INDIVIDUAL CONSUMPTION RISK AND THE WELFARE COST OF BUSINESS CYCLES*

MASSIMILIANO DE SANTIS†

Abstract. This paper measures the welfare gain from removing aggregate consumption fluctuations starting from an economy in which each individual faces both aggregate and idiosyncratic income shocks, and incomplete consumption insurance. We show that, because this welfare gain is a convex function of the overall consumption risk — aggregate plus idiosyncratic — each individual faces, to gauge the magnitude of the gain, it is important to match individuals' overall risk prior to any policy. We also show that the convexity of the welfare gain function increases substantially if individual consumption risk contains a realistic random walk component. While being agnostic about how much consumption risk countercyclical policy can remove, we show that in an economy calibrated to match individuals' overall risk, even removing ten percent of aggregate fluctuations results in a large welfare gain. We also review the previous literature that has found a low gain and argue that their estimates are low because they unrealistically assume that the idiosyncratic shocks to income are transitory. With transitory shocks, individuals can come close to perfect consumption insurance, thus undercutting the need for countercyclical policy.

*I am grateful to the co-editor, Mark Gertler, and three anonymous referees for their help in improving a previous version of the paper. I thank Tim Cogley, Jason Lepore, Bob Lucas, Martine Quinzii, and Aaron Smith for many helpful comments and discussions. I especially thank Louis Makowski, who helped me a lot from the very inception of this project, giving generously of his time.

†Department of Economics, Dartmouth College, 6106 Rockefeller Hall, Hanover, NH 03755. E-mail: massimiliano.desantis@dartmouth.edu.

1. Introduction

This paper measures the welfare gain from removing aggregate consumption fluctuations in an economy with idiosyncratic income shocks and incomplete consumption insurance, hence in an economy where individual risk is substantially greater than per capita risk. In contrast to previous literature that studied the welfare cost of fluctuations in models with idiosyncratic shocks, we focus directly on consumption risk rather than income risk. We show that, because the welfare cost of consumption fluctuations is an increasing and convex function of the overall level of risk that each individual faces, even removing only a small amount of *aggregate* fluctuation can lead to a large welfare gain.

A number of papers have attempted to estimate the cost of aggregate fluctuations. The seminal paper is Lucas (1987) using a representative agent economy. Lucas finds only a very small welfare improvement in removing aggregate consumption fluctuations of the size observed in U.S. data. The welfare improvement is equivalent to about one tenth of a percent of extra consumption at each date during the lifetime of an individual.

Lucas recognized that because of uninsurable idiosyncratic consumption shocks, household risk is larger than per capita risk (see Lucas 1987, p. 29). His estimate of overall household consumption risk (aggregate plus idiosyncratic) is three times per capita risk. Removing this amount of fluctuation would give a welfare gain equal to 6.8% of consumption at each date in Lucas' model.[1] But Lucas discounts this result because stabilization policies can be expected to eliminate no more than a small part of the uninsurable risk borne at the individual level.

As Lucas shows, the welfare gain function is convex in the level of risk faced by an individual, so that a greater fraction of the gain is attributable to removing initial portions of the risk he faces. This raises a natural question: starting from a higher level of individual risk —because of missing markets— what would be the gain from removing *only* aggregate risk? To answer this question, one needs to model explicitly uninsurable income risk, but given the 6.8% found in Lucas' simple calculation, there is potential for a large gain. This brings us to the first main finding of this paper. When idiosyncratic consumption risk is explicitly incorporated in a simple model, we show that the welfare gain function is even more convex than Lucas' function. This implies that removing only aggregate risk results in a much larger welfare gain, very close to Lucas' 6.8%.

Where does the increased convexity come from? It comes from assuming *permanent* idiosyncratic shocks. Lucas (2003) seems to agree on the permanence of such shocks:

[1]See Lucas 1987, Table 2 with risk aversion equal to 10.

> The fanning out over time of the earnings and consumption distributions within a cohort that Angus Deaton and Christina Paxson (1994) document is striking evidence of a sizable uninsurable random walk component in earnings. (p. 10)

A random walk component is crucial to an accurate description of the risk that each individual faces. Unlike transitory shocks, permanent income shocks cannot be insured away with standard saving technologies. As a consequence, individual consumption will inherit the volatility of income and, most importantly, the permanent (random walk) component.

Others also have studied the effect of idiosyncratic income shocks and incomplete markets. But typically, like Lucas, they find that the welfare gain from removing aggregate fluctuations is low (Imrohoroglu 1989, Krusell and Smith 1999, 2002). We will review these papers in more detail below, and we will argue that the reason for their low estimates is that their models do not generate enough volatility and persistence in individual consumption to be consistent with panel data evidence. In fact, in some cases like Imrohoroglu (1989), individuals have smoother consumption than Lucas assumed by matching aggregate U.S. consumption growth. Agents in these models are hit by idiosyncratic income shocks which cannot be fully insured. But because the shocks are not persistent, agents come close to full insurance by borrowing, lending, and/or saving. This allows an individual's consumption to be smoother than his income, and hence undercuts the need for policy. These papers typically do not check how much individual consumption fluctuation their models generated.

Because the welfare gain function is highly convex in the level of overall consumption risk, in quantifying the gain from removing aggregate fluctuations, it is important to match the overall risk each individual faces in the baseline economy prior to any policy, both in terms of volatility and persistence. To our knowledge, no one has previously pointed out that the baseline level of overall risk — aggregate plus idiosyncratic — is important for calculating the welfare gain from removing a marginal unit of aggregate risk.

One feature that distinguishes our work from almost all the previous literature is the way we address the question posed by Lucas. It is clear that if policy is to have a positive effect on welfare, it has to be able to ultimately affect individual consumption risk. But, because there is no generally accepted theory of how macro policy affects this risk, Lucas focused on a simpler thought experiment: he asked what the effect on welfare would be if aggregate consumption variation could be eliminated. Lucas considered one logical possibility, the elimination of short-lived aggregate consumption shocks around a deterministic trend when agents only face aggregate risk. We analyze many other possibilities and find that welfare gains are not trivial in many cases. The analysis of these richer, more realistic, cases does not come at the cost of increased complexity. In fact, as in Lucas, the welfare gains under all scenarios have easy to interpret closed form solutions.

In modeling idiosyncratic risk, we build on the framework of Constantidides and Duffie (1996). Their main no-trade theorem shows that if idiosyncratic income shocks are permanent, agents will not find it useful to trade in stocks and bonds to insure against such shocks. So, idiosyncratic income shocks translate into idiosyncratic consumption shocks. This explains why, relative to most of the literature, we find a much higher benefit from marginally lowering individual risk.

Our analysis is based on the assumption that individual shocks contain a martingale component. That this assumption holds empirically was observed in Lillard and Willis' (1978) pioneering work. Deaton and Paxson (1994), using panel data from three countries (U.S., U.K., and Taiwan), find that earnings and consumption tend to fan out over time within a cohort, implying the presence of a sizable random walk component in earnings. More recently, Meghir and Pistaferri (2004) document the presence of permanent shocks (martingales) to earnings in the PSID data. Beaudry and DiNardo (1991) also document a pattern of history dependence in labor market outcomes: when workers are employed in periods of high unemployment, their entry wages are much lower than the wages of workers employed in periods of lower unemployment, and this difference only disappears slowly over time. Thus, they relate the persistent component to expansions and recessions. Similarly, Storesletten, Telmer, and Yaron (2004) show evidence of a significant persistent component to household earnings using PSID data, which they relate to business cycle variations. They find a negative correlation over time between cross-sectional earnings means and standard deviations in PSID.

Because this evidence for countercyclical variation is still controversial, our baseline model will consider the case in which idiosyncratic risk is *independent* of aggregate risk. The baseline model shows that even in this case, provided aggregate risk is a random walk, removing only aggregate risk can result in substantial welfare gains. Thus our conclusion does not hinge critically on a correlation between the two sorts of risk. We then proceed to show that adding (what we view as) realistic correlation further increases the welfare gain.

How large the welfare gain from countercyclical policy will be, depends on how much overall risk one believes such policy can remove. Because it is an open question how much overall risk can be removed by countercyclical policies, we leave the task of answering this to future research. Atkeson and Phelan (1994) present a model in which removing aggregate fluctuations leaves overall consumption risk unaltered, thus making policy ineffective. On the other extreme, Beaudry and Pages (2001) present a model in which eliminating only aggregate productivity shocks also eliminates all idiosyncratic risk. We will be agnostic about this issue, and hence will measure the welfare gain under a variety of scenarios for aggregate risk, idiosyncratic risk, and the interactions between the two. We find that there is a region of plausible parameters for which removing only 10% of U.S. aggregate consumption variation yields a welfare gain greater than 0.5% of consumption at each date, a level that Lucas (1987) would have considered large.[2] In contrast with some of the literature, we find that if aggregate consumption is a random walk, this large gain does not depend on countercyclical idiosyncratic risk. We also find that if aggregate risk is correlated with idiosyncratic risk, even removing short lived consumption shocks around a deterministic trend yields a large welfare gain. We emphasize that, as in Lucas (1987), our results are based entirely on standard CRRA preferences, which makes our conclusions easy to interpret. It may be that larger gains can be obtained using Epstein-Zin preferences.[3]

Section 2 proceeds by performing Lucas' exercise for several scenarios about the types of risk and interactions among these. Section 3 relates our results to the existing literature. Section 4 concludes the paper.

2. Measuring The Cost of Fluctuations

Assume that individuals' preferences over consumption streams are represented by

$$E\left[\sum_{t=0}^{\infty}\beta^t\frac{(C_t^i)^{1-\gamma}}{1-\gamma}\right],$$

where β is a subjective discount factor and γ is the relative risk aversion coefficient.

Suppose further that the log of individual consumption, $\ln C_t^i$, is the sum of an aggregate plus an idiosyncratic stochastic process, that is:

$$\ln C_t^i = \ln C_t + \ln\delta_t^i, \tag{1}$$

where C_t is per-capita consumption, and δ_t^i is an idiosyncratic shock. In particular, δ_t^i is given by the martingale

$$\delta_t^i = \exp\left\{\sum_{s=1}^{t}\left(\eta_s^i y_s - \frac{y_s^2}{2}\right)\right\},$$

where y_s, $s = 1, \ldots, t$, is the cross sectional standard deviation of consumption growth at time s, known at time t, and η_s^i are idiosyncratic shocks, assumed to have a standard normal $N(0,1)$ distribution.[4] To see that y_t is the cross sectional standard deviation of consumption

[2] On page 29 Lucas says that "...something like one-half of 1 percent of total consumption. As deadweight losses go, *this* is a large number."

[3] See Lucas (2003) for a perspective on the use of this preference structure, and Bansal and Yaron (2004) for a model with Epstein-Zin preferences that stresses long run risks.

[4] Recall that for η normal $E[\exp(\eta k - (k^2/2))] = 1$, which implies that δ^i is a geometric martingale.

growth at time t, consider individual consumption growth between $t-1$ and t:

$$\frac{C_t^i}{C_{t-1}^i} = \frac{\delta_t^i}{\delta_{t-1}^i}\frac{C_t}{C_{t-1}} = \exp\{\eta_t^i y_t - \frac{1}{2}y_t^2\}\frac{C_t}{C_{t-1}}.$$

Therefore conditioning on C_t,

$$y_t^2 = \mathrm{Var}\left(\log\left(\frac{C_t^i/C_t}{C_{t-1}^i/C_{t-1}}\right)\right),$$

i.e., y_t^2 is the cross sectional variance of consumption growth. To fix ideas, it is useful to think about individual consumption as if the aggregates (C_t, y_t^2) are determined first, then the idiosyncratic shocks η_t^i are handed out.

The individual consumption process in (1) can be derived as an equilibrium consumption process in a standard finance economy, which is the essence of Constantinides and Duffie's theorem (1996). Appendix A briefly describes such a model. The Appendix also shows that both asset prices and the welfare cost of consumption fluctuations only depend on the stochastic behavior of C_t and y_t^2.[5]

To be able to calibrate the model and provide quantitative answers to our specific questions, we need to add to the Constantinides and Duffie framework assumptions about the stochastic process for C_t and y_t^2. We now provide a specification of the stochastic process that is simple enough to yield a closed form solution to the welfare gain from reducing aggregate fluctuations. Let $g_{t+1} = \Delta \ln C_{t+1}$. The process for g_t, y_t is as follows:

$$\begin{aligned} g_{t+1} &= \mu + \sigma\eta_{t+1} \\ y_{t+1}^2 &= \overline{y}^2 + b\sigma\eta_{t+1} + \sigma_u u_{t+1}, \end{aligned} \tag{2}$$

where the aggregate shock η_{t+1} is assumed to be i.i.d. with normal distribution $N(0,1)$. Hence σ is the standard deviation (volatility) of consumption growth, and aggregate consumption follows a geometric random walk. The shock u_{t+1} is assumed to be i.i.d. with normal distribution $N(0,1)$, and the parameter b allows for a correlation between the innovation to per capita consumption growth and and the variance of the idiosyncratic shock y_{t+1}^2. If $b = 0$, the two processes are independent; we will focus on this case initially, then move to the more realistic case where $b \neq 0$, calibrating b to existing empirical evidence.

The assumption that per capita consumption follows a random walk with drift is not innocuous in this context as we will see below, but it is a convenient starting point for our analysis of the cost of consumption fluctuations. A theoretical reason for this assumption is offered by Robert Hall (1978), and the presence of a unit root in per capita consumption is consistent with U.S. data. Further, the assumption is almost universal in the consumption

[5]See also Section 2.1 and Appendix B.

based asset pricing literature.[6] A contrasting assumption, considered by Lucas (1987), will be addressed in Section 2.6 below.

2.1. **Measuring the Welfare Gain.**
To compute Lucas' measure in this economy, we need to calculate the value of Δ such that

$$E_0\left[\sum_{t=0}^{\infty}\beta^t[(1+\Delta)C_t^i]^{1-\gamma}\right]=E_0\left[\sum_{t=0}^{\infty}\beta^t(\bar{C}_t^i)^{1-\gamma}\right], \tag{3}$$

where $\{C_t^i\}$ is the consumption stream in the economy with aggregate fluctuations, and $\{\bar{C}_t^i\}$ is the consumption stream in the economy without aggregate fluctuations. By an economy without aggregate fluctuations in this case we mean an economy in which aggregate consumption growth equals expected consumption growth in the economy with aggregate fluctuations. That is, $\bar{C}_{t+1}/\bar{C}_t = e^{\mu+\frac{1}{2}\sigma^2}$ with probability one, where $\bar{C}$ denotes aggregate consumption in the economy without aggregate shocks η_t. In contrast with Lucas' benchmark in which there is only aggregate risk, this is still a risky economy because agents remain subject to uninsurable idiosyncratic income shocks.

We can calculate the two expected utilities in (3). Consider the left side first. Multiplying and dividing by $(C_0^i)^{1-\gamma}$ yields

$$E_0\left[\sum_{t=0}^{\infty}\beta^t[(1+\Delta)C_t^i]^{1-\gamma}\right]=[C_0^i(1+\Delta)]^{1-\gamma}\sum_{t=0}^{\infty}E_0\beta^t\left(\frac{C_t^i}{C_0^i}\right)^{1-\gamma}.$$

In Appendix B we show that

$$[C_0^i(1+\Delta)]^{1-\gamma}\sum_{t=0}^{\infty}E_0\beta^t\left(\frac{C_t^i}{C_0^i}\right)^{1-\gamma}=\frac{[C_0^i(1+\Delta)]^{1-\gamma}}{1-A},$$

where

$$A\equiv A(\theta)\equiv\beta\exp\left\{\left[(1-\gamma)\mu+\alpha\bar{y}^2\right]+\frac{1}{2}\;\left[(1-\gamma)\sigma+\alpha b\sigma\right]^2+\frac{1}{2}\alpha^2\sigma_u^2\right\} \tag{4}$$

with $\alpha=\frac{1}{2}\gamma(\gamma-1)$ and $\theta=(\beta,\gamma,\mu,\sigma,\bar{y}^2,b,\sigma_u)$.

For the economy without aggregate fluctuations (the right side of (3)), the same calculations can be performed, with the assumption that $\eta_t=0$ with probability one. The right side will then become:

$$(C_0^i)^{1-\gamma}\sum_{t=0}^{\infty}E_0\beta^t\left(\frac{\bar{C}_t^i}{C_0^i}\right)^{1-\gamma}=\frac{(C_0^i)^{1-\gamma}}{1-A'},$$

[6] For a recent empirical analysis of aggregate consumption, see Ricardo Reis (2005). The data does not reject a unit root, although it rejects the random walk model in favor of a model in which consumption growth is positively serially correlated. This would increase our welfare gain. Krebs (2003) is an example of a production economy in which per capita consumption is a random walk in equilibrium.

where

$$A' \equiv A(\theta') \equiv \beta \exp\left\{(1-\gamma)(\mu + \frac{1}{2}\sigma^2) + \alpha\overline{y}^2 + \frac{1}{2}\alpha^2\sigma_u^2\right\} \tag{5}$$

and $\theta' = (\beta, \gamma, \mu, \sigma_u, \sigma)$.[7] Therefore, assuming both A and A' are less than unity, Δ is the solution to

$$\frac{[C_0^i(1+\Delta)]^{1-\gamma}}{1-\gamma}\frac{1}{1-A} = \frac{(C_0^i)^{1-\gamma}}{1-\gamma}\frac{1}{1-A'}.$$

So, Δ as a function of the parameters θ is given by

$$\Delta(\theta) = \left(\frac{1-A'}{1-A}\right)^{\frac{1}{\gamma-1}} - 1. \tag{6}$$

Notice that $\gamma > 1$ implies $A > A'$, hence $\Delta > 0$, i.e., the economy without aggregate fluctuations is strictly preferred.[8]

To begin, consider the special case in which $b = 0$ and $\sigma_u = 0$, so that the cross sectional dispersion y_t^2 is constant and equal to $\overline{y}^2$. Here, agents still face idiosyncratic shocks, but the shocks come from a constant distribution with variance $\overline{y}^2$. In this special case,

$$A = \beta \exp\left\{\left[(1-\gamma)\mu + \alpha\overline{y}^2\right] + \frac{1}{2}(1-\gamma)^2\sigma^2\right\}.$$

The convexity of the welfare gain in overall risk should now become clear. Notice that A is increasing in both $\overline{y}$ and σ. Consequently, when the variability of η is removed, the percentage change in utility (as A changes to A') will be larger when $\overline{y}$ and σ are larger. How important this convexity effect is for the magnitude of the gain depends on the parameters in θ and θ', which we estimate below. If either $b < 0$, $\sigma_u > 0$, or both, the welfare gain from removing fluctuations will be strictly greater than in this special case.

2.2. **Benchmark Calibration.**

Table 1 presents the chosen parameters in θ. The mean and the standard deviation of per-capita consumption growth, μ and σ, match the BEA data on real per capita consumption of non-durables and services for the period 1929-1998.

Our benchmark value for $\overline{y}^2$ is well below the cross sectional variation reported in Carrol (1992) or Storesletten, Telmer, and Yaron (2001). In his studies of precautionary saving, Carrol (1992, 1997) uses a level of 10% for the standard deviation of *permanent* shocks to income, after accounting for measurement error in PSID data.[9] We choose this value as our

[7] The presence of σ^2 in the formula for A' comes from the fact that in removing aggregate fluctuations, we equate consumption growth to mean growth in the economy with aggregate shocks, which depends on σ^2.

[8] The parameterizations chosen in our calibration exercises below imply that both A and A' are less than one.

[9] Storesletten, Telmer, and Yaron calibrate their model so that the cross sectional variance of consumption growth is 0.029, i.e. a standard deviation of 17%.

Table 1. Parameter Choices

Parameter	Symbol	Value
Mean consumption growth (%)	μ	1.89
Standard deviation of consunption growth (%)	σ	2.9
Mean idiosyncratic shock (%)	$\overline{y}^2$	$(10\%)^2$
Std. Dev. idiosyncratic shock	σ_u	0.00389
Covariation with aggregate risk	b	0
Risk aversion	γ	2,4
Implied log risk-free rate (%)	r^f	1.4
Subjective discount factor*	β	0.99, 0.95

*The subjective discount factor β is calculated so that we match a risk-free rate of 1.4%, given other parameters.

benchmark level, and we pair it with low values of risk aversion that many economists would agree on, i.e., $\gamma = 2$ and 4.

The benchmark value of σ_u, which represents the amount of variation in y_t^2, are chosen so that with 99% probability the cross sectional variance y_t^2 lies between zero (absence of heterogeneity) and $2\,\overline{y}^2$. When $\overline{y}^2 = 0.01$, this means that $\text{Prob}(0 \leq y_t^2 \leq 0.02) = 0.99$. Stated in terms of cross sectional standard deviation, this implies that with 99% probability the cross-sectional standard deviation of consumption growth will be between 0 and 14%. Notice that modeling the variance (y_t^2) as normally distributed, as opposed to the standard deviation (y_t), reduces the probability mass of values of y_t far from the mean $\overline{y}$. In our example with $\overline{y}^2 = 0.01$, $\text{Prob}(0 \leq y_t \leq 0.10) = \text{Prob}(0.10 \leq y_t \leq 0.14) = 0.499$. All these values are consistent with CEX data and are lower than the magnitudes assumed in Soresletten, Telmer and Yaron (2001).

The only parameter left to calibrate is β. In all cases, we chose the parameter β so that the model matches a risk free rate of 1.4%, a value consistent with time series data on the U.S. three month T-bill. Our assumption of log-normality for C_t and y_t^2 implies that the log risk free rate r_{t+1}^f, known at time t, is given by

$$r_{t+1}^f = \underbrace{-\ln\beta + \gamma\mu}_{\text{intertemporal substitution}} \quad \underbrace{-\widetilde{\alpha}\overline{y}^2 - \frac{1}{2}(\sigma\gamma - \widetilde{\alpha}b\sigma)^2 + \widetilde{\alpha}^2\sigma_u^2}_{\text{precautionary saving}}, \tag{7}$$

with $\widetilde{\alpha} = 0.5\gamma(\gamma+1)$. By contrast, notice that in Lucas' economy $r_{t+1}^f = -\ln\beta + \gamma\mu - \frac{\gamma^2\sigma^2}{2}$. Since σ is only about 3%, σ^2 is very small. So, even with high risk aversion, the precautionary saving term $\frac{\gamma^2\sigma^2}{2}$ is second order, implying an unrealistically large risk free rate when γ is large (the risk free rate puzzle). In our economy, at least within a range of plausible values, the risk free rate is decreasing in risk aversion, i.e. the precautionary saving motive is not second order.

2.3. **Results.**

We perform the welfare calculations presented above for our chosen benchmark parameterizations. Table 2 presents the results. In the left panel, we vary the relative risk aversion γ between 2 and 4, adjusting β to match the low risk free interest rate as described above. Welfare gains are large, 5.1% and 4.2% when γ is 2 and 4 respectively (see row $\Delta_{100\%}$).

TABLE 2. Welfare Gain from Removing Consumption Fluctuations

Welfare Measure	$\gamma = 2$	$\gamma = 4$	$\gamma = 2$	$\gamma = 4$
	$r^f = 1.4\%$		$\beta = 0.96$	
$\Delta_{100\%}$	0.051	0.042	0.017	0.047
$\Delta_{70\%}$	0.046	0.038	0.015	0.043
$\Delta_{50\%}$	0.038	0.031	0.012	0.035
$\Delta_{30\%}$	0.026	0.022	0.008	0.024
$\Delta_{10\%}$	0.010	0.008	0.003	0.009

Notes: γ is the coefficient of relative risk aversion, r^f the risk free interest rate, β the subjective discount factor. In the left panel, pairings of (γ, β) are chosen to match a risk free rate of 1.4%. The right panel presents results for $\beta = 0.96$ and changing risk aversion. $\Delta_{X\%}$ means that only X% of variation in η_t is removed.

These results differ from Lucas in two respects: (i) the inclusion of idiosyncratic risk, and (ii) the assumption that aggregate consumption follows a random walk rather than being trend-stationary. While our main interest is in (i), let us first focus on (ii). U.S. data confirms that aggregate consumption is almost a random walk, and not trend-stationary, so shocks to aggregate consumption have a permanent effect. How much of the large gain is due to (ii)? This can be calculated by setting $\overline{y}^2 = 0$ in our formula for Δ. For $\gamma = 4$, the gain would be 1.6% if policy could eliminate the random walk in aggregate consumption, even in an economy without idiosyncratic risk.

However, as Lucas pointed out, eliminating *all* aggregate variation should be regarded as an upper bound; it is hard to imagine that policy could eliminate all such variation. This skepticism seems all the more warranted when aggregate consumption is a random walk rather than trend-stationary (although there are models in the literature in which policy is so powerful). This motivates considering scenarios in which macro policy can only remove some fraction of aggregate variation. Rows denoted by $\Delta_{X\%}$ signify that only $X\%$ of variation in η_t is removed. With a risk aversion coefficient $\gamma = 2$ and cross sectional variance $y_t^2 = 0$, we obtain a welfare gain of 0.3% if policy can remove only 10% of the variation in aggregate consumption risk *and there is no idiosyncratic risk*; a relatively small number. This sets the stage for (i). As the table shows, the presence of idiosyncratic risk more than triples the welfare gain to 1% (see row $\Delta_{10\%}$). This is more than ten times larger than in Lucas' analysis with only aggregate risk and $\gamma = 10$.

The large welfare gain arises from the convexity of the welfare gain function. The presence of idiosyncratic risk leads each individual to face more total risk in the absence of policy. Hence any marginal decrease in total risk yields a substantially larger welfare gain. Similar results are obtained in the right panel of the Table, where we fix β to the commonly assumed value of 0.96, and vary risk aversion from 2 to 4.[10]

It is important to notice that the high welfare gain in Table 2 does not depend on the correlation between aggregate shocks and the cross sectional standard deviation of consumption growth, i.e. the parameter $b = 0$. We will see below that, once we realistically set $b < 0$, the welfare gain from removing even 10% of aggregate variability will be larger yet, in accord with the convexity of the welfare gain function. As pointed out in Section 2.1, even if $b = 0$, the welfare gain from removing variation in η_t increases with $\overline{y}$. It is of interest to contrast this implication for the welfare gain with the implication of idiosyncratic risk for the equity premium; the presence of idiosyncratic risk can help explain the high equity premium only if the dispersion of idiosyncratic risk is countercyclical (see Mankiw 1986, and Constantinides and Duffie 1996).

Appendix C presents additional estimates of the welfare gain from a parameterization of the model that meets some minimal requirement for consistency with stock market observations. The results are similar to the benchmark calculations.[11]

2.4. **Effect of Cyclical Variation in Idiosyncratic Risk.**

To evaluate the effect of the correlation between aggregate shocks η_t and the cross sectional variance of the idiosyncratic shock $\overline{y}^2$, we calculate the welfare gain assuming $b < 0$: negative aggregate shocks are, on average, accompanied by greater cross sectional heterogeneity in consumption growth. The empirical evidence for significant negative correlation ($b < 0$) is found in Storesletten, Telmer, and Yaron (2004), and Meghir and Pistaferri (2004). The

[10] In all cases, the mean of consumption growth is kept constant, $E_t(C_{t+1}/C_t) = e^{\mu+0.5\sigma^2}$.

[11] The minimal requirement is that the model matches the Sharpe ratio on the S&P 500.

models of Soresletten, Telmer, and Yaron (2001) and Krebs (2003) underscore the importance of the correlation for large welfare gains.[12]

Table 3. Effect of Cyclical Variation in Idiosyncratic Risk on the Welfare Gain

Welfare Measure	$\gamma = 2$	$\gamma = 4$	$\gamma = 2$	$\gamma = 4$
	$b = -0.13$		$b = -0.81$	
$\Delta_{100\%}$	0.056	0.055	0.081	0.110
$\Delta_{70\%}$	0.051	0.050	0.074	0.101
$\Delta_{50\%}$	0.042	0.042	0.061	0.085
$\Delta_{30\%}$	0.029	0.029	0.041	0.059
$\Delta_{10\%}$	0.011	0.011	0.015	0.023

Notes: γ is the coefficient of relative risk aversion, b the regression coefficient of y_t^2 on g_t. In all cases the subjective discount factor β is chosen to match a risk free rate of 1.4%. $\Delta_{X\%}$ means that only X% of variation in η_t is removed.

We calibrate b in two ways. A simple way is to assume that all the variation in y_t^2 (equal to σ_u in the benchmark model) depends on the aggregate shock η_t,

$$y_{t+1}^2 = \overline{y}^2 + b\sigma\eta_{t+1},$$

and set $b = -\sigma_u/\sigma$, where the value of σ_u is taken from Table 2. This generates a value of $b = -0.13$. Results from this exercise are presented on the left panel of Table 3.

We can also get a rough measure of b from the empirical literature cited above. Using NBER business cycle dates and BEA per capita consumption data 1929-2005, consumption growth is about 2.9% during expansions, and -0.8% during contractions. Using the same NBER indicator, Storesletten, Telmer, and Yaron (2004) finds that y_t^2 is $(21\%)^2$ during expansions and $(12\%)^2$ during contractions. Their estimates thus imply a value of $b = -0.81$. The right panel of Table 3 presents the results.

Relative to the results of Table 2 the welfare gain does not increase much for $b = -0.13$, but increases substantially if $b = -0.81$.

2.5. Why is the Potential Gain so High?

Figure 1 summarizes our findings by plotting the welfare gain Δ as a function of σ for different values of $\overline{y}^2$ and b. The Figure allows us to identify the relative contribution of each main

[12]Krebs (2005) points out that the results in Storesletten et al. (2001) and in Krebs (2003) may be biased upward because they used estimates from the *working paper* version of Storesletten, Telmer, and Yaron (2004). The latter authors revised their estimates of the parameter b downward in the published version of their paper. Our calibration is not subject to this criticism.

FIGURE 1. Individual Risk and Welfare Gains

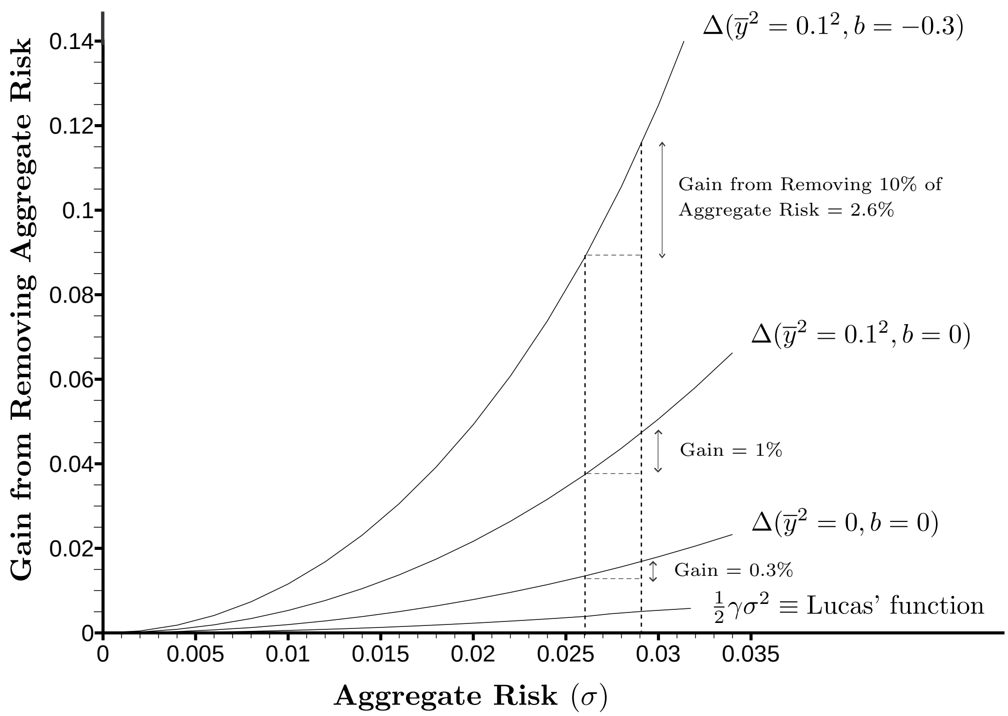

The function $\Delta(\theta)$ is calculated from the definition of $\Delta(\theta)$ in the body of the paper for $\gamma = 4$, $\beta = 0.96$, and other parameters as in the benchmark model. $\frac{1}{2}\gamma\sigma^2$ is Lucas' welfare gain function. $\Delta(\overline{y}^2 = 0)$ indicates absence of heterogeneity, and $\Delta(\overline{y}^2 = 0.1^2)$ is the benchmark model of Table 2. The case $\Delta(\overline{y}^2 = 0.1^2, b = -0.3)$ introduces some correlation between aggregate and idiosyncratic risk.

component of individual risk: the random walk in aggregate consumption, the level of idiosyncratic risk $\overline{y}^2$, and the degree to which idiosyncratic risk depends on aggregate risk—the coefficient b. With a random walk in per capita consumption, the welfare gain Δ is already more convex in σ than with Lucas' trend stationary process, even when there is no cross sectional heterogeneity (see curve $\Delta(\overline{y}^2 = 0, b = 0)$). But notice the increase in convexity as $\overline{y}^2$ increases, illustrated by curve $\Delta(\overline{y}^2 = 0.10^2, b = 0)$. This is why removing only 10% of aggregate fluctuation yields a large welfare gain, as seen in Table 2. Finally, taking account of the fact that shocks to per capita consumption growth are correlated to the permanent idiosyncratic shock ($b < 0$), convexity increases further, leading to an even larger welfare gain.

The Figure makes clear how crucial it is to properly characterize an individual's total risk *prior* to any policy, in order to correctly evaluate the welfare gain from marginally removing some aggregate risk. Removing 10% of aggregate risk is given by the change in σ between the two vertical dashed lines in the Figure (from right to left). The double arrows to the right show the welfare gain from removing the 10% in each case. The gain is highest (2.6%) if, in addition to idiosyncratic risk, we assume a realistic value of b.

2.6. Mean Reverting Shocks to Aggregate Consumption.

The results above rest on the assumption that per capita consumption follows a random walk with drift. Because of the nature of this process, removing aggregate shocks amounts to removing transitory shocks to per capita consumption *growth*, but permanent shocks to the *level* of per capita consumption.

It is worth emphasizing that, like Lucas (1987), we are not describing policies that would remove aggregate risk, but only evaluating the consequences. In this spirit, and because U.S. per capita consumption is very close to a random walk, the scenarios illustrated in Figure 1 are interesting and informative, emphasizing our theme that, in order to accurately evaluate any business cycle policy, it is crucial to first match the level of overall risk faced by individuals prior to policy.

One could take the rather pessimistic view that stabilization policy can only remove consumption shocks that are short-lived. How large is the gain from removing shocks of this type? To answer this question, let us take the extreme view that macro policy can only remove aggregate shocks that last one period. That is, assume aggregate consumption in the absence of policy would follow a trend stationary process with i.i.d. shocks:

$$\ln C_t = \delta + \mu t + \sigma \eta_t. \tag{8}$$

The essence of this process is that if log consumption is hit by a negative shock η_t, it will revert to the linear trend at $t+1$ if there were no subsequent shocks. This means that consumption growth between t and $t+1$ will be greater than average to bring consumption back to trend, i.e, consumption growth is strongly negatively correlated. This is not an accurate description of U.S. consumption growth data, which is positively rather than negatively correlated. But it is a convenient shortcut to analyze the welfare gain from removing very short lived cyclical shocks. This seems to be the path taken in Lucas (1987).

With y_t^2 specified as in (2), the welfare gain in this special case is

$$\Delta_T = \left(\frac{1 - A'_T}{1 - A_T}\right)^{\frac{1}{\gamma - 1}} \exp\left\{\frac{1}{2}\sigma^2\gamma - \alpha b \sigma^2\right\} - 1, \tag{9}$$

where

$$\begin{aligned} A_T \equiv\ & \beta \exp\{(1-\gamma)\mu + \alpha \overline{y}^2 + \frac{1}{2}\alpha^2(b^2\sigma^2 + \sigma_u^2)\} \quad \text{and} \\ A'_T \equiv\ & \beta \exp\{(1-\gamma)\mu + \alpha \overline{y}^2 + \frac{1}{2}\alpha^2\sigma_u^2\}. \end{aligned}$$

The subscript "T" stands for "trend stationary." Apart from a proportionality constant, $(A_T)^t$ is the t^{th} element of the utility function in the economy with trend stationary aggregate shocks, and $(A'_T)^t$ is the t^{th} element in the economy without such shocks.

Notice from the presence of $b^2\sigma^2$ in the formula for A_T that, even if aggregate risk is small and short lived, it has a long lasting effect through its correlation with idiosyncratic risk b. Thus a transitory shock to aggregate consumption is converted into a permanent shock at the individual level.

For reasonable levels of b, the welfare gain is substantial even if policy can remove only these one-period self-adjusting shocks, as Table 4 shows. If $b = -0.81$ (the value of b found by Storesletten et al.), the gain is 1.4% and 4.1% when γ is 2 and 4, respectively. For the U.S., this corresponds to a per capita gain in 2004 dollars of \$386 and \$1,130, respectively. Even if b is set to the conservative value of -0.13, the gain is one order of magnitude larger than Lucas'. With risk aversion $\gamma = 4$, the gain is 0.043%, which corresponds to about \$119 per capita in 2004 dollars, and about \$35 billion in aggregate.

Table 4. Welfare Gain When Aggregate Consumption is Trend Stationary

Welfare Gain Δ_T		
	$\gamma = 2$	$\gamma = 4$
$b = -0.13$	0.14%	0.43%
$b = -0.81$	1.4%	4.1%

Notes: γ is the coefficient of relative risk aversion, b the regression coefficient of y_t^2 on g_t. In all cases the subjective discount factor β is chosen to match a risk free rate of 1.4%.

Notice from the welfare formula above that, in the special case $b = 0$, the factors A_T and A'_T are equal, and hence the welfare gain from removing business cycles equals

$$\Delta_T = e^{\frac{1}{2}\sigma^2\gamma} - 1,$$

which is exactly Lucas' result. In this special case, not only is aggregate risk predictable and short lived, but also has no effect on idiosyncratic risk. Theoretically, it is an interesting special case of the model, but it is hardly realistic since it assumes the substantial permanent shocks at the individual level are completely unrelated to recessions and expansions.

The effect of $\overline{y}^2$ and b on Δ_T is shown in Figure 2. For any given $b < 0$, the convexity of the welfare gain function increases with $\overline{y}^2$; and potential welfare gains are very high: see curves $\Delta_T(\overline{y}^2 = .10^2, b = -.3)$ and $\Delta_T(\overline{y}^2 = .12^2, b = -.3)$. The convexity of the welfare function Δ_T also increases substantially with the absolute value of b as illustrated by the curve $\Delta_T(\overline{y} = .10^2, b = -0.5)$.

Figure 2. Welfare Gains with Trend Stationary Per Capita Consumption

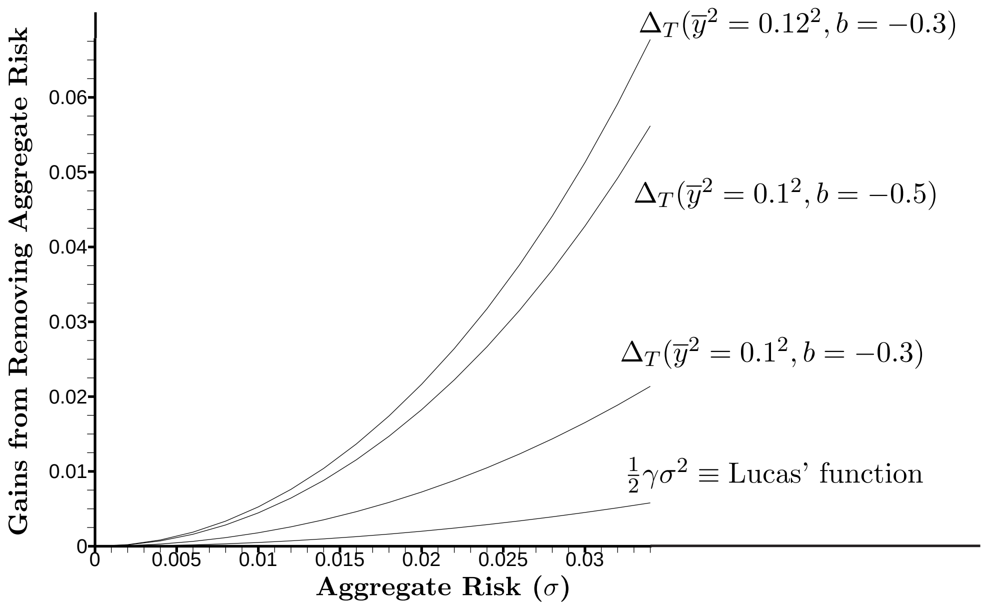

The function $\Delta_T(\theta)$ is calculated from the definition of $\Delta_T(\theta)$ in the body of the paper for $\gamma = 4$, $\beta = 0.96$, and other parameters as in the benchmark model. $\frac{1}{2}\gamma\sigma^2$ is Lucas' welfare gain function, which obtains if $b = 0$. The case $\Delta_T(\overline{y}^2 = 0.1^2, b = -0.3)$ introduces some correlation between aggregate and idiosyncratic risk. Relative to this case, the two higher curves correspond to higher $|b|$, $b = -0.5$, and higher heterogeneity, $\overline{y}^2 = 0.12^2$, respectively.

3. Discussion and Relation to Existing Literature

In this section we review some of the previous literature that studied the welfare cost of business cycles in the context of incomplete consumption insurance. We want to make two basic points. First, in most of the literature, the idiosyncratic income shocks are not persistent. Consequently, individuals can insure themselves by using storage or saving, instruments which allow them to come pretty close to complete insurance. Indeed, in some cases, the level of consumption risk faced by each individual in equilibrium is lower than the per-capita risk assumed by Lucas. This is counterfactual since panel data shows large *permanent* shocks and significant consumption risk. This point explains why most of the previous literature finds a low welfare gain from removing aggregate fluctuations. As we have seen in both Figures 1 and 2, the size of the gain depends crucially on the level of risk faced by each individual in the benchmark economy without policy, not just on the amount of risk policy can remove. If individuals face only a small amount of risk even without policy — because they have self-insured — naturally macro policy can only contribute a small additional benefit.

The second point is, even granted that individuals face significant consumption risk in the absence of macro policy, it remains an open question how effective macro policy can be in reducing their risk. Atkeson and Phelan (1994) present a model in which eliminating aggregate fluctuations does not alter individual risk at all. On the other extreme, Beaudry and Pages (2001) present a model in which aggregate fluctuations are at the heart of individual risk. Hence, removing the business cycle eliminates individual risk completely. This is why, in our simulations, we considered multiple scenarios.

A paper that illustrates the first point is Imrohoroglu (1989). Her work was directly motivated by the fact that individuals face incomplete insurance markets and cannot perfectly insure against idiosyncratic risk. Thus she departed from Lucas' representative agent model. But, even allowing for idiosyncratic shocks, Imrohoroglu finds very low welfare gains from removing aggregate fluctuations. Indeed, in some cases she considers, the welfare gain is even *smaller* than Lucas' estimate. Although each individual is hit by only partially insurable income shocks, she includes a storage technology and costly borrowing in her model, instruments which permit each individual to smooth his consumption more than in Lucas' model, effectively eliminating the need for countercyclical policy. The reason individuals in Imrohoroglu's model can reach a high level of consumption insurance using storage and costly borrowing is that her shocks to income are *transitory.* In terms of our Figures, in Imrohoroglu's model, prior to any policy, agents move to a low welfare gain function by self insuring —recall that decreasing individual risk results in a lower, less convex Δ. Thus, Imrohoroglu's gain calculations can be thought of as involving a movement along the lower curve in Figure 1 on the very flat section close to the origin.[13]

The papers of Krusell and Smith (1999, 2002) also illustrate our first point. They argue, while it is plausible that the average gain for the typical household is low, removing fluctuations may be very beneficial for some fraction of the population. Therefore, they study the distribution of welfare gains from removing fluctuations across a population consisting of infinitely lived agents affected by three shocks: an aggregate productivity shock, an idiosyncratic income shock (the loss of a job), and an idiosyncratic preference shock that affects the subjective discount factor. Both aggregate shocks and idiosyncratic shocks to income are transitory. Individuals have access to a saving technology that converts one unit of consumption into one unit of capital and pays the aggregate marginal product of capital. In equilibrium, patient agents save more than impatient ones. But, because income shocks for the most part are transitory, all agents find it optimal to use the saving technology, thus smoothing consumption. With different wealth levels, the welfare gain from removing the business cycle

[13] In fact, the gain she finds is so low that it must be the case that agents can also insure some of the aggregate shocks before policy.

will differ across the population. Krusell and Smith find that an unborn agent is better off in the economy without business cycles, but the gain is very small, 0.14%. This value does not capture distributional effects; so Krusell and Smith also calculate welfare gains for agents with different wealth levels and employment status, by evaluating the change in utility resulting from transitioning to a new steady state without aggregate fluctuations. They find that for most groups the welfare gains are negative.

Why are the gains in Krusell and Smith so low? In both economies they consider, a cyclical economy without policy and a non-cyclical economy with policy, agents save to self insure against idiosyncratic shocks. Thus most likely, as in Imrohoroglu, Krusell and Smith find that the gain from removing the business cycle is so low because most individuals have already self-insured. In terms of our Figure 1, prior to any policy, by self insuring against idiosyncratic shocks, agents move to a lower welfare gain function. They also insure against the aggregate shock, thus moving to a flatter portion towards the origin.[14] Thus, removing aggregate variation only results in very small gains.

We now turn to the second point that emerges from the literature, that it is an open question how effective macro policy can be in reducing individual risk.

Atkeson and Phelan (1994) argue that the welfare gain from removing aggregate risk is zero, regardless of how much individual risk agents face, because removing aggregate risk leaves individual risk unaltered. Thus, even if an individual finds himself on the steep portion of a highly convex welfare gain function, removing aggregate fluctuations will not change his position on the curve, because aggregate risk is replaced by higher idiosyncratic risk, and hence will not benefit him at all. Atkenson and Phelan construct a simple model in which this can be the case: Suppose the aggregate state can assume two values, high or low, and that the state determines only the probability of becoming unemployed. If aggregate fluctuations are removed, and as a result the unemployment rate equals the *mean* unemployment rate of the economy with aggregate fluctuations, on average agents lose their jobs as many times as before. So their ex-ante income streams are the same in both cases, and removing aggregate shocks does not yield any welfare gain.

On the other extreme, Beaudry and Pages (2001) show that removing aggregate risk can completely remove individual risk.[15] In their model, because workers lack the ability to commit to a firm, aggregate risk causes firms to offer wage contracts that are downwardly rigid — to insure risk averse agents against downward risk, and upwardly flexible — to keep agents from quitting in expansions when labor markets are tight. Because the wage is downward rigid, a

[14] This can be deduced from the saving behavior before and after removing aggregate fluctuations in Krussel and Smith's economy.

[15] This situation, in which $y_t^2 = 0$, is even more extreme than the scenarios we considered.

laid off worker who re-enters the labor force with a lower wage contract will have lower wages than if he had not been laid off. Thus negative wage shocks are persistent.[16] On the other hand, without aggregate productivity shocks, at equilibrium firms would offer a constant wage, and agents would choose to allocate their labor between firms and a home technology to equalize their marginal product in the two sectors (in terms of consumption units), thus completely eliminating income risk. Hence, in this economy, removing aggregate fluctuations removes the *persistent* consumption risk completely. Although an individual employed by a firm may become unemployed because of a re-allocation shock, the firm pays him the consumption value of his marginal product in the household sector, therefore the individual is indifferent about losing his job.

Two papers in the literature that are very close to ours are Storesletten, Telmer, and Yaron (2001) and Krebs (2003). In Storesletten, Telmer, and Yaron (2001), agents live finite lives and are hit by age specific earning shocks which, at least in part, are highly persistent. The aggregate shock determines the variance of the persistent component of the age specific shocks: the variance is higher with low realizations of the aggregate shock ($b < 0$). In their model, removing the aggregate shock means (i) the aggregate shock is set to its unconditional mean, and (ii) the variance of individual shocks is made independent of aggregate shocks, thus reducing individual risk further. Their age specific income process does not contain a random walk, but it does have persistence. In terms of the scenarios we analyzed, their model is closest to the benchmark model with per capita consumption being a random walk.

Krebs (2003) extends the infinite horizon Constantinides and Duffie (1996) framework to include production. Agents face individual specific martingale shocks as in our case, and economy-wide technology shocks that are correlated with aggregate shocks. Thus, removing aggregate shocks also eliminates the cyclicality of the cross sectional dispersion y_t^2. Krebs' model implies that, in equilibrium, per capita consumption follows approximately a random walk. Thus, as in Storesletten et al., removing aggregate fluctuations removes persistent shocks to per capita consumption.[17]

One difference between the above two models and ours is that they focus exclusively on the correlation between aggregate and idiosyncratic shocks. We have shown that when per capita consumption follows a random walk (or close to it), $b < 0$ is not necessary for large welfare gains. Nevertheless the case $b < 0$ is more realistic; hence we view their analyses as basically complementary to ours. The important difference lies in the level of complexity:

[16]This seems to be consistent with the evidence in Bils (1985), and Beaudry and DiNardo (1991, 1995).

[17]Krebs 2005 is another interesting example of how a reduction of individual risk from removing business cycles could come about. The paper focuses on welfare gains from removing cyclical variation in the long-term earning losses of displaced workers.

their models are much more complex, and hence what drives their conclusions is much less transparent. Like Lucas', our model can be solved analytically, while their models cannot, requiring numerical solution procedures. So, even the large welfare gains that they find are hard to interpret. Unlike our Figure 1, there is no simple, convex welfare gain function that they can point to. A fortiori, they cannot compare their conclusions with Lucas' by saying that, once individual risk is explicitly modelled by including a random walk component to individual earnings, the welfare gain function becomes much more convex, hence the gain from policy becomes much larger.

It is worth amplifying on this important difference, which also applies to most of the other papers in the literature. We have addressed Lucas' question by focusing directly on the consumption process; by contrast, the cited literature analyzes production economies. The latter approach is useful to get a concrete feeling of how changes in individual risk could come about. But it also requires assumptions about the workings of a production economy, and hence adds complexity, which forecloses the possibility of simple closed-form solution and analysis.

To summarize, what we learn from extant literature is that when shocks to income are only transitory, individuals can reduce consumption risk using simple storage or saving technologies. This makes any macro policy that would further reduce aggregate risk almost unnecessary. But this is counterfactual since panel data reveals a high level of cross sectional consumption risk, and the presence of a sizable random walk component to individual earnings and consumption. Lucas' argument for considering idiosyncratic risk is that individual risk should be greater than per capita risk because of incomplete markets, not lower. Where the persistent and uninsurable shocks come from we do not know at this time, although Beaudry and Pages offer one story. Even accepting the presence of high consumption risk and a sizable random walk component to earnings shocks, we still cannot say how much of this risk aggregate policy can remove. Given the contrary examples in Atkeson and Phelan and in Beaudry and Pages, the plausible thing to do is to consider multiple scenarios.

4. Concluding Remarks

In this paper we show that, to evaluate the welfare gain from removing aggregate fluctuations, it is essential for a good model to first replicate the actual consumption risk that individuals face in the absence of any policy.

Aggregate consumption data follows nearly a random walk; further, panel data shows there is a sizable individual-specific random walk to income. Accordingly, we have constructed a simple endowment economy that incorporates these features. The model is consistent with the high and persistent consumption risk observed in panel data. Also, in contrast to Lucas', our model is consistent with the market price of risk (maximal Sharpe ratio) implied by the stock market and with a low risk free rate. As in Lucas, the model is kept at the simplest level, so it yields a closed form solution to the welfare gain from removing aggregate risk.

Unlike Lucas, we find that the welfare gain from removing aggregate risk is large. The main reason is that, with CRRA preferences, this welfare gain is a convex function of the overall risk that each individual faces. Since individual risk is larger than per capita risk, we find that the gain from removing only aggregate fluctuations is two orders of magnitude larger than in Lucas' exercise. The welfare gain is large if policy can remove only 10% of unpredictable shocks to per capita consumption growth, independent of the correlation between aggregate and idiosyncratic shocks. The welfare gain also is large if policy can only remove short-lived shocks, provided these shocks are related to individuals' idiosyncratic income shocks.

Lucas suggested that we consider seriously his estimates as an upper bound to the "... marginal social product of additional advances in business cycle theory" (1987, p. 27). If Lucas had used a simple model that included persistent uninsurable shocks to income, he possibly would have reached a very different conclusion. Our analysis suggests that additional advances in business cycle theory may have a large return.

Using a non-parametric model, in the sense that it abstracts from the utility maximization problem of agents, and focusing on asset pricing observations such as the Sharpe ratio, Alvarez and Jermann (2004) also find high costs to consumption fluctuations. But Alvarez and Jermann conclude that removing business cycles would only lead to small welfare improvements. This is because they estimate that only a small fraction of fluctuations in aggregate consumption comes from business cycles frequencies. Alvarez and Jermann's definition of business cycles excludes low frequency, one-of-a-kind events. This is a very different definition than that used by many economists, and also by the NBER: Many aggregate fluctuations are due to one-of-a-kind events like the 1970's oil shocks, which nevertheless may be softened by appropriate macro policies.

Our model shows that the potential social marginal product of advances in business cycle theory — broadly conceived — is large. The important open question is how much individual risk aggregate policy can remove. For addressing it, the work of Beaudry and Pages (2001) and Atkeson and Phelan (1994) seems an interesting starting point. Even assuming no correlation between aggregate and idiosyncratic risk, our baseline scenario shows it is important to know how much stabilization policy can affect the permanent component of aggregate risk; perhaps a model with endogenous growth, in which high frequency fluctuations have permanent effects, can throw light on this.

Appendix A. A model with Uninsurable Individual Risk

Consider a standard finance model, an exchange economy with a single non-durable consumption good and two traded assets, a risk-free discount bond and a risky equity. Bonds are issued at time $t-1$, matured at t, and each bond has a par value of one. We assume the bond is in zero net supply. The risky equity (whose net supply we normalize to be one) pays dividend D_t and has ex-dividend price P_t. Each consumer i is endowed with labor income I_t^i and consumes C_t^i at time t. Aggregate labor income is I_t, and aggregate consumption is $C_t = I_t + D_t$. It is assumed that $I_t + D_t > 0$ for all times t. There is an infinite set of distinct consumers denoted by $\mathcal{A}$. At time t, consumer i holds a portfolio of shares of the risky asset θ_t^i and of the bond b_t^i. The time t budget constraint is:

$$C_t^i + \theta_t^i P_t + b_t^i \leq I_t^i + \theta_{t-1}^i (P_t + D_t) + b_{t-1}^i R_t^f, \tag{10}$$

where R_t^f denotes the return on a bond issued at $t-1$. Consumers have homogeneous preferences represented by a time-separable von Neumann-Morgenstern utility function with constant relative risk aversion coefficient γ and a constant subjective discount factor β. At time 0, each consumer maximizes

$$E\left[\frac{\sum_{t=0}^{\infty} \beta^t (C_t^i)^{1-\gamma}}{1-\gamma} \mid \mathcal{F}_0\right] \tag{11}$$

subject to the sequence of budget constraints (10) by choosing a sequence $(\theta^i, b^i, C^i) \equiv (\theta_t^i, b_t^i, C_t^i),\ t = 0, 1, 2, \ldots$

An equilibrium is a security price and bond return process (P, R^f), and strategies $\{(\theta^i, b^i, C^i) : i \in \mathcal{A}\}$ for the consumers such that

(i) (θ^i, b^i, C^i) maximizes (11) subject to (10)
(ii) markets clear, i.e., $\sum_{i\in\mathcal{A}} \theta_t^i = 1$ and $\sum_{i\in\mathcal{A}} b_t^i = 0$ for all t.

Market clearing implies that $\sum_{i\in\mathcal{A}} C_t^i = C_t \equiv I_t + D_t$ for all t.

An equilibrium price process for the risky asset will satisfy the following condition for all $i \in \mathcal{A}$:

$$P_t = E\left[\beta \left(\frac{C_{t+1}^i}{C_t^i}\right)^{-\gamma} (P_{t+1} + D_{t+1}) | \mathcal{F}_t\right], \tag{12}$$

where the expectation is taken conditional on $\mathcal{F}_t$, the information set at time t.

Labor income (I_t^i in (10)) is defined by

$$I_t^i = \delta_t^i C_t - D_t$$

The process δ_t^i is the following martingale:

$$\delta_t^i = \exp\left\{\sum_{s=1}^{t} \left(\eta_s^i y_s - \frac{y_s^2}{2}\right)\right\},$$

where y_t is the cross sectional standard deviation of consumption growth, and it depends on aggregates at t. The aggregates (C_t, y_t^2) are determined first, then the idiosyncratic shocks η_t^i are handed out, where η_t^i is assumed to be standard normal $N(0,1)$. With this income process agents do not find it useful to trade in stocks nor bonds: the two instruments do not provide insurance against the martingale shocks to income. Constantinides and Duffie (1996) show the existence of a no-trade equilibrium for this labor income process, i.e. each agent just consumes his income ($C^i = D_t + I_t^i = \delta^i C_t$). Individual consumption growth then is:

$$\frac{C_{t+1}^i}{C_t^i} = \frac{\delta_{t+1}^i}{\delta_t^i}\frac{C_{t+1}}{C_t} = \exp\{\eta_{t+1}^i y_{t+1} - \frac{1}{2}y_{t+1}^2\}\frac{C_{t+1}}{C_t}. \tag{13}$$

Therefore, given normality of η_t^i,

$$y_{t+1}^2 = \mathrm{Var}\left(\log\left(\frac{C_{t+1}^i/C_{t+1}}{C_t^i/C_t}\right)\right),$$

i.e., y_t^2 is the cross sectional variance of consumption growth.

Appendix B. Solving for Δ with Uninsurable Individual Risk

We derive the welfare gain for the general case in which $(1-\phi L)\ln C_t = \mu + \delta t + \sigma\eta_t$. From the left hand side of (3), multiply and divide by $(C_0^i)^{1-\gamma}$:

$$[C_0^i(1+\Delta)]^{1-\gamma}\sum_{t=0}^{\infty} E_0\beta^t\left(\frac{C_t^i}{C_0^i}\right)^{1-\gamma} = [C_0^i(1+\Delta)]^{1-\gamma}\sum_{t=0}^{\infty} E_0\left[\beta^t\exp\left\{(1-\gamma)c_t^i - c_0^i\right\}\right].$$

The generic element in the sum can be written as

$$\begin{aligned} E_0\beta^t\exp\left\{(1-\gamma)\sum_{s=1}^{t}\Delta c_s^i\right\} &= E_0\beta^t\exp\left\{(1-\gamma)(\ln C_t - \ln C_0) + (1-\gamma)\sum_{s=1}^{t}(\ln\delta_s^i - \ln\delta_{s-1}^i)\right\} \\ &= E_0\beta^t\exp\left\{(1-\gamma)(\ln C_t - \ln C_0) + (1-\gamma)\sum_{s=1}^{t}(\eta_s^i y_s - \frac{1}{2}y_s^2)\right\}, \end{aligned}$$

where the last equality is obtained using the definition of δ_t^i.

Making use of the law of iterated expectations

$$\begin{aligned} &E_0\beta^t\exp\left\{(1-\gamma)\sum_{s=1}^{t}g_s + (1-\gamma)\sum_{s=1}^{t}(\eta_s^i y_s - \frac{1}{2}y_s^2)\right\} = \\ &\quad E_0\left[\beta^t\exp\left\{(1-\gamma)\sum_{s=1}^{t}g_s\right\}E_0\left[\exp\left\{(1-\gamma)\sum_{s=1}^{t}(\eta_s^i y_s - \frac{1}{2}y_s^2)\right\}|g_1,\ldots,g_t,y_1,\ldots,y_t\right]\right]. \end{aligned}$$

Given that the η_s^i are normally distributed, i.i.d. variables, we can compute the expectation conditional on $g_1,\ldots,g_t,y_1,\ldots,y_t$.[18] The expectation

$$E_0\left[\exp\left\{(1-\gamma)\sum_{s=1}^{t}(\eta_s^i y_s - \frac{1}{2}y_s^2)\right\}|g_1,\ldots,g_t,y_1,\ldots,y_t\right] = \exp\left\{\frac{1}{2}(\gamma-1)\gamma\sum_{s=1}^{t}y_s^2\right\}. \tag{14}$$

[18] Recall that if $\ln X$ is distributed as $N(\mu,\sigma^2)$, $Ee^X = e^{\mu+0.5\sigma^2}$.

Therefore the generic element

$$(15)\quad E_0\left[\beta^t \exp\left\{(1-\gamma)\sum_{s=1}^{t}\Delta c_s^i\right\}\right] = E_0\left[\beta^t \exp\left\{(1-\gamma)(\ln C_t - \ln C_0) + \frac{1}{2}(\gamma-1)\gamma\sum_{s=1}^{t} y_s^2\right\}\right].$$

Notice that we can compute this expectation for all t, as it is the mean of a log-normally distributed random variable.[19]

Under the random walk, this yields:

$$(16)\qquad A^t \equiv \left[\beta \exp\left\{\left[(1-\gamma)\mu + \alpha\overline{y}^2\right] + \frac{1}{2}\left[(1-\gamma)\sigma + \alpha b\sigma\right]^2 + \frac{1}{2}\alpha^2\sigma_u^2\right\}\right]^t,$$

where $\alpha = \frac{1}{2}\gamma(\gamma-1)$. Notice this function is increasing and convex in $\overline{y}^2$ and b if $\gamma \geq 1$. This implies that the welfare function Δ will be increasing and convex in $\overline{y}^2$ and b.

For the economy without aggregate fluctuations, the generic element in the sum will be:

$$(17)\qquad A' \equiv \beta \exp\left\{(1-\gamma)(\mu + \frac{1}{2}\sigma^2) + \alpha\overline{y}^2 + \frac{1}{2}\alpha^2\sigma_u^2\right\}.$$

With trend stationary consumption, the expectation is,

$$\begin{aligned} A(t) &\equiv \left[\beta \exp\{(1-\gamma)\mu + \alpha\overline{y}^2 + \frac{1}{2}\alpha^2\sigma_u^2 + \frac{1}{2}\alpha^2 b^2\sigma^2\}\right]^t \times \exp\left\{\frac{1}{2}\left[(1-\gamma)^2\sigma^2 + \alpha b\sigma^2(1-\gamma)\right]\right\} \\ &\equiv A_{TS}\exp\left\{\frac{1}{2}\left((1-\gamma)^2\sigma^2 + \alpha b\sigma^2(1-\gamma)\right)\right\} \end{aligned}$$

And for the economy without aggregate fluctuations,

$$A'(t) \equiv \left[\beta \exp\{(1-\gamma)\mu + \alpha\overline{y}^2 + \frac{1}{2}\alpha^2\sigma_u^2\}\right]^t \exp\{\frac{1}{2}(1-\gamma)\sigma^2\}.$$

Appendix C. Other Parameterizations and Asset Prices

Since the problem of pricing risky assets is closely related to the question of assessing the costs of instability, it is useful to calibrate our model so that it meets some minimal requirement for consistency with stock market observations. This minimal requirement is the Hansen and Jagannathan bound (1991), which says that the maximal Sharpe ratio of an economy should be greater than or equal to the largest observed Sharpe ratios (such as the one on the S&P 500).[20]

In Table 5, we calculate the measure Δ for a set of parameters θ and θ' consistent with the Hansen and Jagannathan bound. As well known in the asset pricing literature, this typically

[19] Make the substitution $X = (1-\gamma)(\ln C_t - \ln C_0) + \frac{1}{2}(\gamma-1)\gamma\sum_{s=1}^{t} y_s^2$, and use the result of footnote above.

[20] The Sharpe ratio of a risky asset s is defined as the excess return per unit of volatility, $E_t(R_{t+1}^s - R_{t+1}^f)/\sigma(R_{t+1})$, where R_{t+1}^s is the return on asset s, and R_{t+1}^f is the return on the risk free asset. The maximal Sharpe ratio is the maximum Sharpe ratio that a model can generate.

involve relative risk aversion values greater than 4. Our values of γ are on the high end of plausible values, but are not extreme and some estimates do exceed 10.[21]

Table 5. Welfare Gain from Removing Consumption Fluctuations

Welfare Measure	$\gamma = 10$ $\overline{y}^2 = .061^2$	$\gamma = 7$ $\overline{y}^2 = .10^2$	$\gamma = 5$ $\overline{y}^2 = .136^2$
$\Delta_{100\%}$	0.074	0.036	0.023
$\Delta_{70\%}$	0.069	0.033	0.021
$\Delta_{50\%}$	0.059	0.028	0.017
$\Delta_{30\%}$	0.043	0.019	0.012
$\Delta_{10\%}$	0.018	0.08	0.004

Notes: γ is the coefficient of relative risk aversion, and $\overline{y}^2$ is the cross sectional variance of consumption growth. Pairings of $(\gamma, \overline{y}^2)$ are chosen to match the Sharpe Ratio of the S&P 500. Thus the right panel presents results for lower values of risk aversion using the intermediate level of average cross sectional standard deviation, $\overline{y} = 10\%$. $\Delta_{X\%}$ means that only $X\%$ of variation in η_t is removed.

We pair $\gamma = 10,\ 7,$ and 5 with $\overline{y}^2 = 0.00372,\ 0.01,$ and 0.0184 respectively so that higher levels of risk aversion are paired with lower levels of idiosyncratic risk. These levels of variance correspond to standard deviations of 6.1%, 10%, and 13.6% respectively. As shown in De Santis (2005), these parameterizations can generate a low risk-free interest rate (1.4%) and the high Sharpe ratio observed for the S&P 500. The intuition is that, when there are permanent idiosyncratic shocks, the precautionary saving motive is strong, which resolves the risk free rate puzzle and generates a high price of risk (the maximal Sharpe ratio).

As before, $b = 0$ —cross sectional dispersion independent of aggregate shocks— and values of σ_u are chosen so that with 99% probability the cross-sectional variance y_t^2 lies between zero (absence of heterogeneity) and $2\,\overline{y}^2$. In the extreme case in which $\overline{y}^2 = 0.0184$, the standard deviation of consumption growth will be between 0 and 19% with 99% probability.[22]

With a risk aversion coefficient $\gamma = 10$ and cross-sectional variance equal to $(6.1\%)^2$, we obtain a welfare gain of 1.8% by removing only 10% of aggregate variation. The gain when $\gamma = 5$ and only 10% of aggregate fluctuations are removed is 0.4%. Notice that the welfare gain is not always greater than in Table 2.

[21]See for example Parker and Julliard (2005).

[22]Storesletten et al. model $\overline{y}^2$ as two state Markov process, and find that the high variance is 21%. Our value is not only lower, but much less frequent.

References

[1] F. Alvarez and U. Jermann. Using asset prices to measure the cost of business cycles. *Journal of Political Economy*, 112(6):1223–1256, December 2004.

[2] F. Atkeson and C. Phelan. *Reconsidering the Cost of Business Cycles with Incomplete Markets*, pages 187–207. Cambridge, MA: MIT Press, 1994.

[3] Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59:1491–1509, 2004.

[4] Paul Beaudry and John DiNardo. The effect of implicit contracts on the movements of wages over the business cycle: Evidence from micro data. *Journal of Political Economy*, 99(4):665–688, 1991.

[5] Paul Beaudry and John DiNardo. Is the behavior of hours worked consistent with implicit contract theory? *Quarterly Journal of Economics*, 110:743,768, 1995.

[6] Paul Beaudry and Carmen Pages. The cost of business cycles and the sabilization value of unemployment insurance. *European Economic Review*, 45:1545–1572, 2001.

[7] M. Bils. Real wages over the business cycle: Evidence from panel data. *Journal of Political Economy*, 93:668–689, 1985.

[8] Christopher D. Carrol. The buffer-stock theory of saving: Some macroeconomic evidence. *Brookings Papers on Economic Activity*, (2):61–156, 1992.

[9] Christopher D. Carrol. Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly Journal of Economics*, 112(1):1–55, 1997.

[10] G. Constantinides and D. Duffie. Asset pricing with heterogeneous consumers. *Journal of Political Economy*, 104:219–40, 1996.

[11] Massimiliano De Santis. Interpreting aggregate stock market behavior: How far can the standard model go? Dartmouth College, manuscript, August 2005.

[12] Angus Deaton and Christina Paxson. Intertemporal choice and inequality. *The Journal of Political Economy*, 102(3):437–467, 1994.

[13] A. Imrohoroglu. The cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political Economy*, 97:1364–1383, 1989.

[14] Tom Krebs. Growth and welfare effects of business cycles in economies with idiosyncratic human capital risk. *Review of Economic Dynamics*, 6:846–868, 2003.

[15] Tom Krebs. Job displacement risk and the cost of business cycles. manuscript, May 2005.

[16] P. Krusell and A. A. Jr. Smith. On the welfare effects of eliminating business cycles. *Review of Economic Dynamics*, 2:245–272, 1999.

[17] P. Krusell and A. A. Jr. Smith. Revisiting the welfare effects of eliminating business cycles. Working Paper, Carnegie Mellon University, 2002.

[18] Lee A. Lillard and Robert Willis. Dynamic aspects of earning mobility. *Econometrica*, 46(5):985–1012, Sep 1978.

[19] Robert E. Jr. Lucas. *Models of Business Cycles*. Blackwell, New York, 1987.

[20] Robert E. Jr. Lucas. Macroeconomic priorities. *American Economic Review*, (4):1–14, 2003. AEA Presidential Address.

[21] Costas Meghir and Luigi Pistaferri. Income variance dynamics and heterogeneity. *Econometrica*, 72(1):1–32, 2004.

[22] Jonathan Parker and Christian Julliard. Consumption risk and the cross section of stock returns. *Journal of Political Economy*, 113(1):185–222, February 2005.

[23] Ricardo Reis. The time-series properties of aggregate consumption: Implications for the cost of fluctuations. Princeton University, manuscript, 2005.

[24] K. Storesletten, C.I. Telmer, and A. Yaron. The welfare cost of business cycles revisited: Finite lives and cyclical variations in idiosyncratic risk. *European Economic Review*, 45:1311–1339, 2001.

[25] K. Storesletten, C.I. Telmer, and A. Yaron. Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112:695–717, 2004.