	−0.025	0.020	1.048	1.878	2.197
$\epsilon = 0$, $\beta =$ low	3.685	0.992	0.546	0.155	0.417	1.201	1.518
$\epsilon = 0$, $\beta =$ middle	2.785	0.612	0.243	−0.051	0.658	1.489	1.807
$\epsilon = 0$, $\beta =$ high	2.023	0.347	0.080	0.000	1.036	1.876	2.196

The effect on agents who are borrowing-constrained can be seen in this table; if they are impatient and unemployed, the reduction in risk amounts to close to a 4% increase in utility measured in consumption equivalents. This number is 500 times larger than the number provided by Lucas and it is an answer of sorts to the main query in this paper: are there agents in the population who would really gain a lot from the elimination of cycles? The answer is yes, if we think of 4% as a large number.

The table also shows that those in the top one-thousandth of the wealth distribution gain over 2% if they are patient. Neither the very poorest nor the very richest represent large groups. Especially the losers among the poor need to be at, or very close to, the borrowing constraint in order to lose significantly; just a little bit of wealth goes a long way toward lowering the utility losses from the risk. The gains among the richest are fall off somewhat less rapidly; there is a larger group of big winners in this group.

The next table shows the long-run wealth distribution in the U.S. data and in the economies with and without cycles.

The Distribution of Wealth

	Pct. of wealth held by top					Pct. with	Gini
	1%	5%	10%	20%	30%	wealth < 0	coefficient
Cycles	24	54	72	87	91	11	0.81
No cycles	26	60	79	95	99	31	0.90
Data	30	51	64	79	88	11	0.79

This table shows that inequality increases significantly; the interest rate goes up, making the rich richer, and the lowering of individual risk makes the poor less concerned about holding low levels of assets—they engage in less precautionary savings. In an economy without discount-factor heterogeneity, these effects would not appear nearly as strongly (or at all). The effect of removing uninsurable individual risk in an economy with discount-factor heterogeneity, at least of the persistent kind we consider here, is to move close to the complete-markets outcome, which we know will be rather extreme. It would not be degenerate here since the discount factors do vary stochastically, but close enough that the effects on inequality of eliminating cycles are very large.

The next three graphs show how different measures of inequality (the Gini coefficient for the wealth distribution, the fraction of consumers with negative wealth, and the fraction of wealth held by the top 20% of wealthholders) increase over time during the transition from the economy with cycles to the economy without cycles.

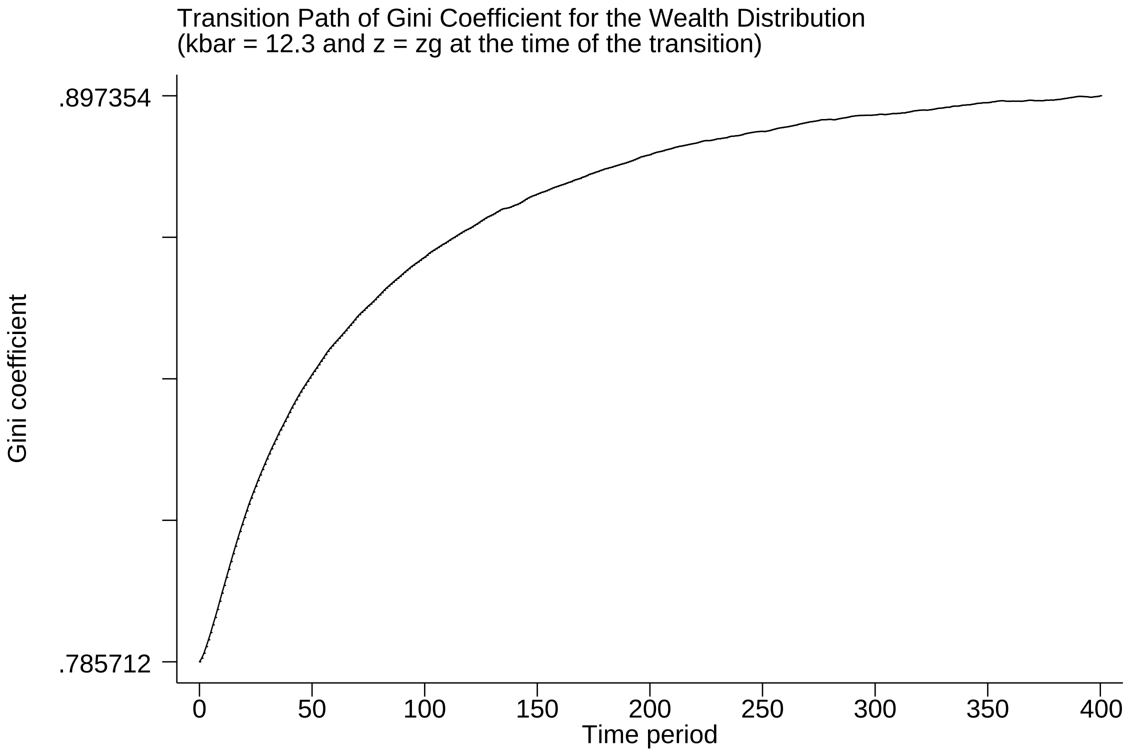
Transition Path of Gini Coefficient for the Wealth Distribution
(kbar = 12.3 and z = zg at the time of the transition)
.897354
.785712
Gini coefficient
0
50
100
150
200
250
300
350
400
Time period

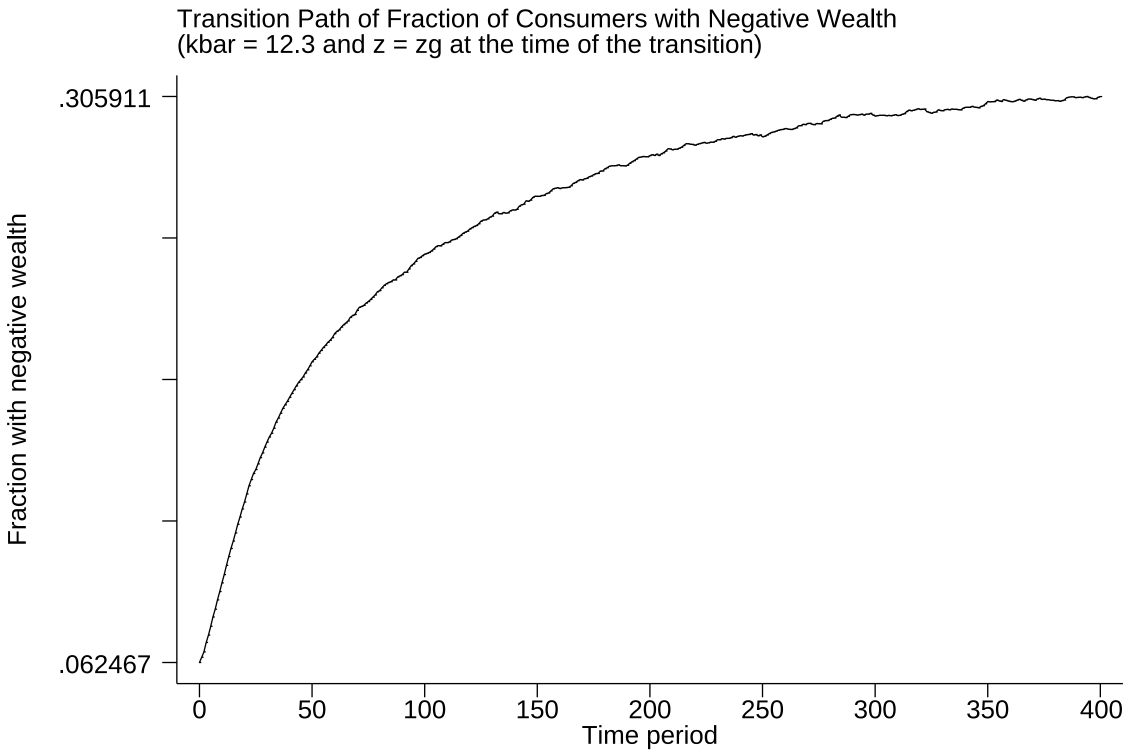
Transition Path of Fraction of Consumers with Negative Wealth
(kbar = 12.3 and z = zg at the time of the transition)
.305911
.062467
Fraction with negative wealth
0
50
100
150
200
250
300
350
400
Time period

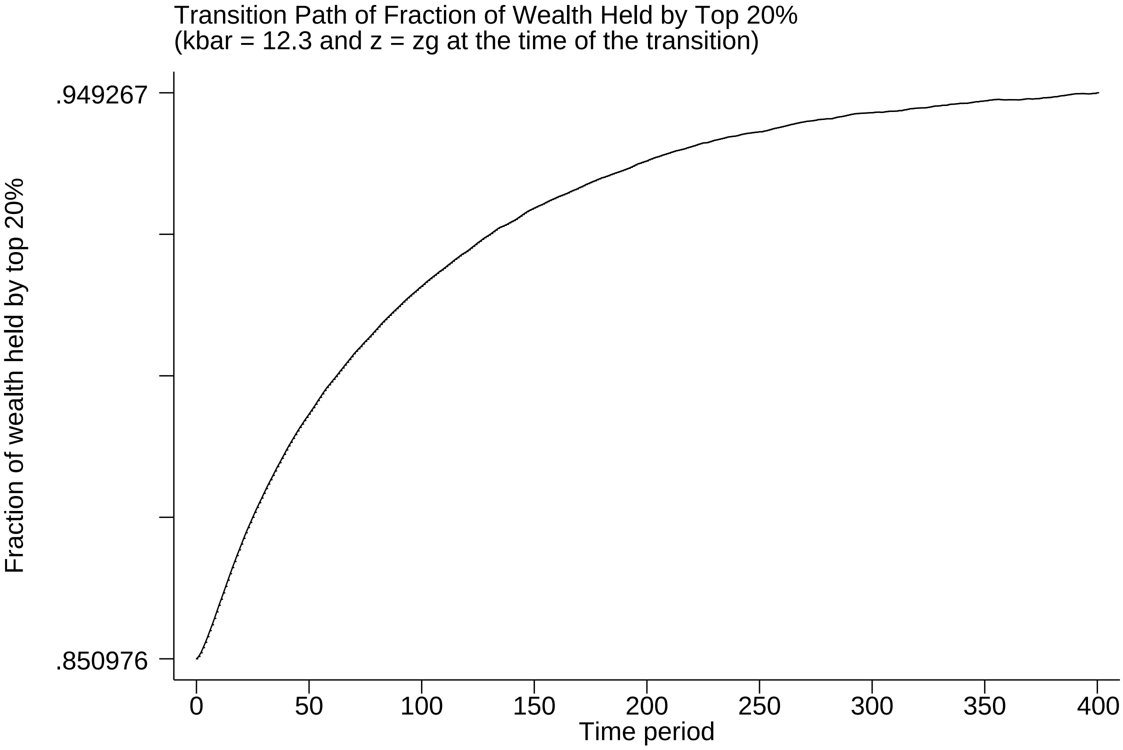

Not only does wealth inequality increase, but, as we discussed earlier, consumption inequality also increases, by roughly 10% (both the cross-sectional variance of log consumption and the Gini for consumption increase by about 10%). Wealth and consumption inequality increase in spite of the *decrease* in earnings inequality (the coefficient of variation of the earnings distribution is almost 20% lower in the economy without cycles).[14] In addition, *welfare* inequality also increases. We measure welfare inequality by asking each consumer how much better (or worse) off he would be (expressed, as before, in terms of a percentage consumption equivalent) if he were to be given the wealth and employment status of an "average" consumer: say, an employed consumer who holds per capita wealth. This metric avoids making welfare comparisons between consumers who have different preferences because of differences in discount rates. We use the cross-sectional standard deviation of these consumption equivalents to summarize the dispersion in welfare. This standard deviation increases by about 8%, from 17.3% on average in the economy with cycles to 18.7% in the economy without cycles.

Finally, the next two graphs show the evolution of the total capital stock starting from above and from below the steady state, respectively. In each graph, there are two curves. The top curve represents the transition path that results if the individual employment process are not altered as cycles are removed—the Atkeson and Phelan approach. These graphs give higher savings for precautionary reasons, since individual income has higher variation than for the economy where we integrate out the aggregate risk.

[14]By contrast, the Gini for income, which includes both labor earnings and interest income, is roughly the same (about 0.25).

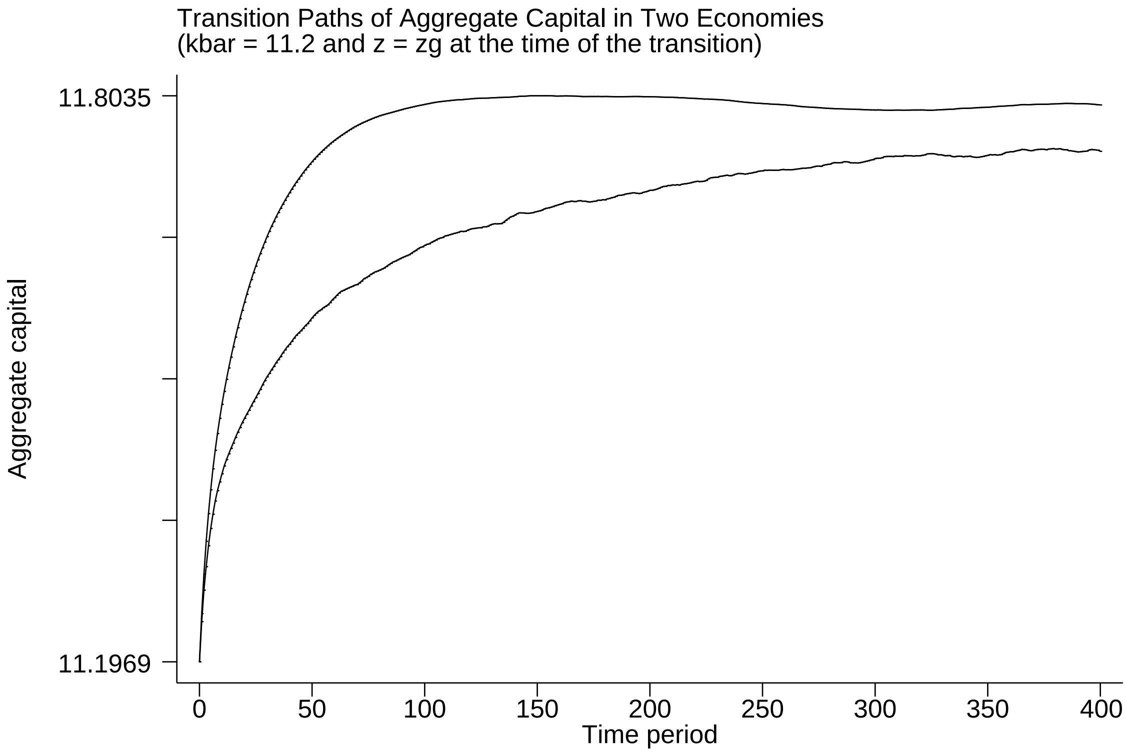

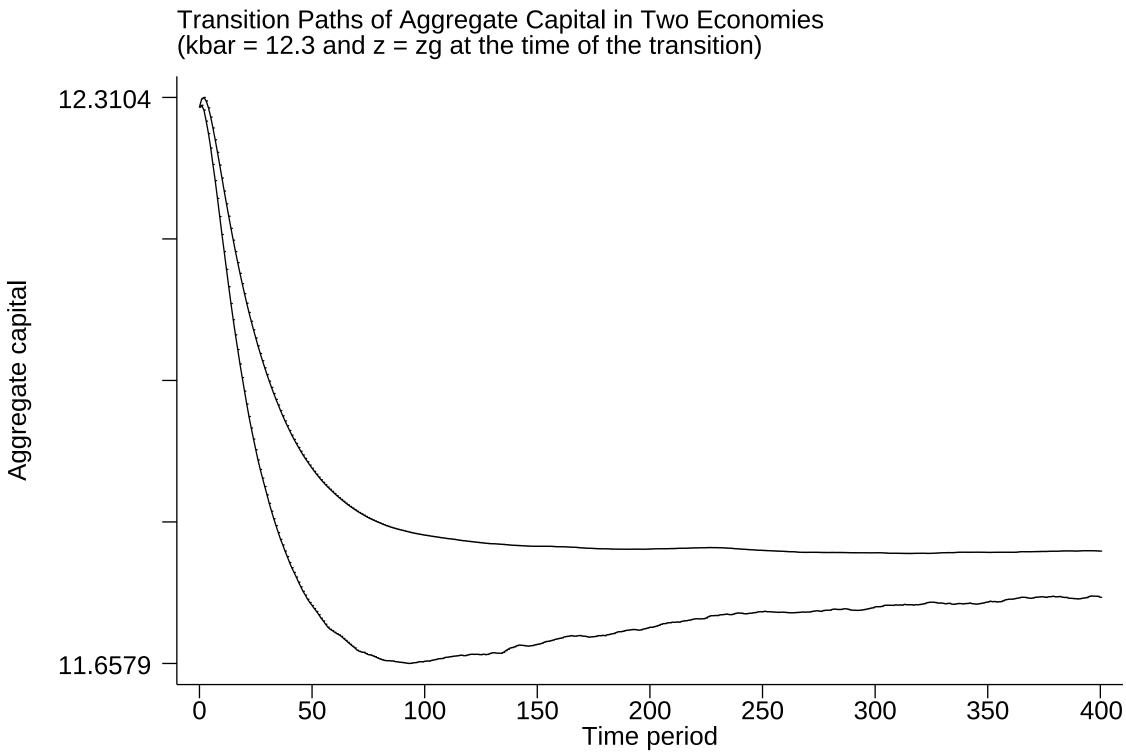

We see that the transition is more rapid coming from above, even nonmonotone, as it undershoots the new steady state and then comes back up. Along these paths, agents consume the large amounts of initial capital, amounts that are "excessive" given the new and less risky individual employment process. Starting from a lower-than-steady-state capital stock, the transition path upward is much slower.

4.2 The calibration with short- and long-term unemployment

To be completed.

5 Summary

We find that the average gains from eliminating cycles are significantly higher than in Lucas's (1987) representative-agent calculations. However, these gains are still quite low in absolute terms: about 0.1% of average consumption. We do find significant differences among agents in how they are affected by eliminating cycles. The largest gains are recorded in the poorest, most impatient group: close to 4% of average consumption. The gain for this group is mainly from lowering the risk: the employment process now has less risk after the elimination of the part of that risk that had aggregate origins, and there is also less wage (and interest) rate risk, since the aggregate productivity does not fluctuate. The very richest also gain substantially: over 2%. The effect on the richest is due to the increase in the interest rate that comes from a fall in precautionary saving, which in turn comes from the lower individual income risk. On the other hand, close to 65% *lose* from eliminating cycles. This large middle class typically are reasonably well insured in terms of wealth—so they do not benefit much from a reduction in risk—and they are employed. The disadvantage for the employed is that their wages fall as a result of a fall in aggregate savings. Finally, we note that the effect on long-run wealth inequality from eliminating aggregate uncertainty is rather drastic in our economy: the Gini coefficient goes up from just under 0.8 to around 0.9. This effect is due to the discount factor heterogeneity that underlies the realism of our initial wealth distribution: with less individual uninsured risk, wealth holdings of agents with different time preference rates become more extreme as the patient lend to the impatient early on.

6 References

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

Atkeson, A., and C. Phelan (1994), Reconsidering the Costs of Business Cycles with Incomplete Markets, in: S. Fischer and J. Rotemberg, eds., *NBER Macroeconomics Annual 1994*, 187–207.

Attanasio, O., and S. Davis (1996), Relative Wage Movements and the Distribution of Consumption. *Journal of Political Economy* 104:6, 1227–1262.

Barlevy, Gadi (2001), The Costs of Business Cycles under Endogenous Growth. Manuscript.

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

Bils, M. (1985), Real Wages over the Business Cycle: Evidence from Panel Data. *Journal of Political Economy* 93, 668–89.

Castro, R. (2001), Economic Development and Growth in the World Economy. Manuscript.

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 N. Kocherlakota (1997), A Microfoundation for Incomplete Security Markets. Federal Reserve Bank of Minneapolis Working Paper #577.

Díaz-Giménez, J., V. Quadrini, and J.-V. Ríos-Rull (1996), Measuring Inequality: Facts on the Distribution of Earnings, Income, and Wealth. *Federal Reserve Bank of Minneapolis Quarterly Review* (forthcoming).

Dolmas, J. (forthcoming), Risk Preferences and the Welfare Cost of Business Cycles, *Review of Economic Dynamics*.

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.

Krebs, T. (2002), Growth and Welfare Effects of Business Cycles in Economies with Idiosyncratic Human Capital Risk. Manuscript.

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

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

Krusell, P. & A.A. Smith, Jr. (1998), Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy* 106:5, 867–896.

Lucas, R.E., Jr. (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. & E.C. Prescott (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., C. Telmer, and A. Yaron (2001), The Welfare Costs of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk, *European Economic Review*, 45:7, 1311–1339.

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

Appendix I: Computation

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 475 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 475 time periods because both aggregate employment and the aggregate productivity shock have settled down to their steady-state values after 125 periods and because the transition path of aggregate capital—which converges much more slowly than aggregate employment and productivity—is, in some cases, nonmonotonic during the first 125 periods of transition (and monotonic thereafter).

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 backwards starting in period 125 to calculate agents' decision rules during the first 125 periods of transition. We therefore calculate 125 different decision rules and value functions, one for each period. Specifically, in period 125, the agent takes as given the current values of aggregate capital, employment, and productivity and the value of aggregate capital in period 126. (Recall that the consumer takes as given a time path for aggregate capital. In addition, at each time period in the backwards 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 125, 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 126.) The solution to the agent's problem in period 125 is a decision rule and a value function, each of whose arguments are the individual state variables. We then move back to period 124, in which period the agent takes as given the values of aggregate capital, employment, and productivity in period 124 as well as the period-125 value function (this function determines the future value of the agent's savings decision in period 49). We continue iterating backwards 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 and the period-t value function do depend implicitly (by means of the backwards 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.

Appendix II: More tables

Average Utility Gains by Wealth Group ($\bar{k} = 11.2$, $z = z_b$)

	Utility gain in percentage consumption							
	< 1	1–5	5–25	25–50	50–75	75–95	95–99	> 99
All	0.845	0.279	0.082	−0.043	−0.088	0.232	1.078	1.675
$\epsilon = 1$	0.276	0.172	0.057	−0.051	−0.091	0.232	1.075	1.677
$\epsilon = 0$	1.516	0.686	0.267	0.044	−0.046	0.236	1.110	1.656

Average Utility Gains by Wealth Group ($\bar{k} = 12.3$, $z = z_g$)

	Utility gain in percentage consumption							
	< 1	1–5	5–25	25–50	50–75	75–95	95–99	> 99
All	0.269	0.179	0.041	−0.059	−0.079	0.234	1.025	1.558
$\epsilon = 1$	0.247	0.168	0.038	−0.061	−0.080	0.233	1.027	1.557
$\epsilon = 0$	0.503	0.327	0.105	−0.010	−0.053	0.244	0.975	1.594

Average Utility Gains by Wealth Group ($\bar{k} = 12.3$, $z = z_b$)

	Utility gain in percentage consumption							
	< 1	1–5	5–25	25–50	50–75	75–95	95–99	> 99
All	0.413	0.216	0.061	−0.050	−0.074	0.230	1.014	1.564
$\epsilon = 1$	0.276	0.179	0.047	−0.058	−0.078	0.230	1.018	1.557
$\epsilon = 0$	0.836	0.452	0.183	0.020	−0.035	0.229	0.981	1.645

Utility Gains for Different Types of Agents ($\bar{k} = 11.2$, $z = z_b$)

	Wealth percentile						
Type of agent	constr.	0.005	0.05	0.5	0.95	0.995	0.999
$\epsilon = 1$, $\beta =$ low	0.619	0.497	0.307	0.087	0.433	1.180	1.515
$\epsilon = 1$, $\beta =$ middle	0.305	0.206	0.054	−0.087	0.668	1.459	1.796
$\epsilon = 1$, $\beta =$ high	0.108	0.044	−0.020	0.016	1.032	1.831	2.170
$\epsilon = 0$, $\beta =$ low	3.828	1.665	0.748	0.193	0.422	1.178	1.515
$\epsilon = 0$, $\beta =$ middle	2.916	1.159	0.406	−0.023	0.655	1.456	1.796
$\epsilon = 0$, $\beta =$ high	2.128	0.763	0.189	0.000	1.019	1.829	2.170

Utility Gains for Different Types of Agents ($\bar{k} = 12.3$, $z = z_g$)

	Wealth percentile						
Type of agent	constr.	0.005	0.05	0.5	0.95	0.995	0.999
$\epsilon = 1$, $\beta =$ low	0.560	0.346	0.242	0.098	0.390	1.086	1.352
$\epsilon = 1$, $\beta =$ middle	0.262	0.087	0.008	−0.077	0.629	1.369	1.637
$\epsilon = 1$, $\beta =$ high	0.089	0.003	−0.004	0.043	1.009	1.757	2.026
$\epsilon = 0$, $\beta =$ low	1.558	0.627	0.408	0.163	0.381	1.084	1.351
$\epsilon = 0$, $\beta =$ middle	1.076	0.312	0.134	−0.041	0.618	1.367	1.636
$\epsilon = 0$, $\beta =$ high	0.714	0.131	0.034	0.024	0.998	1.755	2.025

Utility Gains for Different Types of Agents ($\bar{k} = 12.3$, $z = z_b$)

	Wealth percentile						
Type of agent	constr.	0.005	0.05	0.5	0.95	0.995	0.999
$\epsilon = 1$, $\beta =$ low	0.637	0.368	0.255	0.100	0.393	1.088	1.388
$\epsilon = 1$, $\beta =$ middle	0.323	0.104	0.019	−0.074	0.627	1.364	1.667
$\epsilon = 1$, $\beta =$ high	0.128	0.009	−0.005	0.038	0.996	1.741	2.046
$\epsilon = 0$, $\beta =$ low	3.035	0.869	0.532	0.197	0.384	1.086	1.388
$\epsilon = 0$, $\beta =$ middle	2.277	0.508	0.232	−0.017	0.615	1.362	1.667
$\epsilon = 0$, $\beta =$ high	1.641	0.269	0.086	0.021	0.983	1.739	2.045

Appendix III: The integration principle applied to the baseline model

The 4x4 Markov transition matrix for (ϵ, z), in its calibrated form, reads

$$\begin{pmatrix} 0.8507 & 0.1159 & 0.0243 & 0.0091 \\ 0.1229 & 0.8361 & 0.0021 & 0.0389 \\ 0.5833 & 0.0313 & 0.2917 & 0.0938 \\ 0.0938 & 0.3500 & 0.0313 & 0.5250 \end{pmatrix},$$

with ordering (e, g), (e, b), (u, g), (u, b). This implies probabilities of finding a job, conditional on last period's employment status and on the aggregate shocks in the current and in the last period. These job-finding probabilities can be ranked from least to most luck as follows: the probability of becoming employed next period ($\epsilon' = 1$) is,

1. conditional on $\epsilon = 0$, $z = g$, and $z' = b$, $0.0313/0.1250 = 0.2504 \equiv \bar{i}_1$;
2. conditional on $\epsilon = 0$, $z = b$, and $z' = b$, $0.3500/0.8750 = 0.4000 \equiv \bar{i}_2$;
3. conditional on $\epsilon = 0$, $z = g$, and $z' = g$, $0.5833/0.8750 = 0.6666 \equiv \bar{i}_3$;
4. conditional on $\epsilon = 0$, $z = b$, and $z' = g$, $0.0938/0.1250 = 0.7504 \equiv \bar{i}_4$;
5. conditional on $\epsilon = 1$, $z = g$, and $z' = b$, $0.1159/0.1250 = 0.9272 \equiv \bar{i}_5$;
6. conditional on $\epsilon = 1$, $z = b$, and $z' = b$, $0.8361/0.8750 = 0.9555 \equiv \bar{i}_6$;
7. conditional on $\epsilon = 1$, $z = g$, and $z' = g$, $0.8507/0.8750 = 0.9722 \equiv \bar{i}_7$; and
8. conditional on $\epsilon = 1$, $z = b$, and $z' = g$, $0.1229/0.1250 = 0.9832 \equiv \bar{i}_8$.

Thus, our idiosyncratic shock i_t can end up in 9 relevant subintervals, defined by the cutoff values $\bar{i}_1$–$\bar{i}_9$, in each period. Let us take an example: if $i \in [\bar{i}_6, \bar{i}_7)$, the agent would have been employed currently only if the aggregate and idiosyncratic shocks leading up to the last period made the agent employed in that period and the current aggregate state would have been bad (independently of what the state was last period).

For a realized sequence of idiosyncratic shocks $\{i_s\}_{s=1}^t$ one can then compute an average employment outcome in period t by brute-force averaging across all $\{z_s\}_{s=1}^t$ sequences (appropriately weighted by probabilities): given each such sequence of aggregate shocks, together with an employment status in period 0, the employment outcomes in all periods up to and including t are known: they follow applying the cutoff values above in every time period.

The resulting employment process will have long memory in terms of the idiosyncratic shocks: one generally needs to know *all* prior values of i_s, $s < t$, in order to know what the average employment outcome is at t. However, it is possible to represent the new employment process recursively. To this end, let P_{gt} denote the probability (or fraction of the time) that, among all possible outcomes of the aggregate process, (i) the individual would have been employed in time t, given his initial (time-0) employment status and an initial (time-0) value for the aggregate state AND (ii) the aggregate state at time t would have been good. Similarly define P_{bt} as the probability that the agent would have been employed in t jointly with a bad aggregate state in that period. Letting π_t denote the probability of a good aggregate state in period t given z_0, these definitions imply that $\pi_t - P_{gt}$ is the probability that the agent would have been unemployed in t jointly with a good aggregate state in t and similarly that $1 - \pi_t - P_{bt}$ is the probability that the agent would have been unemployed in t jointly with a bad aggregate state in t. The key insight now is that the variables $P_t \equiv (P_{gt}, P_{bt})$ summarize all there is to know from history in order to know the expected (average) value for employment in period $t+1$ given a value for i_{t+1}. I.e., P_t summarizes all the relevant knowledge about $\{i_1, i_2, \ldots, i_t\}$. This representation is possible because the joint underlying process for employment and the aggregate state is first-order Markov.

The recursive structure needs to update P_t into P_{t+1} given a value for i_{t+1}, and it needs to assign a value for the average employment outcome in period $t+1$ conditional on the state variable P_t summarizing the individual's idiosyncratic history and i_{t+1}, the new idiosyncratic shock. The latter is easy: the average value of employment across the aggregate shock outcomes will be $\epsilon_{t+1}^{w/o} = P_{g,t+1} + P_{b,t+1}$, because g and b are disjoint outcomes.

The updating from P_t into P_{t+1} given i_{t+1} requires more care. The variable i_{t+1} can fall into any one of the 9 regions defined by the cutoffs $\bar{i}_1$–$\bar{i}_9$ above, and each case implies a different joint employment/aggregate state probability in period $t+1$. These are:

1. $i_{t+1} \in [0, \bar{i}_1)$: $P_{g,t+1} = \pi_{t+1}$ and $P_{b,t+1} = 1 - \pi_{t+1}$;
2. $i_{t+1} \in [\bar{i}_1, \bar{i}_2)$: $P_{g,t+1} = \pi_{t+1}$ and $P_{b,t+1} = P_{gt}\pi_{b|g} + (1-\pi_t)\pi_{b|b}$
3. $i_{t+1} \in [\bar{i}_2, \bar{i}_3)$: $P_{g,t+1} = \pi_t$ and $P_{b,t+1} = P_{gt}\pi_{b|g} + P_{bt}\pi_{b|b}$;
4. $i_{t+1} \in [\bar{i}_3, \bar{i}_4)$: $P_{g,t+1} = P_{gt}\pi_{g|g} + (1-\pi_t)\pi_{g|b}$ and $P_{b,t+1} = P_{gt}\pi_{b|g} + P_{bt}\pi_{b|b}$;
5. $i_{t+1} \in [\bar{i}_4, \bar{i}_5)$: $P_{g,t+1} = P_{gt}\pi_{g|g} + P_{bt}\pi_{g|b}$ and $P_{b,t+1} = P_{gt}\pi_{b|g} + P_{bt}\pi_{b|b}$;
6. $i_{t+1} \in [\bar{i}_5, \bar{i}_6)$: $P_{g,t+1} = P_{gt}\pi_{g|g} + P_{bt}\pi_{g|b}$ and $P_{b,t+1} = P_{bt}\pi_{b|b}$;
7. $i_{t+1} \in [\bar{i}_6, \bar{i}_7)$: $P_{g,t+1} = P_{gt}\pi_{g|g} + P_{bt}\pi_{g|b}$ and $P_{b,t+1} = 0$;
8. $i_{t+1} \in [\bar{i}_7, \bar{i}_8)$: $P_{g,t+1} = P_{bt}\pi_{g|b}$ and $P_{b,t+1} = 0$; and
9. $i_{t+1} \in [\bar{i}_8, 1]$: $P_{g,t+1} = P_{b,t+1} = 0$.

One needs to spell through these carefully to see that they are correct. We have, and for verification we have also (i) simulated this process for various draws of the idiosyncratic process $\{i_t\}_{t=1}^T$ (where T is large) and, based on the resulting $\{P_t\}_{t=1}^T$ sequence, computed the associated employment outcomes and (ii) made sure that the resulting values are replicated for the same $\{i_t\}_{t=1}^T$ draws by a brute-force averaging across aggregate shock processes. They do.[15]

In the long run, π_t goes to 1/2 in our economy. The wage process in this ergodic state can be simulated easily by drawing $\{i_t\}_{t=1}^T$ sequences; the histogram in the text shows the distribution across values of the employment outcomes after removing aggregate shocks.

[15] Well, now they do.