**Preliminary and Incomplete**

# The Welfare Cost of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk*

Kjetil Storesletten[†] Chris I. Telmer[‡] Amir Yaron[§]

June 2000

**Abstract**

This paper investigates the welfare costs of business cycles in a heterogeneous agent, overlapping generations economy which is distinguished by idiosyncratic labor market risk. Aggregate variation arises both in terms of aggregate productivity shocks and countercyclical variation in the volatility of idiosyncratic shocks. Based on both aggregate data and microeconomic data from the Panel Study on Income Dynamics, we find that the welfare benefits of eliminating aggregate productivity shocks are small, but the benefits of eliminating countercyclical variation in the cross-sectional variance are large. Our results, therefore, are consistent with the increasingly popular notion that distributional effects are an important aspect of understanding the welfare cost of business cycles.

JEL: E21, D31

*Keywords:* business cycles, idiosyncratic risk.

---

*We are grateful to Marianne Baxter, Chris Geczy, Joao Gomes, Per Krusell, José Victor Ríos-Rull, Tony Smith and Stan Zin for helpful comments and discussions. This work has benefited from the support of NSF grant SES-9987602.

[†]IIES, Stockholm University; storeslk@iies.su.se

[‡]GSIA, Carnegie Mellon University; chris.telmer@cmu.edu

[§]Wharton School, University of Pennsylvania and NBER; yarona@savage.wharton.upenn.edu

# 1 Introduction

The notion that an important macroeconomic policy objective should be the mitigation of business cycles is a familiar one, both in academic and non-academic circles. In his influential monograph, Lucas (1987) challenged this orthodoxy by demonstrating, in a representative agent environment, that the welfare gains associated with eliminating variation in aggregate consumption are miniscule. Subsequent work has challenged Lucas' finding on a number of grounds, ranging from statistical assumptions regarding aggregate consumption, to assumptions about preferences, to the basic assumptions underlying the representative agent paradigm. Our paper follows along this latter tack. The main idea is that there are important distributional effects associated with aggregate variation and that these distributional effects lie at the heart of why individuals care about business cycles. The primary effects we focus on are cyclicality in the cross-sectional variance of idiosyncratic risk — the idea that individual-specific shocks become more volatile during aggregate downturns — and how this interacts with life-cycle consumption and savings decisions. What we mean by a 'business cycle,' therefore, goes beyond aggregate productivity shocks to include shocks to the cross-sectional distribution. As a result, an important aspect of the welfare benefits we associate with the removal of business cycles involves how much agents value the removal of countercyclical heteroskedasticity in the idiosyncratic shocks they face.

The specifics of our analysis are as follows. We begin by providing evidence that the focal point of our study — persistent idiosyncratic risk in the labor market, distinguished by countercyclical heteroskedasticity — is a robust feature of U.S. data. Borrowing from previous work, we provide both qualitative evidence as well as quantitative estimates of a specific time series process. Next, we embed these estimates in an overlapping generations (OLG) model in which agents face both aggregate and idiosyncratic sources of uncertainty. The latter arises in the form of persistent shocks to individual labor market productivity. The model's financial market structure, in conjunction with the stochastic process for idiosyncratic shocks and the fact that lives are finite, limits the extent to which agents can collectively pool idiosyncratic risk. As a result, the cyclical component in each individual's earnings process — an increase in volatility during aggregate downturns — is manifested in their consumption process, thereby leading to potential welfare benefits associated with the elimination of business cycles. The finite life, OLG structure of our model — one thing which distinguishes our work from previous studies — plays a crucial role along these lines. Following previous work (Storesletten, Telmer, and Yaron (1999, 2000)), we argue that both finiteness as well as a life cycle pattern in idiosyncratic risk are important elements of the mapping between income and consumption and, consequently, of an assessment of how business cycles interact with heterogeneity.

There are a number of reasons to believe that this framework may lead to sizeable welfare benefits where previous work (discussed below) has found little. The conditional distribution of idiosyncratic earnings shocks and aggregate productivity shocks suggests an amplified cost of aggregate shocks, since aggregate downturns are necessarily coincident with an increase in idiosyncratic risk. Our estimates of the magnitude of this increase are quite large, thus suggesting a potentially large welfare effect. Similarly, the unconditional distribution of the idiosyncratic shocks in our economy is leptokurtic (*i.e.*, the regime shifting nature of the conditional distribution leads to a fat-tailed unconditional distribution). Because agents in our economy dislike kurtosis, and because the elimination of business cycles leads to normally distributed idiosyncratic shocks, this also represents a potential avenue for welfare effects. Finally, the economics of our model are such that, even if the cyclical heteroskedasticity in earnings has a relatively minor impact on its own unconditional distribution, it can have an important impact on the process for consumption and, therefore, for welfare. This is because decision rules for consumption depend not on the unconditional distribution, but on the conditional distribution. The consumption decision rules, therefore, are the main conduit through which variation in the *conditional* distribution of income affects welfare, something which is primarily related to the *unconditional* distribution of consumption. That these decision rules are sensitive to countercyclical heteroskedasticity in idiosyncratic shocks is a fairly robust feature of a growing body of papers, including Alvarez and Jermann (1997), Costain (1999), Constantinides and Duffie (1996) Mankiw (1986), Storesletten, Telmer, and Yaron (1999) and others. What remains to be seen is how important this sensitivity is for welfare.

The methodology with which we measure welfare effects, and the results we find, are as follows. All welfare calculations explicitly incorporate the transition from a world with business cycles to a world without business cycles. As a result, we report both a distribution of welfare gains across ages as well as an average which is constructed with weights corresponding to relative cohort size. Our baseline model is an economy with both aggregate productivity shocks as well cyclical variation in the variance of idiosyncratic shocks. To begin with we eliminate only the aggregate productivity shocks. The welfare effects of doing so are quite small, being equivalent to a 0.08% increase in an agent's consumption. Next, we eliminate the variation in the variance of the idiosyncratic shocks, reformulating the process so that it has the same unconditional variance as in our baseline economy. The welfare costs in this case are quite large, especially in light of previous work. We find that an unborn agent is better off by 1.17% and that, on average, the population is better off by 0.81%. Finally, we ask to what extent "general equilibrium effects" — changes in aggregate savings behavior and market clearing prices — are the driving force. By conducting experiments where aggregate quantities and prices are not allowed to change we conclude that the bulk of our results are not attributable to general equilibrium effects, but instead to the way in which cyclical heteroskedasticity affects consumption decisions.

In terms of related work, our study follows a long line of papers on the welfare costs of business cycles. Most closely related are studies which focus on the interaction between business cycles, heterogeneous agents and uninsurable labor market risk, including Atkeson and Phelan (1994), Beaudry and Pages (1999), İmrohoroğlu (1989), and Krusell and Smith (1999). What distinguishes us is an OLG framework and the way in which we model and calibrate the process which drives heterogeneity. We focus on earnings as opposed to unemployment, and calibrate our model using microeconomic data from the Panel Study on Income Dynamics (PSID). The latter leads to substantially more volatile and persistent idiosyncratic shocks with a greater dependence on aggregate variation. Our approach is also distinct in that the lion's share of the welfare gains we identify are due to the elimination of aggregate fluctuations in the variance of idiosyncratic shocks, not just aggregate productivity shocks. This approach is similar in spirit to İmrohoroğlu (1989), but departs from the approach of Atkeson and Phelan (1994) and Krusell and Smith (1999), who assume that the cyclical distribution of idiosyncratic risk is impervious to changes in the aggregate technological process.

## 2 Evidence

Our question is inherently one about how income and consumption variation across households is related to aggregate, time series variation. In order for this question to have empirical content, we require that the cross-sectional distribution of labor income exhibit the following three properties. First, there must be an interesting cross-sectional distribution to begin with. In our framework this translates to the requirement that individual-specific income must contain an idiosyncratic component which is of sufficient variability. Second, the innovations associated with this process must be quite persistent. Without persistence our theory turns out to closely resemble a theory in which heterogeneity is absent, thereby rendering our question moot. Finally, for the reasons outlined above, the volatility of these innovations must increase during aggregate downturns. In this section we summarize previous work which argues that these are robust features of U.S. data.

Our data correspond to the years 1969-1992 and are obtained from the Panel Study on Income Dynamics (PSID). We examine data at the household level and define annual labor market earnings to equal wages and salaries plus various transfer payments (*e.g.*, unemployment insurance, workers compensation, transfers from non-household family members, and so on). Transfers are included so as to provide measures of idiosyncratic risk *net* of the risk sharing mechanisms which these payments implicitly represent, mechanisms which are absent from our model. We construct a sequence of 22 overlapping panels each with a three year time horizon. For example, our '1975 panel' features three time series observations on each household, covering the years 1975, 1976 and 1977. This feature of our dataset — its departure from the standard longitudinal approach of requiring that all households be represented in all years — is distinct and has a number of important advantages. It provides sufficient time series information and at the same time features a relatively large sample size, a limited degree of survivorship bias and, most importantly, a stable cross-sectional distribution of age across time. The latter is particularly important for our questions, which relate aggregate shocks to idiosyncratic shocks in a life cycle context (*i.e.*, one cannot learn much about labor market shocks from a group of retirees). This comes at a cost, of course, which in our case is primarily related to measurement error and the loss of information associated with restricting ourselves to three time series observations. For more explicit details and more information on the pros and cons of our empirical methodology, we refer the reader to Storesletten, Telmer, and Yaron (1999).

The time series process we are interested in is as follows. Define $u_t^h$ as the idiosyncratic component of the earnings process of an $h$ year old household, at date $t$. We specify the process for $u$ as a mean zero ARMA process — a process with both a transitory and a persistent component — with heteroskedastic innovations in the persistent component:

$$
\begin{aligned}
u_{it}^h &= z_{it}^h + \varepsilon_{it} \\
z_{it}^h &= \rho z_{i,t-1}^{h-1} + \eta_{it} \quad ,
\end{aligned}
\tag{1}
$$

where household age, $h$, is made explicit only when the conditional distribution of a variable depends upon it. We assume that $\varepsilon_{it} \sim \text{Niid}(0, \sigma_\varepsilon^2)$, $\eta_{it} \sim \text{Niid}(0, \sigma_\eta^2(Y_t))$ and,

$$
\begin{aligned}
\sigma_\eta^2(Y_t) &= \sigma_H^2 \text{ if aggregate expansion at date } t \\
&= \sigma_L^2 \text{ if aggregate contraction at date } t \quad ,
\end{aligned}
$$

where $Y_t$ denotes aggregate income. Our methodology will admit any definition of what constitutes an aggregate expansion and contraction. For the results we present, we will use National Income and Product Account (NIPA) data, and define an expansion (contraction) as a year in which growth in real U.S. GNP per capita is above (below) its average over our sample.

Before presenting estimates of the parameters of equation (1) an informal graphical analysis is helpful. Figure 1 reports cross-sectional moments by age and time based on our PSID panel. The top panel — which reports the age profile of the cross-sectional variance — is informative for the magnitude of the parameter $\rho$. The graph demonstrates that over the working years earnings dispersion increases, loosely speaking, linearly. Given the process in equation (1), linearly increasing cross-sectional variance is associated with the value $\rho = 1$. In Storesletten, Telmer, and Yaron (2000) we examine this relationship in more detail including, among other things, 'fixed-effects' in equation (1) and information contained in the autocovariances of $u$. We argue that values for $rho$ close to unity are a very robust implication of our PSID data.

The bottom panel of Figure 1 pools our data across ages instead of time and, at a very informal level, suggests that countercyclical heteroskedasticity is a striking feature of the data. The correlation of the detrended mean and the standard deviation is $-0.74$. The obvious weakness here is inference based on a very limited number of business cycles, with at most five occurring between 1969 and 1992. The econometric methodology to which we turn next attempts to overcome this by exploiting the interaction between age, time and cross-sectional variance, and the fact that we can condition on the macroeconomic history of the U.S. over a much longer time span.

## 2.1 Estimation

Figure 1 suggests that what we are after — evidence that idiosyncratic shocks are both persistent and countercyclically heteroskedastic — may be consistent with data from the PSID. We now turn to a more formal econometric interpretation of the data, which is a summary of results in Storesletten, Telmer, and Yaron (1999).

The essence of our methodology is that we make assumptions about initial conditions and interpret the process in equation (1) as being finite. This allows us to exploit the fact that we can associate an age, $h$, with each observation in our panel and, conditional on macroeconomic data from the National Income and Product Accounts (NIPA), estimate the parameters conditional on an agent-specific 'macroeconomic history.' A simple example helps to illustrate this point.

Ignoring the transitory shocks, $\varepsilon_{it}$, suppose that there were only three generations: young, middle aged and old. Suppose also that the economy is in an expansion at the current time, but was in a recession during the previous two years. Finally, suppose that we only observe data dated at the current time, period $t$. The population cross-sectional variances of the idiosyncratic processes, $u$, for each generation are,

$$\begin{aligned} \text{young}: \; E(u_{it}^1)^2 &= \sigma_H^2 \\ \text{middle aged}: \; E(u_{it}^2)^2 &= \sigma_H^2 + \rho^2 \sigma_L^2 \\ \text{old}: \; E(u_{it}^3)^2 &= \sigma_H^2 + \rho^2 \sigma_L^2 + \rho^4 \sigma_L^2 \end{aligned}$$

Thus, given sufficiently many observations on $u$ for each generation, we can identify all three parameters *without any time series observations on individual agents.* The key piece of information we are exploiting is how the cross-sectional variance at date $t$ varies across age cohorts and how this interacts with what is essentially a cohort-specific macroeconomic history which is known at date $t$. Note also that our methodology is perfectly well defined for any values of $\rho$, including those greater than unity. That is, we do not suffer from the litany of issues which arise in the context of inference with potentially non-stationary time series processes. This comes at a cost, of course, in the form of assumptions regarding initial conditions. In Storesletten, Telmer, and Yaron (1999) and Storesletten, Telmer, and Yaron (2000) we argue that our results are robust to a number of variations along these lines, including first differencing, quasi differencing and the incorporation of 'fixed effects.'

In Exhibit 1 we reproduce GMM-based estimates from Storesletten, Telmer, and Yaron (1999) which are based on a generalization of this methodology, applied to our 1969-1992 panel. The analog of the above example — which exploits 3 years of macroeconomic information using only 1 year of microeconomic information — is that we incorporate data going back to 1910 in spite of having PSID data which only begins in 1969.

**Exhibit 1: Idiosyncratic Risk: Parameter Estimates**

| | $\rho$ | $\sigma_H^2$ | $\sigma_L^2$ | $\sigma_\varepsilon^2$ |
|---|---|---|---|---|
| Estimate | 0.916 | 0.037 | 0.181 | 0.025 |
| Standard Error | 0.009 | 0.007 | 0.033 | 0.007 |

Our estimates confirm what is suggested by Figure 1. Idiosyncratic shocks appear to be quite persistent, with an estimated autocorrelation coefficient of 0.92. The evidence in favor of countercyclical heteroskedasticity is also striking. Our estimates imply that the conditional standard deviation of the persistent shocks increases by 126% from expansion to contraction.

There are a number of reference points which are suggestive of the robustness of these findings. In relation to persistence and the relative magnitude of the persistent and transitory shocks, papers by Abowd and Card (1989), Hubbard, Skinner, and Zeldes (1994) and MaCurdy (1982) find results which are quite similar to ours. In addition, in Storesletten, Telmer, and Yaron (2000) we estimate $\rho$ using a variety of alternative methods, including an examination of the implications for the cross-sectional distribution of consumption, and find convincing evidence in support of values of $\rho$ between 0.95 and unity. In relation to the countercyclical heteroskedasticity, comparable studies are more difficult to find. Heaton and Lucas (1996) find evidence of much smaller effects than we do but, as we argue in Storesletten, Telmer, and Yaron (1999), this is to some extent predictable given differences in statistical methodology and data. One corroborating piece of evidence which we find convincing involves how changes in the *conditional* variance manifest themselves in changes in the *unconditional* cross-sectional variance. In Storesletten, Telmer, and Yaron (1999) we show that the implications of Exhibit 1 — relatively large cyclical changes in the conditional variance — are not necessarily inconsistent with the relatively modest changes in the cross-sectional variance documented in Figure 1. That is, given the aggregate dynamics and the demographic structure of our economy, the variance in the variance which we observe in Figure 1 *implies*, to some extent, the striking results presented in Exhibit 1.

To make this last point more concrete, it is instructive to consider the above three-generation example. Given generations of equal size, and given the process for aggregate dynamics which we employ below, the estimates in Exhibit 1 imply a (time series) standard deviation in the percentage changes in the *conditional* standard deviation of the cross-sectional distribution of 52 percent. The volatility in the percentage changes in the associated *unconditional* cross-sectional standard deviation is only 27 percent. Moreover, this effect will be strengthened by additional generations because the older a cohort is, the less the cross-sectional variance associated with that cohort will vary over time (*i.e.*, with age, there is more 'averaging' within the distribution). The three-generation example, therefore, understates what is relevant for our data and our model, where there are 43 working-age generations.

## 3 Theory

We now embed the evidence of the previous section into a stationary overlapping generations (OLG) model. The motivation for a life cycle framework arises from previous work where we've argued that finite lives and a life-cycle pattern in idiosyncratic risk are important for understanding the mapping between income and consumption. The main idea is that the infinite horizon abstraction affords theoretical agents far greater ability to use financial markets to self-insure against idiosyncratic risk than actual agents seem to enjoy. Given this, a life cycle framework seems important for our current question, which presumes that idiosyncratic risk is an important aspect of consumption outcomes.

Our model is essentially a one-asset version of the framework developed in Storesletten, Telmer, and Yaron (1999), which is itself a variant of Ríos-Rull's (1994) model, designed to incorporate idiosyncratic risk. Agents are indexed by their age, $h$, where $h \in \mathcal{H} = \{1, 2, \ldots, H\}$. Each of the $H$ age cohorts consists of a large number of atomistic agents who face uncertain lifetimes with maximum length of $H$ years. Each year a new cohort of agents are born and some positive fraction of each existing cohort dies. We use $\phi_h$ to denote the unconditional probability of surviving up to age $h$, with $\phi_1 = 1$, and use $\xi_h = \phi_h/\phi_{h-1}$, $h = 2, 3, \ldots, H$, to denote the probability of surviving up to age $h$, conditional on being alive at age $h-1$. The fraction of the total population attributable to each age cohort is fixed over time at $\varphi_h$ and the population grows at rate $\vartheta$.

Each individual agent is characterized by a preference ordering over consumption distributions, an endowment process, and an asset market position. Preferences for an unborn agent are represented by,

$$E \sum_{h=1}^{H} \beta^h \phi_h u(c_h) \quad , \tag{2}$$

where $\beta$ denotes the utility discount factor, $c_h$ denotes the consumption of an $h$ year old agent, $u$ is the standard twice differentiable, strictly concave utility function and the expectation is assumed to be conditional on the state of the economy prior to birth.

Agents begin working at age 22 (or $h = 1$) and, conditional on surviving, retire at age 65 (or $h = 43$). After retirement they must finance consumption entirely from an existing stock of assets. Prior to retirement an agent of age $h$ receives an annual endowment, $n_h$, of an age-specific amount of labor hours (or, equivalently, productive efficiency units) which they supply inelastically to an aggregate production technology. Individual labor income is then determined as the product of hours worked and the market clearing wage rate.

We adopt the following process for the logarithm of hours worked,

$$\log n_h = \kappa_h + z_h \quad , \tag{3}$$

where

$$z_h = \rho z_{h-1} + \eta_h \quad , \quad \eta_h \sim \mathrm{N}(\sigma_\eta^2(Z)/2, \sigma_\eta^2(Z)) \quad ,$$

$Z$ is an aggregate productivity shock and $\kappa_h$ is a parameter used to match the cross sectional distribution of mean age-earnings profile. Idiosyncratic shocks, therefore, are comprised of both a transitory and a persistent component, following the time series model in section 2. Countercyclical heteroskedasticity arises in the form of time variation in the variance of the innovations to the persistent process. More specifically,

$$\begin{aligned} \sigma_\eta^2(Z) &= \sigma_H^2 \text{ if } Z \geq E(Z) \\ \sigma_\eta^2(Z) &= \sigma_L^2 \text{ if } Z < E(Z) \ . \end{aligned}$$

The notion of countercyclical heteroskedasticity is simply the condition that $\sigma_H < \sigma_L$.

Output in this world is produced by an aggregate technology to which individuals rent their labor services and capital. The production function takes the form,

$$Y = Zf(K, N) \quad , \tag{4}$$

where $K$ and $N$ represent per capita capital and labor, respectively, $Y$ represents per capita output, and $Z$ is a technology shock restricted to lie in a finite set $\mathcal{Z}$. Given aggregate consumption, $C$, and the rate of depreciation on aggregate capital, $\delta$, the law of motion for aggregate capital can be written,

$$K' = Y - C + (1 - \delta)K \quad . \tag{5}$$

The financial market structure allows agents to trade in only one asset, shares of ownership in the risky aggregate technology. Each agent's choice problem is one-dimensional: given knowledge of their idiosyncratic status, they simply choose an amount of assets to accumulate, which we label $a_h$. Asset holdings are restricted to lie in a set $\mathcal{A}$.

The state of the economy can now be represented as a pair $(\mu, Z)$, where $\mu$ is a measure defined over an appropriate family of subsets of $S = (\mathcal{H} \times \tilde{\mathcal{Z}} \times \mathcal{A})$, where $\tilde{\mathcal{Z}}$ is the product space containing all possible idiosyncratic shocks. In words, $\mu$ is simply a distribution of agents across ages, idiosyncratic shocks, and capital holdings (wealth). The aspect of $\mu$ which is somewhat non-standard is that, because of the aggregate uncertainty in our economy, it must evolve stochastically over time (*i.e.*, $\mu$ belongs to some family of distributions over which there is defined yet another probability measure). We therefore use $G$ to denote the law of motion of $\mu$, the cross sectional distribution of the economy,

$$\mu' = G(\mu, Z). \tag{6}$$

This characterization of the state of the economy implies that prices for the risky rate of return on capital and wages can be expressed as $R(\mu, Z)$ and $W(\mu, Z)$ respectively. The timing convention we use is that consumption-saving decision decision are made at the end of the period and market returns are paid the following period at the realized capital rental rate. Thus, the decisions of an agent of age $h$ are constrained by

$$\begin{aligned} c_h + a'_{h+1}\xi_{h+1} &\leq a_h R(\mu, Z) + n_h W(\mu, Z) \\ a'_{h+1} &\geq \underline{a} \quad \text{and} \quad a_{H+1} \geq 0 \ , \end{aligned} \tag{7}$$

where $a_h$ denotes beginning of period asset (or capital) holdings and $a'_{h+1}$ denotes end of period asset holdings. The term $\xi_{h+1}$ — the conditional probability of surviving to age $h+1$ given that one survives to age $h$ — is a convenient way to represent perfect annuity markets, something we incorporate so as to focus on labor market risk only. The idea is simply that, because one may not survive to capture the benefit of saving, the sacrifice in terms of current consumption is reduced in a manner which is actuarially fair.

Denoting the value function of an agent of age $h$ as $V_h$, an agent's choice problem can be represented as,

$$V_h(z_h, a_h, Z, \mu) = \max_{a'_{h+1}} \left\{ u(c_h) + \beta \frac{\phi_{h+1}}{\phi_h} E\left[ V'_{h+1}(z'_{h+1}, a'_{h+1}, Z', G(\mu, Z)) \mid z_h, Z \right] \right\} \tag{8}$$

subject to equations (7).

An equilibrium can now be represented as a collection of stationary price functions, $R(\mu, Z)$ and $W(\mu, Z)$, a set of cohort specific value functions and decision rules, $\{V_h, a'_{h+1}\}_{h=1}^{H}$, and a law of motion for $\mu$, $\mu' = G(\mu, Z)$, such that the firm's profit maximization problem is satisfied,

$$\begin{aligned} R(\mu, Z) &= Zf_1(K, N) - \delta + 1 \\ W(\mu, Z) &= Zf_2(K, N) \ , \end{aligned}$$

aggregate quantities result from individual decisions,

$$\begin{aligned} K &= \int_S a \ d\mu \\ N &= \int_S n \ d\mu \ , \end{aligned}$$

agents' optimization problems are satisfied given the law of motion for $(Z, \mu)$ (so that $\{V_h, a'_{h+1}\}_{h=1}^{H}$ satisfy problem (8)), and the law of motion, $G$, is consistent with individual behavior.

# 4 Calibration

We interpret one period in our model as corresponding to one year of calendar time. The aggregate production technology is Cobb-Douglas:

$$Y = ZK^{\theta}N^{1-\theta} \ ,$$

We set $\theta$ equal to 0.4 (which corresponds to capital's share of national income being 40%, *cf.* Cooley and Prescott (1995)) and allow for a 7.8% annual depreciation rate on the aggregate capital stock. The technology shocks, $Z$, follow a first-order Markov chain with parameter values chosen so that theoretical aggregate *output* matches several important features of observed, aggregate U.S. output. We focus on output — as opposed to consumption — as the anchor for calibrating our aggregate economy because it provides for a conservative measure of business cycle fluctuations in consumption. By following the business cycle tradition of matching the variability in Solow residuals, our theoretical aggregate consumption process is too smooth by a factor of two, excessively smooth consumption being a common feature of business cycle models.

The average growth rate in GNP per capita is chosen to be 1.5% per year. The transition probabilities for $Z$ are chosen so that the expected duration of a 'business cycle' is 6 years, whereas the possible realizations of $Z$ are chosen to match the variability in theoretical and observed aggregate output. The end result is a two state Markov chain for the aggregate shock with $Z \in \{-0.057, 0.057\}$ and probability of remaining in the current state of 2/3.

Turning to the characteristics of individual agents, preferences are identical (up to age-dependent mortality risk) and are described by equation (2). We parameterize the period utility function with the standard isoelastic specification,

$$u(c) = \frac{c^{1-\alpha} - 1}{1 - \alpha} \ .$$

For the time being, we set the risk aversion parameter, $\alpha$, to 2 and the utility discount factor, $\beta$, to 0.95. The latter is chosen in order to generate an aggregate capital to output ratio of 2.9, which is within the range of standard estimates based on U.S. data (*cf.* Cooley and Prescott (1995)). The choice of $\alpha$ is arbitrary, but seems a sensible starting place in addition to representing a rough consensus of what is typical in the literature.

The demographic structure of our economy is calibrated to correspond to several simple properties of the U.S. work force. Agents are 'born' at age 22, retire at age 65 and are dead by age 80. 'Retirement' is defined as having one's labor income drop to zero and having to finance consumption from an existing stock of assets. Mortality rates are chosen to match those of U.S. females in 1991 and population growth is set to 1.0%.

The process for idiosyncratic labor supply, equation (3), is implemented as a discrete approximation to the autoregressive time series model and is parameterized using our point estimates from Exhibit 1. The age dependent intercept terms, $\kappa_h$, are chosen so that, on average, our theoretical age-earnings profile matches that of the P-SID. The persistent process, a regime-shifting autoregression for $z_h$ with $\sigma_H = 0.425$, $\sigma_L = 0.192$ and $\rho = 0.916$, is approximated with a 19-state Markov chain where the conditional heteroskedasticity is generated by variation in the transition probabilities, as opposed to the potential realizations of $z_h$. Further details are provided in Storesletten, Telmer, and Yaron (1999). Finally, in our baseline economy we disallow short-selling of the risky technology (i.e. there is no borrowing) — thereby setting $\underline{a} = 0$.

In terms of computing equilibrium we use methods developed in Storesletten, Telmer, and Yaron (1999), which are finite-lived extensions of work by Castañeda, Díaz-Giménez, and Ríos-Rull (1994), den Haan (1994) and, in particular, Krusell and Smith (1998). Specifically, we approximate $\mu$, the distribution of capital across agents and age, with a finite number of moments and characterize the transition function $G(\mu, Z)$ corresponding to these moments. Given such a parameterization, agents decisions are derived from a finite dynamic programming problem. Using the decision rules we simulate the economy and derive an empirically determined distribution of agent specific capital which can in turn be compared to the conjectured functional form for that distribution. An equilibrium is reached when the empirical process coincides with the parameterized process. We verify that in such an equilibrium agent's perceptions about $\mu$ and $G(\mu, Z)$ are very precise in the sense that their Euler equation errors are very small.

We derive the stochastic stationary equilibrium for our benchmark economy as well as for each of the alternative economies we study. As noted earlier, in computing the welfare gains of moving from our benchmark economy to one of the alternative economies, we solve for an equilibrium transition path. Having characterized the stochastic stationary equilibrium in each economy, we conjecture the matrices characterizing the cross-sectional distribution of capital and age $G(\mu, Z)$ during the transition and a time period by which the economy settles into the alternative economy's stochastic stationary equilibrium. These parameters are iterated upon until the equilibrium transition path is attained. Further details on the computations are available in Storesletten, Telmer, and Yaron (1999).

## 5 Quantitative Results

We start by describing the basic features of our baseline economy. We then describe how we compute welfare gains. That is, how much will individuals be willing to pay in order to move to an alternative economy? We then proceed to describe our welfare results, as well as various experiments designed to isolate some of the economics underlying our results.

Our baseline economy is broadly consistent with several key features of the aggregate U.S. economy. Consumption is about 70% of the magnitude of output (which equals 50% plus about 20% for government which is not explicitly modeled). Moreover, in Table 1 we demonstrate that our parameterization for technological shocks generates realistic stochastic behavior for various (endogenous) aggregate quantities in our model. Doing so is important as a check on our quantitative analysis. Table 1 shows that the variability in theoretical aggregate consumption is about 1/2 of the data, but that aggregate output has the same volatility as the data. The same basically holds for investment. The contemporaneous correlation of aggregate output and investment with consumption match those of the data, but their autocorrelations are somewhat smaller. Relatively low autocorrelation, however, is a necessary implication of our two-state aggregate shock process averaging one 'cycle' every 6 years. In other words, given our simple process for aggregate variation, we can match either autocorrelation or some notion of 'cyclicality,' but not both. We choose the latter, largely because of the methodology used in the last section to measure countercyclical, cross sectional variation.

Our baseline economy is also consistent in many respects with key cross-sectional moments of the consumption distribution. The annual cross-sectional variance of consumption growth in our model is 0.029, which is close to what is observed in the CEX (for example Souleles (1999) finds the cross-sectional variance of consumption growth to have a mean of .023). Moreover *changes* in this cross-sectional consumption growth measure do not fluctuate very much. The mean and variance are 0.0001 and 0.0007 respectively, both on the order of what Brav, Constantinides, and Geczy (2000) find in calibrating their economy. Moreover, a key channel by which these authors calibrate their economy is the time series regression of returns onto the variance of idiosyncratic risk in consumption growth. The corresponding regression result in the model is -0.157, which is consistent with the parameter range Brav, Constantinides, and Geczy (2000) find in their analysis. Finally, in Storesletten, Telmer, and Yaron (2000) we remove business cycles from our analysis by explicitly controlling for cohort effects and show that a model without business cycles but otherwise with essentially the same features as the one here accounts for the empirical rise in consumption inequality across age. Hence, in spite of endowing our economy with substantial income shocks, the resulting distribution of consumption is not entirely inconsistent with the limited amount of evidence which exists.

## 5.1 Welfare Comparisons Across Alternative Economies

To derive our welfare measure we ask what are the welfare gains from reducing various shocks present in our baseline economy. Our approach for obtaining welfare comparisons is standard. We denote the value function for an $h$-year-old agent living in economy $A$ as $V_h^A(\cdot)$. Similarly, $V_h^B(\cdot)$ denotes this agent's value function, should they live in an economy with an alternative stochastic process for income shocks, say economy $B$. The welfare gain associated with moving from economy $A$ to economy $B$ is measured as the proportional compensation required to increase consumption permanently in economy $A$, in a way that will make the agent indifferent between the two economies.

With a slight abuse of notation, a few auxiliary definitions will help describe our procedure for calculating welfare gains. Based on $\mu$, define $\Phi(\tilde{z}, a|\mu, Z)$ as the conditional density of idiosyncratic shocks (in economy A) and asset holdings *given* aggregate shock $Z$ and the cross-sectional measure $\mu$. Also, let $P(\mu|Z)$ be the probability density of $\mu$ given $Z$ and let $F(\cdot)$ be the density function for aggregate shocks $Z$, again for economy A. The welfare gain for age-cohort $h$ can now be expressed as the number $\psi_h$ which results in average utility being equated:

$$\int\int[\int_{\tilde{Z}\times\mathcal{A}}\left\{V_h^A(\cdot, Z, \mu; \psi_h) - V_h^B(\cdot, Z, \mu)\right\}\Phi(\tilde{z}, a|\mu, Z)d\tilde{z}da]P(\mu|Z)d\mu F(Z)dZ \quad = \quad 0 \tag{9}$$

where $V_h^A(\cdot, Z, \mu; \psi_h)$ solves

$$V_h^A(\cdot, Z, \mu; \psi_h) = \max_{a'_{h+1}} \left\{ u(c[1+\psi_h]) + \beta \frac{\phi_{h+1}}{\phi_h} EV_{h+1}^A(\cdot, Z, \mu; \psi_h) \right\} \quad (10)$$

subject to the budget constraint (7), and the proportional change $\psi_h$ in consumption is held equal across all agents of age $h$. Finally our overall measure for welfare, $\psi$ is just the weighted sum of $\psi_h$ by $\varphi_h$ the cohorts size of those who are alive.

Before continuing to discuss our results a few additional words about our computations are in order. Our benchmark economy varies over time so any comparison to an alternative economy involves integrating over the welfare gains starting from alternative states in the benchmark economy. This is reflected in the integration over $\mu$ and $Z$ in our definition of welfare gains in equation (9). Another aspect of our welfare computation is the fact that our analysis takes explicit account of the equilibrium transition path to any of the alternative economies under considerations. Therefore, when we report *average* welfare gains they actually pertain to a weighted average of the welfare gains during the transitions for the different cohorts. The welfare gains reported for a *new-born* reflect the gain that would be associated without considering the transition.

## 5.2 The Cost Of Business Cycles

Overall, we find large welfare gains from removing aggregate productivity shocks and business cycle variation in idiosyncratic shocks. The welfare gains are 0.82% for the 'average' agent and 1.32% for new-borns.

Understanding the economics behind our results is somewhat involved. In particular there are a number of conceivable forces at work. The primary candidates we wish to analyze (i) elimination of aggregate productivity shocks, (ii) elimination of cyclical variation in the variance of idiosyncratic shocks (iii) 'general equilibrium effects:' the impact of reduction in aggregate savings on market clearing prices, and (iv) the role of finite lives. In order to decompose our overall results into components attributable to each of these effects, we conduct a number of additional experiments, each involving an economy in which progressively fewer of the effects are at work. In simple terms, we first eliminate aggregate productivity shocks $Z$, then business cycle variation in the variance of idiosyncratic shocks, and then general equilibrium effects. The differences we find each step of the way constitute the contribution associated with the factor most recently removed. We conduct these successive experiments for our OLG economy as well as for an otherwise identical infinitely lived agent economy. Differences across the infinitely lived agent economy and the OLG economy shed light on the role finite lives play in our analysis.

Exhibit 2 summarizes our experiments, the successive effects discussed above, and the corresponding welfare gains associated with each step:

**Exhibit 2: Summary of Experiments Underlying Welfare Decompositions**

| *Notation for Welfare Gain* | *Gain (%) in OLG Economy* | | *Gain (%) in ∞ Lived Economy* | *Eliminated* Effects | *Comments* |
|---|---|---|---|---|---|
| | *Avg.* | *New Born* | | | |
| $\psi_{a,b}$ | 0.08% | 0.15% | 0.03% | $Z_t = E(Z)\ \forall t$ | Comparison includes transition. |
| $\psi_{b,c}$ | 0.82% | 1.32% | 1.05% | $\sigma_H = \sigma_L = \bar{\sigma}$ | Comparison with prices of economy B |
| $\psi_{b,d} - \psi_{b,c}$ | −0.01% | −0.15% | −0.01% | $\sigma_H = \sigma_L = \bar{\sigma}$ (of transition) | General Equilibrium effects |
| $\psi$ | 0.89% | 1.32% | 1.07% | - | Total Welfare Gain $\psi = \psi_{a,b} + \psi_{b,d}$ |

Consider the welfare gain of moving from our benchmark economy (denoted economy $A$) to an economy in which there are no aggregate productivity shocks nor are there any cyclical variation in idiosyncratic shocks (denoted economy $D$). First, $\psi$ represents this overall welfare gain we are attempting to decompose. The idea behind this decomposition is quite simple. Each of the economic factors we have highlighted play some role in generating the overall welfare gain $\psi$. The first row in Exhibit 2, $\psi_{a,b}$, represents the welfare gain of moving from the benchmark economy $A$ to an economy in which aggregate productivity shocks, $Z$'s, have been removed (denoted economy $B$). In this last economy the aggregate productivity shock is set to its unconditional mean. Nonetheless, in the background we maintain the dependency of idiosyncratic shocks on business cycle fluctuations. Hence, $\psi_{a,b}$ provides information on the *direct* effect of eliminating aggregate productivity shocks without the compounding effect of altering the business cycle structure of idiosyncratic shocks. In a similar fashion $\psi_{b,c}$ denotes the welfare gain of moving from economy $B$ to an alternative economy $C$ – one without business cycle variation in the variance of idiosyncratic shocks. As described earlier, in this economy we set the variance of idiosyncratic shocks to its unconditional level. In this alternative economy we also maintain the prices agents faced in economy $B$. The idea is straightforward. Eliminating business cycle variations in idiosyncratic shocks will likely cause a reduction in precautionary savings which in turn can effect market clearing prices. $\psi_{b,c}$ provides a measure without any such compounding effects. Finally, $\psi_{b,d} - \psi_{b,c}$ is essentially a residual effect, describing the general equilibrium effects of savings on prices.

The results in Exhibit 2 indicate that removing only aggregate productivity shocks lead to small welfare gains ($\psi_{a,b}$). Although these welfare gains are slightly larger than what Lucas (1987) and İmrohoroğlu (1989) found, they are still quite small and reflect the fact that aggregate fluctuations in output do not present large risks as per-capita consumption does not vary significantly. This is consistent with much of the previous literature. On the other hand, the welfare gains of reducing business-cycle variation in the variance of idiosyncratic shocks, $\psi_{b,c}$, are large and in our OLG economy amount to 0.82% on average. The corresponding welfare gain for the newborns, who face this idiosyncratic risk throughout their life, is substantially larger at 1.32%. Finally, the general equilibrium effects, $\psi_{c,d}$ turn out to to be negative and quite small on average. Overall, our results suggest that eliminating business cycle shocks can be quantitatively important provided that it is feasible to eliminate business cycle variation in the variance of idiosyncratic risk.

#### 5.2.1 Welfare Gains over the Life Cycle

Figure 2 decomposes the welfare gains we present in Exhibit 2 by age. The dash line decomposes the welfare gains of removing aggregate productivity shocks (corresponding to $\psi_{a,b}$). The solid line provides the age profile of welfare gains of removing also the cyclical variation in idiosyncratic risk, but without accounting for the general equilibrium effects – prices are those of the economy *with* cyclical variation in idiosyncratic risk but *without* aggregate productivity shocks. Thus, this line corresponds to $\psi_{b,c}$. Finally, the solid-starred line in Figure 2 presents $\psi_{b,d}$ – the age profile of the welfare gains stemming from removing business-cycle variation in idiosyncratic risk from an economy without fluctuations in aggregate productivity, but where the general equilibrium effects of changing prices during the transition are accounted for.

It is clear that for each experiment the welfare gains are substantially larger for those who have many years left to live. The graph also shows that the small welfare gains attributable to eliminating aggregate productivity shocks are uniformly small for all age groups. However, the welfare gains of eliminating business cycle variation in idiosyncratic risk are very large for the young and approach zero by retirement. The intuition for this is quite simple. The young hold most of their wealth in terms of human capital and are susceptible to large shocks relative to their wealth. The old on the other hand have no more human capital risk and have accumulated financial wealth and can effectively buffer shocks. The fluctuation in their wealth is solely due to fluctuation in the risky return on capital which is relatively small. Hence moving to the alternative economy is not a big attraction for the old.

The difference between the solid line and the solid-starred line in Figure 2 represents the age profile of the general equilibrium effects ($\psi_{b,d} - \psi_{b,c}$), and reveals that the general equilibrium effects of eliminating business cycle variation in idiosyncratic risk are quite different for different age groups. In particular, the general equilibrium effects are negative for the *young* and are positive for the *old*. The explanation for these age specific general equilibrium effects lies, again, in the respective composition of agents wealth and the direction in which average prices have changed. In the e-conomy with homoskedastic conditional variance agents save less and the resulting equilibrium is one with a higher, on average, return on capital and lower wages. These changes in prices effect the old and the young differently. The old who hold large financial wealth gain from the rise in the return on capital, while the young lose, on average, due to the decline in wages.

#### 5.2.2 The Role of Finite Lives

In order to assess the importance of finite lives for the welfare calculations we report, we also analyze an economy which has infinitely lived agents but otherwise is similar to the benchmark OLG economy. The difference between the welfare gains from the infinitely lived agent economy and the corresponding gains from the finitely lived agent economy can be attributed to the role finite lives have in our economy. The *average* welfare gain under the infinitely lived economy is 0.67% *larger* than the corresponding number in the OLG economy. At first blush this result may appear to contradict our intuition that the gains from eliminating business cycles ought to be larger for finitely lived agents. However, because in the OLG economy retirees do not work and, therefore, are not subject to any idiosyncratic variation risk (except the small variation in returns) and because they constitute a sizable cohort the welfare gains are downward biased. If one appropriately compares the welfare gains in the infinitely lived agent economy to the average welfare gains for those who work in the OLG economy the expected welfare rankings align up closely.

#### 5.2.3 Consumption and Income Variability

The welfare gains we find above are large relative to the corresponding figures other authors have found. We find it instructive to provide an additional context for the welfare gains we report. Specifically, we want to know how much of the welfare gains we have found are attributable to the income process we 'endow' into our model. In Appendix A we provide an analytical counterpart to Lucas (1987) calculation using our income process as consumption. We first show that the welfare gains of removing counter-cyclical variation in consumption depend on the skewness and kurtosis of the consumption processes in the *heteroskedastic* and *homoskedastic* economy. We then provide our analytical calculation of removing counter-cyclical variation, assuming log-normality for all variables in an *i.i.d* economy.

We find that if consumption were indeed equal to income and we would eliminate the business cycle variation in the variance of idiosyncratic shocks, the welfare gain will be very small for the case of $\gamma = 2$ but can be quite large for risk aversions close to 3 and can be negative for risk aversion close to 1. The reason is simple. The skewness and kurtosis of the *level* of consumption is larger in the *heteroskedastic* economy than it is in the *homoskedastic* economy. Agents dislike large kurtosis but prefer the larger positive skewness. For low risk aversion the skewness factor dominates and negative welfare gains arise. For slightly larger risk aversions ($\gamma > 2$) the kurtosis channel dominates which result in large positive welfare gains. Of course the small welfare gains we find for risk aversion of 2 do not contradict our earlier findings as consumption in our model is neither $i.i.d$ nor exogenous and various agents are borrowing constrained. Overall, our analytical analysis points to the fact that risk aversion is quantitatively important in our analysis and that the endogeneity of consumption can yield substantially larger welfare gains than the fictious experiments based on exogenous consumption.

# 6 Conclusions

Our main finding is that agents value the elimination of business cycles by a significant amount. The primary force driving these welfare gains is the elimination of business cycle fluctuations in the distribution of idiosyncratic risk. This channel for producing welfare gains from business cycles is different than alternative stories being put forth for delivering large welfare gains from business cycles in that it still relies on the notion of consumption smoothing. In essence, agents prefer the *homoskedastic* economy since there is a smaller chance for extreme bad events that occur in downturns.

A drawback of our analysis lies in the exogenous way we introduce idiosyncratic risk over the business cycle. The issue of whether business cycle variation in the distribution of idiosyncratic risk will be eliminated once aggregate shocks are removed can not be answered structurally within our context. In our analysis we went ahead and measured what the gains would be had such variations be eliminated. At a minimum the large welfare gains we find suggest that elaborating on this issue is an important future research agenda. As mentioned earlier, Beaudry and Pages (1999) and Gomes, Greenwood, and Rebelo (1999) analyze the unemployment process in a detailed way. Indeed Beaudry and Pages (1999) find significant welfare gains although smaller than what we have found. Reconciling our earnings process within such a framework and asking policy questions seems a natural and direct avenue for progress on this important question.

## References


Abowd, J. and D. Card, (1989), On the covariance structure of earnings and hours changes, *Econometrica* 57, 411–455.

Alvarez, F. and U. Jermann, (1997), Asset pricing when risk sharing is limited by default, Unpublished working paper, University of Chicago.

Atkeson, A. and C. Phelan, (1994), Reconsidering the costs of business cycles with incomplete markets, in Stanley Fischer and Julio Rotemberg eds., *NBER Macroeconomics Annual 1994*.

Beaudry, P. and C. Pages, (1999), The cost of business cycles and the stabilization value of unemployment insurance, manuscript, University of British Columbia.

Brav, A., G. M. Constantinides, and C. Geczy, (2000), Asset pricing with heterogeneous consumers and limited participation: Empirical evidence, Working paper, The Wharton School, University of Pennsylvania.

Castañeda, J., J. Díaz-Giménez, and J. V. Ríos-Rull, (1994), Unemployment dynamics, aggregate fluctuations, and income distribution, Working paper, University of Pennsylvania.

Constantinides, G. M. and D. Duffie, (1996), Asset pricing with heterogeneous consumers, *Journal of Political Economy* 104, 219–240.

Cooley, T. F. and E. C. Prescott, (1995), *Economic Growth and Business Cycles*, chapter 1, 1–38, Princeton University Press, New Jersey, In *Frontiers of Business Cycle Research*, T. Cooley, editor.

Costain, J., (1999), Multiple ....., Unpublished manuscript, Universitat Pompeu Fabra.

Deaton, A. and C. Paxson, (1994), Intertepmoral choice and inequality, *Journal of Political Economy* 102, 437–467.

den Haan, W., (1994), Heterogeneity, aggregate uncertainty and the short term interest rate: a case study of two solution techniques, Working paper, University of California at San Diego.

Gomes, J., J. Greenwood, and S. Rebelo, (1999), Equilibrium unemployment, Manuscript, University of Rochester.

Heaton, J. and D. J. Lucas, (1996), Evaluating the effects of incomplete markets on risk sharing and asset pricing, *Journal of Political Economy* 104, 443–487.

Hubbard, R. G., J. Skinner, and S. P. Zeldes, (1994), The importance of precautionary motives in explaining individual and aggregate savings, *Carnegie Rochester Conferance Series on Public Policy* 40, 59–126.

İmrohoroğlu, A., (1989), Cost of business cycles with indivisibilities, *Journal of Political Economy* 97, 1364–1383.

Krusell, P. and A. A. Smith, (1998), Income and wealth heterogeneity in the macroeconomy, forthcoming, *Journal of Political Economy.*

Krusell, P. and A. A. Smith, (1999), The welfare effects of eliminating business cycles, *Review of Economic Dynamics* 2, 245–272.

Lucas, R. E., (1987), *Models of Business Cycles*, Basil Blackwell.

MaCurdy, T. E., (1982), The use of time series processes to model the error structure of earnings in a longitudinal data analysis, *Journal of Econometrics* 18, 83–114.

Mankiw, N. G., (1986), The equity premium and the concentration of aggregate shocks, *Journal of Financial Economics* 17, 211–219.

Ríos-Rull, J. V., (1994), On the quantitative importance of market completeness, *Journal of Monetary Economics* 34, 463–496.

Souleles, N., (1999), Household securities purchases, transactions costs, and hedging motives, Working paper, The Wharton School, University of Pennsylvania.

Storesletten, K., C. I. Telmer, and A. Yaron, (1999), Asset pricing with idiosyncratic risk and overlapping generations, GSIA working paper no. 1998-E226, Carnegie Mellon University.

Storesletten, K., C. I. Telmer, and A. Yaron, (2000), Consumption and risk sharing over the life cycle, Unpublished working paper, Carnegie Mellon University.

# Appendix A
## Welfare Gains in Autarkic Economies

In this appendix we provide intuition for some of the sources underlying the welfare gains for removing counter cyclical variation (CCV). We also analytically redo Lucas (1987) calculation under the assumption that there is no risk sharing in our economy and consumption is equal to income. We show that the welfare gains are closely related to probabilities of facing extreme events and depend on the skewness and kurtosis of the consumption distribution. It should be noted, however, that in spite of the fact that our analysis below is calibrated to income process, the welfare gains we compute analytically do not provide a bound on the welfare gains that can arise from consumption. This is due to the fact that the welfare gains are a relative statement about expected utilities when consumption is endogenous, and also in the model there are people who are borrowing constrained something we abstract from in this appendix.

### A.1 Higher Moments and Welfare Gains

In this section we want to provide some intuition for the role different moments in the distribution of consumption have on welfare gains. Let $f(c)$ and $g(c)$ be the probability density function for consumption under the *heteroskedastic* and *homoskedastic* economies respectively. For our analysis to be meaningful we impose the following conditions across the two economies,

$$\begin{aligned} \bar{c} &\equiv E_f(c) = E_g(c) \\ Var_f(c) &= Var_g(c) \end{aligned} \tag{11}$$

where the subscript in the mean and variance denotes the distribution (p.d.f) with which expections are taken. These two restrictions imply that the unconditional *mean* and *variance* of the *level* of consumption are equal across the two economies.

The welfare gains to removing CCV are based on comparing the expected utilities across the two economies. The expected utility under the two economies can be approximated by a Taylor expansion around $\bar{c}$ as follows,

$$\begin{aligned} E_i(u(c)) \approx\ & u(\bar{c}) + u'(\bar{c})E(c-\bar{c}) + \frac{1}{2}u''(\bar{c})E_i(c-\bar{c})^2 + \\ & \frac{1}{6}u'''(\bar{c})E_i(c-\bar{c})^3 + \frac{1}{24}u''''(\bar{c})E_i(c-\bar{c})^4 \qquad i = f, g \end{aligned}$$

where $E_i$ denotes again expectations with respect to the p.d.f's $f$ and $g$. Using our assumptions above and noting that the terms multiplying $u'''$ and $u''''$ are simply the population skewness and kurtosis respectively, it follows that the difference in welfare across the two economies depends on

$$E_g(c) - E_f(c) \approx \frac{1}{6}u^{'''}(\text{Skewness}_g - \text{Skewness}_f) + \frac{1}{24}(\text{Kurtosis}_g - \text{Kurtosis}_f) \quad (12)$$

In the case of CRRA utility function $u^{'''} = -\gamma(-\gamma-1)c^{-\gamma-2}$ while $u^{''''} = -\gamma(-\gamma-1)(-\gamma-2)c^{-\gamma-3}$. The first term is positive while the second term is negative for all $\gamma > 0$. Since $(\text{Skewness}_g - \text{Skewness}_f) < 0$ and $(\text{Kurtosis}_g - \text{Kurtosis}_f) < 0$ there are two potential offsetting forces – the excess kurtosis in the *heteroskedastic* economy being an undesired feature, while the positive skewness contributing positively. As we show below it turns out that the differences in skewness are relatively small, and for $\gamma > 2$ the kurtosis term dominates – which imply positive welfare gains. In summary, the analysis above should make it evident that once the mean and variance are equated across the two economies the remaining sources for welfare gains must show up through differences in higher moments. Our environment is obviously richer in the sense that agents are differentially borrowing constrained across economies and the degree to which they are able to smooth consumption is endogenous.

## A.2 Welfare Gains Computed Analytically

We start our analysis with the *heteroskedstic* economy (i.e. the economy with time varying variance of income). As in the previous section, we assume for simplicity that agents are infinitely lived and have CRRA utility function. Let $\sigma^2$ denote the *unconditional* variance of log-consumption and let the cross sectional variance of log-consumption be $\lambda$ times larger in a *recession* than it is in a *boom* – both events can occur with probability of 1/2 and are i.i.d across time. It is trivial then to show that the distribution of consumption follows,

$$\log{(c)} \sim \left\{ \begin{array}{ll} N\left(-\frac{1}{\lambda+1}\sigma^2, \frac{2}{\lambda+1}\sigma^2\right) & \text{in a boom} \\ N\left(-\frac{\lambda}{\lambda+1}\sigma^2, \frac{2\lambda}{\lambda+1}\sigma^2\right) & \text{in a recession} \end{array} \right\} \quad (13)$$

where the means for*log-consumption* are varying across recessions and booms so that the process satisfies $E(c) = E(c|\text{recession}) = E(c|\text{boom}) = 1$. The unconditional variance of the *level* of consumption then follows,

$$Var(c) = \frac{1}{2}[exp(\frac{2}{\lambda+1}\sigma^2) - 1] + \frac{1}{2}[exp(\frac{2\lambda}{\lambda+1}\sigma^2) - 1] \quad (14)$$

We now proceed to specify the consumption process in the *homeskedastic* economy. As in the previous section for our welfare comparisons to be meaningful we must ensure that the unconditional mean and variance in the homoskedastic economy are equal to their respective quantities in the heteroskedastic economy. We therefore specify the following process for consumption in the *homoskedastic* economy,

$$\log(c) \sim N(-\frac{\delta\sigma^2}{2}, \delta\sigma^2) \tag{15}$$

where $\delta > 1$ is a parameter that will scale the unconditional variance of log-consumption in the *heteroskedastic* economy to match the variances in levels across the two e-conomies. That is $\delta$ is the solution to

$$exp(\frac{2}{\lambda+1}\sigma^2) + exp(\frac{2\lambda}{\lambda+1}\sigma^2) = 2exp(\delta\sigma^2). \tag{16}$$

In order to derive the welfare gains of removing counter cyclical variation we need to calculate the expected utility under the two postulated consumption processes. The expected utility in the homoskedastic case is given by

$$\begin{aligned} E\{u(c)\} &= \frac{1}{1-\gamma} E \exp\left((1-\gamma)\log(c)\right) \qquad (17) \\ &= \frac{1}{1-\gamma} \exp\left(-\frac{\delta\sigma^2}{2}(1-\gamma) + (1-\gamma)^2 \frac{\delta\sigma^2}{2}\right) \\ &= \frac{1}{1-\gamma} \exp\left(\frac{\gamma}{2}(\gamma-1)\delta\sigma^2\right) \end{aligned}$$

since $E\{(1-\gamma)\log(c)\} = -\frac{\delta\sigma^2}{2}(1-\gamma)$ and $var\left((1-\gamma)\log(c)\right) = \delta\sigma^2(1-\gamma)^2$.

The expected utility in the heteroskedastic case is then given by

$$\begin{aligned} E\{u(c)\} &= \frac{1}{2}E\{u(c) \mid \text{boom}\} + \frac{1}{2}E\{u(c) \mid \text{recession}\} \qquad (18) \\ &= \frac{1/2}{1-\gamma}\left(\exp\left(\frac{\gamma}{2}(\gamma-1)\frac{2}{\lambda+1}\sigma^2\right) + \exp\left(\frac{\gamma}{2}(\gamma-1)\frac{2\lambda}{\lambda+1}\sigma^2\right)\right) \end{aligned}$$

Let the *welfare gain* of moving from an economy with heteroskedastic process for consumption to a homoskedastic one be defined as the multiplicative increase in consumption, $1+\psi$, required to make an agent be equally well off under the two processes. Thus, $\psi$ is implicitly defined by

$$\begin{aligned} E\{u(c(1+\psi)) \mid \text{heteroskedastic}\} &= E\{u(c) \mid \text{homoskedastic}\} \\ &\quad \text{so that} \\ \psi &= \left(\frac{E\{u(c) \mid \text{heteroskedastic}\}}{E\{u(c) \mid \text{homoskedastic}\}}\right)^{\frac{1}{\gamma-1}} - 1. \end{aligned}$$

Using equations (17) and (18), the welfare gain is computed as,

$$\psi = \left[\frac{\frac{1/2}{1-\gamma}\left(\exp\left(\frac{\gamma}{2}(\gamma-1)\frac{2}{\lambda+1}\sigma^2\right)+\exp\left(\frac{\gamma}{2}(\gamma-1)\frac{2\lambda}{\lambda+1}\sigma^2\right)\right)}{\frac{1}{1-\gamma}\exp\left(\frac{\gamma}{2}(\gamma-1)\delta\sigma^2\right)}\right]^{\frac{1}{\gamma-1}} - 1$$

$$= \left[\frac{1}{2}\exp(\frac{\gamma}{2}(\gamma-1)\sigma^2(\frac{\lambda}{\lambda+1}-\delta)+\frac{1}{2}\exp(\frac{\gamma}{2}(\gamma-1)\sigma^2(\frac{2\lambda}{\lambda+1}-\delta)\right]^{\frac{1}{\gamma-1}} - 1$$

Following our estimates in Exhibit 1 we set $\lambda = 4.8$. The welfare gains for different risk aversions, given different values for $\sigma^2$ (and therefore different values for $\delta$ solving (16)) are:

**Exhibit A.1**
**Welfare Gains of Removing Counter Cyclical Variation**

| | Welfare Gains ($\psi$) | | | | | |
|---|---|---|---|---|---|---|
| $\sigma^2$ | .1 | .2 | .3 | .4 | .5 | .621 |
| $\delta$ | 1.022 | 1.043 | 1.063 | 1.085 | 1.1055 | 1.1315 |
| $\gamma = 1.5$ | -.11% | -.41% | -.89% | -1.57% | -2.42% | -3.75% |
| $\gamma = 2.0$ | 0% | 0% | 0% | 0% | 0 % | 0 % |
| $\gamma = 2.5$ | 0.92% | 0.22% | 2.06% | 3.56% | 5.42% | 8.17% |
| $\gamma = 3.0$ | 2.51% | 0.63% | 5.51% | 9.40% | 14.01% | 20.68% |

where the highest value, $\sigma^2 = .621$ follows from $.621 = \sigma_\eta^2/\left(1-\rho^2\right)$ using the parameters we estimated in Exhibit 1. Figure **A**.1 displays the welfare gain $\psi$ for $\sigma^2 = .5$ for a continuous range of $\gamma$. Around $\gamma = 2$ the welfare gain is very small and initially increases non-linearly with $\gamma$ greater than 2, after which the rise in welfare is dominated by the kurtosis and pretty much rises linearly with $\gamma$.

**Figure A.1**
**Welfare Gains verus Risk Aversion: Analytical Computations**

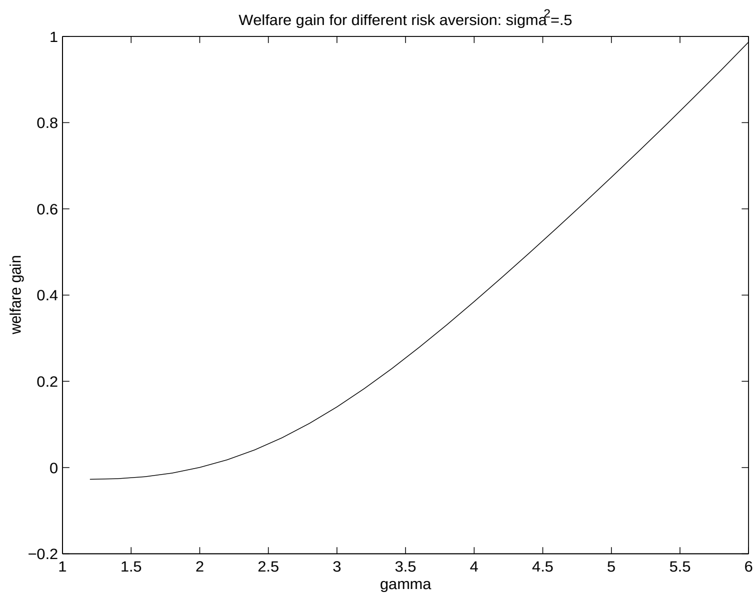


The figure depicts the expected 60 year transition of the capital stock when changing the heteroskedastic conditional variance shocks to homoskedastic shocks. Each point is the capital stock $n$ years after the transition as the average across different simulations (starting from different initial distributions). The capital stock falls by 2.8/during the first 30 years and settles in the steady state thereafter. Wage rate decrease by 1.1% on average and the interest rate ruse by 0.27% on average.

## Table 1
**Aggregate Moments: Baseline Economy**

*Panel A: Population Moments, Baseline Theoretical Economy*

| | Std Dev | Autocorrelation | Correlation with Output |
|---|---|---|---|
| Output | 0.022 | 0.18 | 1.00 |
| Investment | 0.045 | 0.20 | 0.76 |
| Consumption | 0.078 | 0.20 | 0.98 |

*Panel B: Sample Moments, Detrended U.S. Economy, 1955-1997*

| | Std Dev | Autocorrelation | Correlation with Output |
|---|---|---|---|
| Output | 0.022 | 0.52 | 1.00 |
| Investment | 0.085 | 0.36 | 0.85 |
| Consumption | 0.018 | 0.62 | 0.91 |

U.S. sample moments are based on annual NIPA data obtained from the Bureau of Economic Analysis, 1955-1997. Each series was detrended by applying the Hodrick-Prescott filter with a smoothing parameter of 100 to the natural logarithm of the deflated series. Theoretical moments are also based on logarithms and are computed as sample averages of a long simulated time series.

## Figure 1
## Cross-Sectional Moments by Age and Time

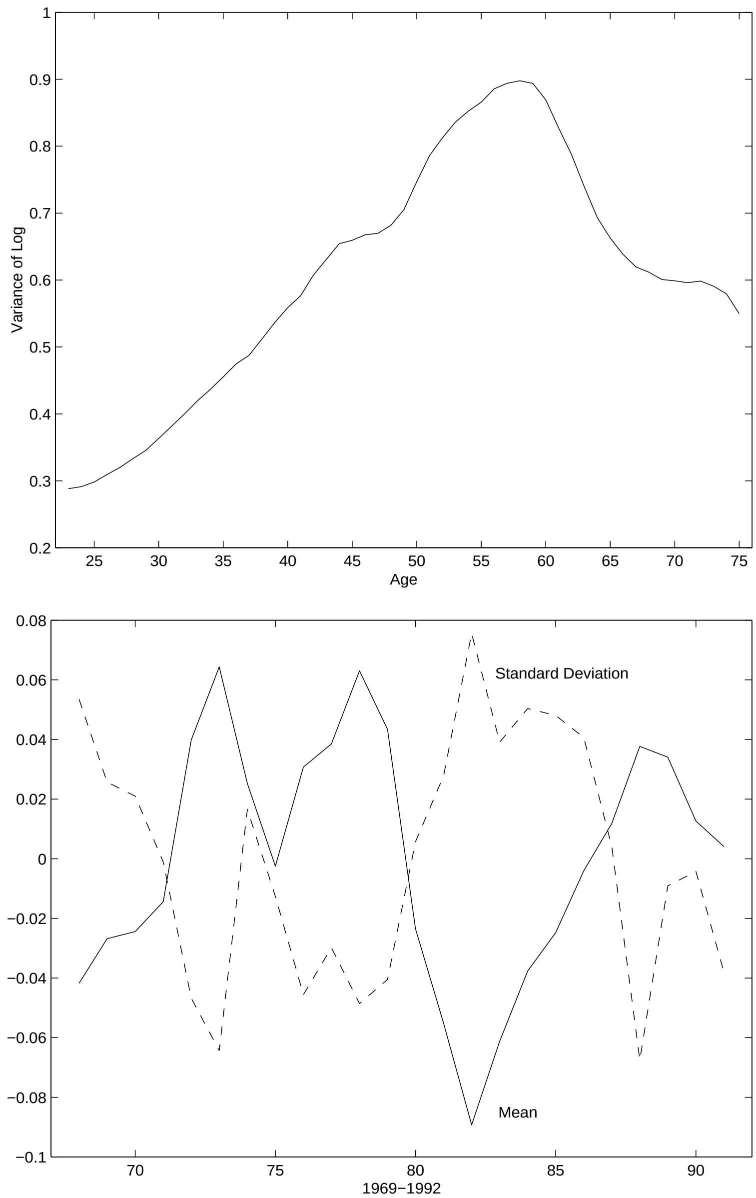


The top panel graphs the age profile of the cross-sectional variance of earnings, based on PSID data which has been pooled over the years 1969-1992. This graph controls for 'cohort effects' using the methods outlined in Deaton and Paxson (1994), and described in more detail in Storesletten, Telmer, and Yaron (2000). The bottom panel represents the (linearly detrended) cross-sectional mean and the cross-sectional standard deviation of PSID earnings, for the years 1969-1992. The standard deviation is additively scaled for graphical reasons.

## Figure 2
### Welfare Gains By Age

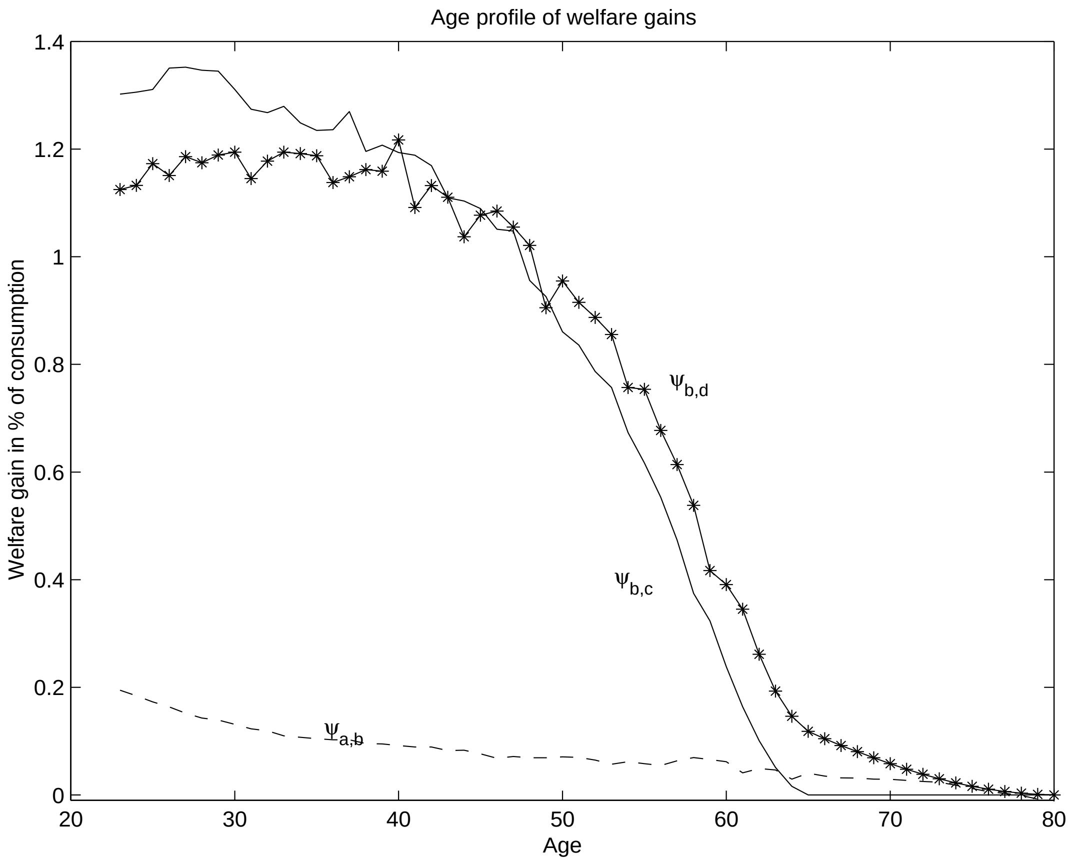


$\psi_{a,b}(h)$ depicts the age-profile of welfare gains of removing aggregate productivity shocks, while keeping the business cycle variation in individual risk.

$\psi_{b,c}(h)$ depicts the age-profile of welfare gains of removing business cycle variation in individual risk, but abstracting from general equilibrium effects of the transition.

Finally, $\psi_{b,d}(h)$ depicts the age-profile of welfare gains of removing business cycle variation in individual risk when general equilibrium effects of the transition have been accounted for.

**Figure 3**
**Welfare Gains By Age: General Equilibrium Effects**

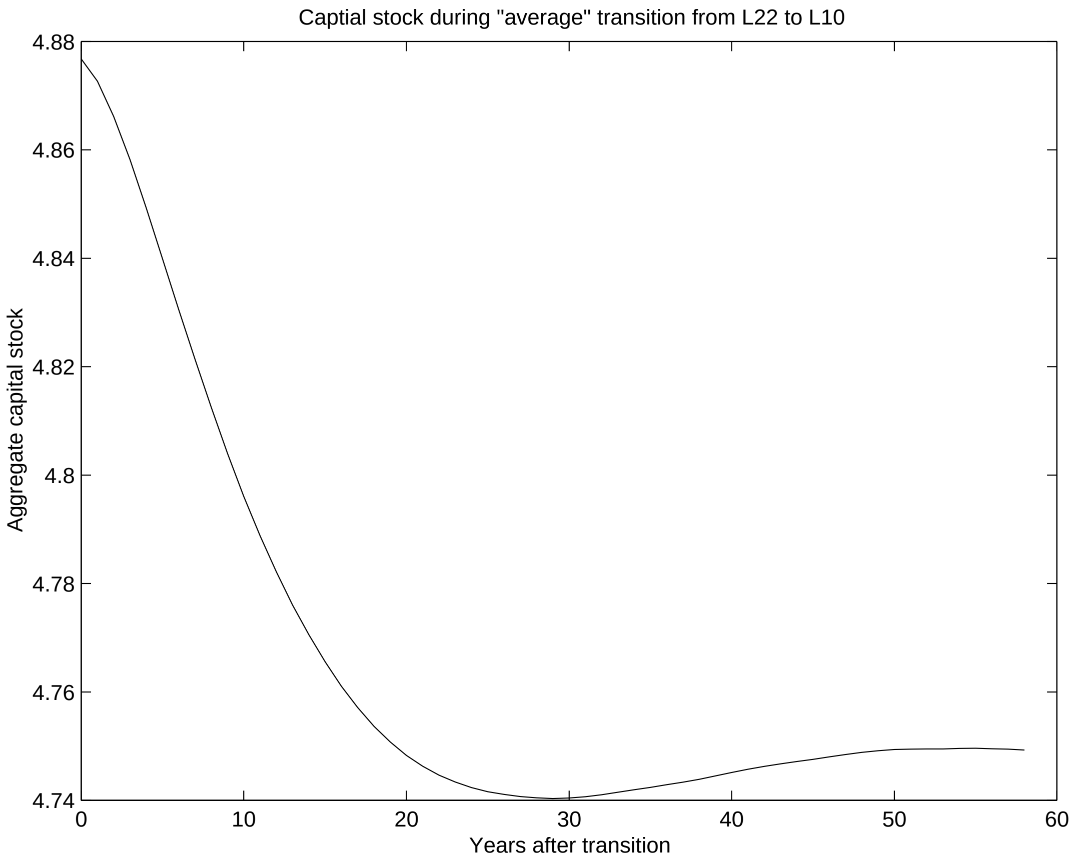


The figure depicts the expected 60 year transition of the capital stock when changing the heteroskedastic conditional variance shocks to homoskedastic shocks. Each point is the capital stock $n$ years after the transition as the average across different simulations (starting from different initial distributions). The capital stock falls by 2.8 % during the first 30 years and settles in the steady state thereafter. Wage rate decrease by 1.1% on average and the interest rate ruse by 0.27% on average.