Ambiguity Aversion, Asset Prices, and the Welfare Costs of Aggregate Fluctuations*

Irasema Alonso†
Yale University

and

Jose Mauricio Prado‡
IMT Lucca Institute for Advanced Studies

January 27, 2008

Abstract

Under the hypothesis that aggregate U.S. consumption is random and, more importantly, viewed as *ambiguous* by consumers, we examine the implications for asset prices and for how consumption fluctuations influence consumer welfare. First, considering a simple, Mehra-Prescott-style endowment economy with a representative agent facing consumption fluctuations calibrated to match U.S. data, we study to what extent ambiguity aversion can deliver asset prices that are consistent with data: a high return on equity and a low return on riskfree bonds. For some configurations of preference parameters – a discount factor, a degree of relative risk aversion, and a measure of ambiguity aversion – we find that it can. We then use these parameter configurations to investigate how much consumers would be willing to pay to reduce endowment fluctuations to zero, thus delivering a Lucas-style welfare cost of fluctuations. These costs turn out to be very large: consumers can be willing to pay over 10% of consumption in permanent terms.

*We are grateful to Per Krusell for very insightful comments. We also thank Hal Cole, David DeJong, and participants in seminars and conferences at IIES, University of Pittsburg, the 2007 Society of Economic Dynamics Meeting, and the 2006 North American Summer Meeting of the Econometric Society. Prado gratefully acknowledges financial support from Jan Wallander's and Tom Hedelius' Research Foundation. Any errors are, of course, ours.

†Corresponding Author; Department of Economics, Yale University, Box 208268, New Haven, CT 06520-8268, USA; e-mail: irasema.alonso@yale.edu

‡IMT Lucca Institute for Advanced Studies, Piazza S. Ponziano 6, 55100 Lucca, Italy; e-mail: jm.prado@imtlucca.it

1 Introduction

The typical assumption in the macroeconomic literature is to consider aggregate U.S. consumption to be random. We take this view here, but we also entertain the assumption that consumers are uncertain about the randomness, as opposed to just viewing it as risky. More precisely, we assume that consumers display *ambiguity aversion* toward aggregate dividends. Ambiguity aversion, which is a way of formalizing preferences that are consistent with the Ellsberg paradox, captures a form of violation of Savage's axioms of subjective probability. Instead, consumers behave as if a range of probability distributions is possible and as if they are averse toward the "unknown". With the typical parameterized representation of ambiguity aversion, consumers have minmax preferences, thus maximizing utility based on the worst possible belief within some given set of feasible beliefs. Thus, the "pessimistic" beliefs are endogenous and fully rational and the axiomatic underpinnings behind ambiguity aversion allows straightforward welfare analysis.

Under ambiguity aversion, consumers evaluate macroeconomic fluctuations differently. First, asset prices will reflect the perceived ambiguity. In particular, a small amount of randomness gives *first-order effects* on utility if there is ambiguity about this randomness. Thus, ambiguity aversion contrasts the standard model, where risk aversion only leads to second-order effects on utility. Second, the implied costs of small "business cycles", i.e., of random consumption fluctuations, also are of first order. In this paper, we investigate the asset prices and welfare costs of business cycles in more detail: we derive asset prices and discuss their quantitative nature, and we assess the implied welfare costs of fluctuations.

The first step in our work is to look at asset pricing in a simple Mehra-Prescott-style endowment economy. Here, we demonstrate how one can obtain larger equity premia by assuming ambiguity aversion, along with low riskfree rates. The key parameter in the model is the amount of ambiguity aversion, but it interacts nonlinearly with other parameters, such as with the coefficient of relative risk aversion. There is no direct evidence of which we are aware that allows us to calibrate the ambiguity parameter, but we show a range of calibrations that match the average returns on risky and riskless assets. This range of calibrations includes one without ambiguity aversion but, as has been pointed out in the literature, the implied risk aversion is very high, and the subjective discount rate has to be negative.[1] With a sufficient amount of ambiguity, however, the required risk aversion parameter is much lower, and a positive subjective

[1] For a survey on the equity premium, see Kocherlakota (1996). Aiyagari (1993) shows how incomplete markets together with costly asset trading both increases the equity premium and lowers the riskfree rate. DeJong and Ripoll (2007) examine asset prices from the perspective of consumers with time-inconsistent preferences; they do not find major effects.

discount rate can be accommodated. Moreover, we use additional return data—the volatility of the short-term bond return and the average return on long-run bonds—to suggest that ambiguity aversion really does allow the model to match the overall return data better.

The second step of the work is to ask, given that preferences are such that the model asset prices are in line with the data, how consumers then assess the fluctuations from a welfare point of view. Thus, we redo the Lucas-style calculation whereby one asks by how much the representative-consumer utility would rise (expressed as a permanent increase in consumption) if all fluctuations around the trend were eliminated. The answer, in the economy with ambiguity, naturally depends also on the amount of ambiguity: since ambiguity is a form of "worry" about random fluctuations, the elimination of the randomness would eliminate the worry, and consumers would be better off as a result. We show that the costs of business cycles can be sizable, indeed significantly over 10% of consumption in permanent terms.

The gains from eliminating fluctuations can be thought of in two parts in this model: the gains from making consumption smooth rather than random, or fluctuating, and the gains from eliminating the pessimism that ambiguity aversion leads to when consumption is random. Interestingly, due to the interplay between these two effects, the gains from eliminating fluctuations, can be decreasing in the coefficient of risk aversion unless the discount factor is sufficiently low. This is because the gain from removing pessimism is higher, the less curved is period utility: a removal of pessimism moves probability mass from low-consumption to high-consumption states, which amounts to a maximal gain when utility is linear.

As explained, thus, the approach here is to use the observed asset prices as a way of selecting the ambiguity parameter. That is, we use the first step in our work as a calibration, and then do the Lucas (1987) calculation based on it. In this sense, our approach is similar to that pursued by Alvarez and Jermann (2004), who find a way of computing the welfare costs of cycles directly based on information about the market price of risk. The main difference between their work and ours is that whereas they do not discuss specific forms of preferences, our main purpose here is to present an explicit parametric class of preferences that, with an axiomatic underpinning and back-up evidence at least on the experimental level, can deliver the observed asset prices without implausible risk aversion or discounting. Moreover, their focus is on marginal costs of fluctuations; our results, in contrast, are about their total costs, and for such calculations more knowledge of preferences is needed. Closer to our approach in that sense is Tallarini (2000), who uses Epstein-Zin preferences (see Epstein and Zin (1989)), thus distinguishing between the intertemporal rate of substitution and risk aversion in a way that we do not do here, in order to simultaneously address asset prices and the welfare costs of cycles. Alvarez and Jermann also make a distinction between stock-market risk and consumption risk that is important and that we for simplicity ignore here; moreover, the present paper also uses the riskfree rate, as well as suggests other features of bond returns, to restrict preferences.

Our finding of a large quantitative effect is very significant compared to what the rest of the literature has found; Lucas found about a tenth of 1%, and most studies since have explored different preference settings or a different stochastic representation than that used by Lucas, but most of these studies report numbers that are much lower than those we find. For example, Bansal and Yaron (2004) emphasize how "long-run risk" can influence asset prices significantly, and Croce (2006) builds further on this hypothesis by incorporating production and by looking at the implied welfare costs of this risk, which turn out to be large.[2] It should be emphasized in this context that the present study uses the Mehra-Prescott data and does not alter preferences beyond the introduction of ambiguity and is in this sense conservative; the introduction of long-run risk and Epstein-Zin preferences would presumably allow even larger welfare costs of risk (though perhaps it is less clear whether this risk ought to be called business-cycle risk).

A few studies explore incomplete markets, as do Krusell and Smith (1999); in an economy with realistically modeled wealth heterogeneity, they find that eliminating aggregate risk results in an average welfare gain of about 0.1 of 1% of consumption, although there is significant diversity of outcomes across subgroups in the population.[3] Related work by Storesletten, Telmer, and Yaron (2001) obtains larger estimates of the gains from removing aggregate shocks in an overlapping-generations economy with incomplete markets economy and productivity heterogeneity but only by about 5 times larger than do Krusell and Smith. Finally, Abel (2002) examines how asset prices can be understood from the perspective of "pessimism"; he does not propose an axiomatization behind this, nor behind the other possibility he considers ("doubt", amounting to a mean-preserving spread in the perceived distribution of asset returns), nor does he discuss the link to the welfare costs of consumption fluctuations, but his setting is formally close to the present one.

Section 2 describes the basic model, Section 3 the asset pricing, and Section 4 the elimination of fluctuations. Section 5 concludes.

[2]For recent accounts, see Reis (2006) and Otrok et al. (2004); see also Lucas (2003) for a survey discussion.

[3]Mukoyama and Şahin (2006) give evidence of larger effects when workers also differ in productivity levels and unemployment probabilities.

2 The economy

As in Mehra and Prescott (1985), we use a variation of the Lucas (1978) pure exchange economy. Production is exogenous: the economy has a tree that pays dividends every period.[4] The dividend grows at a random rate. This rate has a two-state support given by (λ_1, λ_2) and follows a first-order Markov process. The transition probabilities are given by $\phi_{ss'}$ – the probability of going to state s' if today's state is s, with $s, s' = 1, 2$.

When the consumer is ambiguous about these probabilities, he perceives them to be

$$\Phi(v) = \begin{pmatrix} \phi_{11} - v_1 & \phi_{12} + v_1 \\ \phi_{21} - v_2 & \phi_{22} + v_2 \end{pmatrix},$$

where $v_s \in [-a, a]$ $(s = 1, 2)$ with restrictions on a such that all probabilities are in [0,1]. The parameter a measures the amount of ambiguity in the economy.

Preferences of an ambiguity-averse consumer are given by the maxmin formulation

$$V_t(s^t) = u(c(s^t)) + \beta \min_{\pi \in \Pi_{s^t}} E_\pi V_{t+1}(s^{t+1})$$

where c is consumption, $u(c)$ is the period utility function, and Π_{s^t} is a set of transition probability laws given the history s^t today.

Aversion to ambiguity is captured by the "minimization" part in the utility formulation above: the consumer behaves with pessimism, i.e., he assumes the worst possible probability distribution. For an axiomatic foundation for this preference formulation, see Gilboa and Schmeidler (1989) for the static setting and Epstein and Wang (1994) and Epstein and Schneider (2003) for a multiperiod setting.

3 Asset pricing

We begin with a simple example, based on logarithmic preferences and *iid* shocks, that illustrates the key connections between preferences and asset prices, and we then proceed to the more general model: a CRRA period utility function and serially correlated shocks. Then we calibrate the economy and report the model's performance.

3.1 A simple example

Consider an ambiguity-averse representative agent with a logarithmic period utility function and discount factor β. In this section, shocks are *iid* and symmetric, i.e., $\phi_{ss'} = \phi = 0.5$ for all (s, s').

[4]The setting employed here is similar to a one-country version of that in Alonso (2007). There, asset prices are determined from an intertemporal production technology that is linear.

There are an equity share that is competitively traded and a riskless bond that is in zero net supply. We denote by b and e the consumers' bond and equity holdings, respectively.

The representative agent holds the tree and, thus, his consumption every period is the dividend of the tree. A log-period utility function together with the assumption of *iid* and symmetric shocks imply that p, the price of the tree, will be linear in d, the dividend, and independent of the state: $p(d) = \hat{p}d$.

The ambiguity-averse consumer puts a higher weight on the bad outcome than what is warranted by the objective probability; that is, he becomes pessimistic because he is worried about that outcome and does not know its probability.

We assume that $\lambda_1 > \lambda_2$ so the bad outcome is state 2 – the outcome where the dividend is low. The objective probability of this state is 0.5 but the consumer's belief is $\phi(v) = 0.5 + v$ and he chooses v from the set $v \in [-a, a]$.

The problem of the representative agent with wealth today given by w is

$$V(w) = \max_{e} \log(w - p(d)e) + \beta \min_{v \in [-a,a]} [(\phi - v)V(w_1') + (1 - \phi + v)V(w_2')]$$

subject to

$$w_1' = [\lambda_1 d + p(d\lambda_1)]e,$$

$$w_2' = [\lambda_2 d + p(d\lambda_2)]e.$$

Here, for ease of notation, we have excluded the bond (since bond holdings have to be zero in equilibrium). Moreover, the budget constraint, $c + p(d)e + q(d)b = w$ where $w = (d + p(d))e_{-1} + b_{-1}$ (e_{-1} and b_{-1} are equity and bond holdings chosen in the previous period), has been substituted away. The Euler equation for equity is

$$p(d)u'(d) =$$

$$\beta((\phi - v)[\lambda_1 d + p(\lambda_1 d)]u'(\lambda_1 d) + (1 - \phi + v)[\lambda_2 d + p(\lambda_2 d)]u'(\lambda_2 d)).$$

Clearly, p is linear in d (a constant times d), whenever $u'(c) = c^{-\alpha}$ (here, $\alpha = 1$). Since the period utility is logarithmic, the price of equity does not depend on beliefs because the payoff and the inverse of marginal utility (u') are proportional to λd so the payoff times marginal utility is the same in both states. Thus, $p(d) = \frac{\beta}{1-\beta}d$ solves the Euler equation above: the price of equity is independent of ϕ and v.

Trivially here, since $e = 1$ in equilibrium, $w_1' = \frac{\lambda_1 d}{1-\beta}$ and, $w_2' = \frac{\lambda_2 d}{1-\beta}$, then $V(w_1') > V(w_2')$, so the solution for v is a corner, i.e., $v = a$. We show in another paper (Alonso and Prado (2007)) that v can be an interior solution when the economy is populated by both ambiguity-averse and standard consumers.

The Euler equation for bonds similarly gives

$$q(d)u'(d) = \beta[(\phi - a)u'(\lambda_1 d) + (1 - \phi + a)u'(\lambda_2 d)].$$

We see that q depends on beliefs:

$$q = \beta\left[\left(\frac{\phi}{\lambda_1} + \frac{1-\phi}{\lambda_2}\right) + a\left(\frac{1}{\lambda_2} - \frac{1}{\lambda_1}\right)\right].$$

The higher is a – the more ambiguity aversion there is in the economy – the higher is the belief that the bad state will occur, and the higher is the present value of one unit tomorrow, since the probability weight placed on the state with a high marginal utility is higher.

The expected return on equity, ER_e, is given by

$$ER_e = \frac{\phi\lambda_1 + (1-\phi)\lambda_2}{\beta},$$

and it is independent of the belief. The return on bonds, R_b, decreases when ambiguity aversion increases, because $R_b = \frac{1}{q}$.

The equity premium in this economy is

$$ER_e - R_b = \frac{\phi\lambda_1 + (1-\phi)\lambda_2}{\beta} - \frac{\lambda_1\lambda_2}{\beta((1-\phi)\lambda_1 + \phi\lambda_2 + a(\lambda_1 - \lambda_2))}.$$

Since ϕ=0.5, the equity premium in this economy is

$$ER_e - R_b = \frac{\lambda_1 + \lambda_2}{2\beta} - \frac{\lambda_1\lambda_2}{\beta(0.5(\lambda_1 + \lambda_2) + a(\lambda_1 - \lambda_2))}.$$

When ambiguity is most extreme, i.e., when $a = 0.5$, the equity premium becomes

$$\frac{\lambda_1 - \lambda_2}{2\beta}.$$

Using $\lambda_1 = 1.02$, λ_2=1.01 and $\beta = 0.98$, the equity premium is 0.5%, and this is 200 times larger than the equity premium for the same parameter values when $a = 0$ – the standard model. Though this is an example, and not a calibration exercise, it illustrates that the effect of ambiguity on asset prices/returns can be substantial.

3.2 Serial correlation and CRRA utility function

We now assume that the period utility is $u(c) = \frac{c^{1-\alpha}-1}{1-\alpha}$ and that the shocks are serially correlated.

The problem of the representative agent with wealth today given by w and today's shock s is

$$V_s(w) = \max_e u(w - p_s(d)e) + \beta \min_{v_s \in [-a,a]} [(\phi_{s1} - v_s)V_1(w_1') + (\phi_{s2} + v_s)V_2(w_2')],$$

subject to

$$w_1' = [\lambda_1 d + p_1(d\lambda_1)]e,$$

$$w_2' = [\lambda_2 d + p_2(d\lambda_2)]e.$$

The Euler equation for equity is

$$p_s(d)u'(d) =$$

$$\beta\{(\phi_{s1} - v_s)[\lambda_1 d + p_1(\lambda_1 d)]u'(\lambda_1 d) + (\phi_{s2} + v_s)[\lambda_2 d + p_2(\lambda_2 d)]u'(\lambda_2 d)\}.$$

The price of equity is still linear in d, and is now given by

$$p_s(d) = k_s d$$

where

$$k_s = \beta[(\phi_{s1} - v_s)\lambda_1^{1-\alpha}(1 + k_1) + (\phi_{s2} + v_s)\lambda_2^{1-\alpha}(1 + k_2)],$$

for $s = 1, 2$.

Solving explicitly for k_1 and k_2, we obtain:

$$k_1 = \frac{\beta(\phi_{11} - v_1)\lambda_1^{1-\alpha} + \beta(\phi_{12} + v_1)\lambda_2^{1-\alpha} + \beta^2(\phi_{12} + v_1 - \phi_{22} - v_2)(\lambda_1\lambda_2)^{1-\alpha}}{1 - \beta(\phi_{22} + v_2)\lambda_2^{1-\alpha} - \beta(\phi_{11} - v_1)\lambda_1^{1-\alpha} - \beta^2(\phi_{12} + v_1 - \phi_{22} - v_2)(\lambda_1\lambda_2)^{1-\alpha}},$$

and

$$k_2 = \frac{\beta(\phi_{21} - v_2)\lambda_1^{1-\alpha} + \beta(\phi_{22} + v_2)\lambda_2^{1-\alpha} + \beta^2(\phi_{12} + v_1 - \phi_{22} - v_2)(\lambda_1\lambda_2)^{1-\alpha}}{1 - \beta(\phi_{22} + v_2)\lambda_2^{1-\alpha} - \beta(\phi_{11} - v_1)\lambda_1^{1-\alpha} - \beta^2(\phi_{12} + v_1 - \phi_{22} - v_2)(\lambda_1\lambda_2)^{1-\alpha}},$$

Thus, wealth next period is:

$$w_1' = \lambda_1 d(1 + k_1),$$

and

$$w_2' = \lambda_2 d(1 + k_2).$$

The solution for v_s is still a corner (i.e., $v_s = a$ for $s = 1, 2$). To show this, one needs to verify that $V_1(w_1') > V_2(w_2')$. Since $V_s(w_s') = A_s(w_s')^{1-\alpha}$ (where A_s is defined in the next section), one needs to show that $A_1[\lambda_1 d(1 + k_1)]^{1-\alpha} > A_2[\lambda_2 d(1 + k_2)]^{1-\alpha}$, which is straightforward.

The price of the bond is given by

$$q_s(d) = \beta\left[\phi_{s1}\frac{1}{\lambda_1^\alpha} + \phi_{s2}\frac{1}{\lambda_2^\alpha} + a\left(\frac{1}{\lambda_2^\alpha} - \frac{1}{\lambda_1^\alpha}\right)\right]$$

for $s = 1, 2$.

The conditional expected return on equity is

$$ER_s^e = \frac{\phi_{s1}[\lambda_1 d + p_1(\lambda_1 d)] + \phi_{s2}[\lambda_2 d + p_2(\lambda_2 d)]}{p_s(d)}$$

for $s = 1, 2$, and the unconditional expected return on equity, ER^e, is

$$\pi ER_1^e + (1 - \pi) ER_2^e,$$

where the invariant probability π solves

$$\pi = \phi_{11}\pi + \phi_{21}(1 - \pi).$$

Therefore,

$$ER^e = \pi \frac{\phi_{11}[\lambda_1 d + p_1(\lambda_1 d)] + \phi_{12}[\lambda_2 d + p_2(\lambda_2 d)]}{p_1(d)} + (1-\pi) \frac{\phi_{21}[\lambda_1 d + p_1(\lambda_1 d)] + \phi_{22}[\lambda_2 d + p_2(\lambda_2 d)]}{p_2(d)} =$$

$$\pi \frac{\phi_{11}\lambda_1(1 + k_1) + \phi_{12}\lambda_2(1 + k_2)}{k_1} + (1 - \pi) \frac{\phi_{21}\lambda_1(1 + k_1) + \phi_{22}\lambda_2(1 + k_2)}{k_2}.$$

The expected return on the bond, R^b, is given by

$$\pi \frac{1}{q_1} + (1 - \pi) \frac{1}{q_2} =$$

$$\frac{1}{\beta} \left(\frac{\pi}{\phi_{11} \frac{1}{\lambda_1^\alpha} + \phi_{12} \frac{1}{\lambda_2^\alpha} + a\left(\frac{1}{\lambda_2^\alpha} - \frac{1}{\lambda_1^\alpha}\right)} + \frac{1 - \pi}{\phi_{21} \frac{1}{\lambda_1^\alpha} + \phi_{22} \frac{1}{\lambda_2^\alpha} + a\left(\frac{1}{\lambda_2^\alpha} - \frac{1}{\lambda_1^\alpha}\right)} \right).$$

Finally, the equity premium is given by

$$ER^e - R^b.$$

3.3 Calibration and evaluation of asset prices

As in Mehra and Prescott (1985), we now select the parameters of the model so that the average growth rate of per capita consumption, the standard deviation of the growth rate of per capita consumption, and the first-order serial correlation of this growth rate, all with respect to the model's stationary distribution, match the sample values for the U.S. economy between 1889-1978.

The values of the parameters are $\phi = 0.43$ (where $\phi_{11} = \phi_{22} = \phi$ and $\phi_{12} = \phi_{21} = (1 - \phi)$), $\lambda_1 = 1.054$, and $\lambda_2 = 0.982$.

Figure 1 shows the return on the riskfree asset, the expected return on equity, and the equity premium for $\beta = 0.95$, $\alpha = 2$, and the ambiguity parameter a in a range between 0 and 0.43.

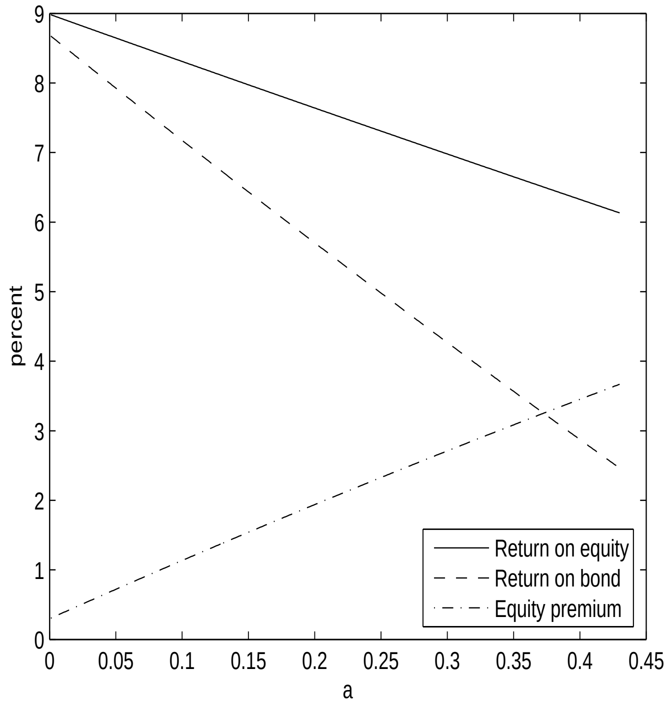

Figure 1: The role of ambiguity (a) for asset returns

The equity premium increases with the amount of ambiguity in the economy. When a is higher, the consumer puts higher probability on the bad state: the future looks worse so returns on assets are lower. For example, for $a = 0.3$, the return on the bond is 4.27%, the expected return on equity is 6.98%, and the resulting equity premium is 2.71%. As a comparison, the largest equity premium that Mehra and Prescott were able to obtain was 0.35%.

Figure 2 shows the return on the riskfree bond, the expected return on equity and the equity premium for $\beta = 0.95$, $a = 0.2$, and a range of α between 0 and 10.

The equity premium is monotonically higher as α increases. If consumers are more risk-averse, they require a higher premium to hold equity. Note, for example, that for $\alpha = 8$, the riskfree return is 4.72%, the expected return on equity is 8.77%, and the resulting equity premium is 4.05%.

The returns on bonds and equity increase in α up to a point, and beyond that point, the returns decrease. If the economy grows over time, as is the case on average here, a higher α implies that consumption today is increasingly more valuable than consumption in the future, since the marginal utility is then decreasing rapidly over

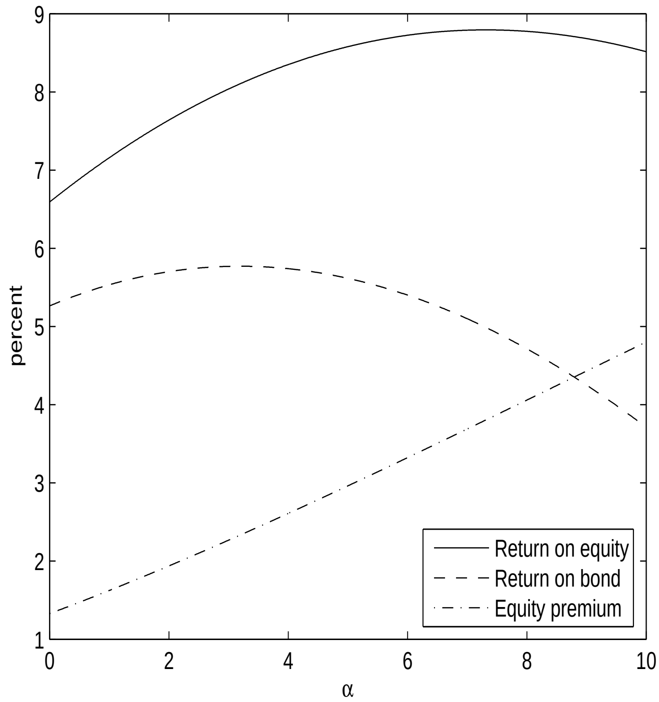

Figure 2: The role of risk aversion (α) for asset returns

time. As a result, the interest rate increases in α. In the two-state economy here, resources grow in state 1 ($\lambda_1 > 1$) but shrink in state 2 ($\lambda_2 < 1$). For lower values of α the effect of positive growth in state 1 dominates the effect of a negative growth in state 2, since it is more likely, and asset returns increase in α. However, passed a threshold level, the effect of negative growth in state 2 dominates over the positive effect in state 1 and, as a result, asset returns start decreasing on α. In other words, when α is high enough, agents care more about the downside risk of the dividend growth rate (λ_2) than about the average dividend growth rate.

3.4 Bond-return volatility, long-run bond returns, and ambiguity aversion

As indicated above, there are many combinations of the risk aversion, subjective discount factor, and ambiguity aversion parameters which match the equity premium and the riskfree rate. We could therefore potentially use other targets, e.g., the volatility of the short-term bond return or the average return on a long-term bond, to select the ambiguity aversion parameter, a.

Using the bond prices derived in subsection 3.2, the volatility of the bond return,

σ_b, can be calculated as

$$\sigma_b = \pi\left(\frac{1-q_1}{q_1}\right)^2 + (1-\pi)\left(\frac{1-q_2}{q_2}\right)^2 - \left(\pi\frac{1-q_1}{q_1} + (1-\pi)\frac{1-q_2}{q_2}\right)^2.$$

To derive the price of the long-term bond return, also called a console (a bond which pays a fixed coupon indefinitely), let us use the prices of contingent claims giving 1 unit in state $j = 1, 2$ tomorrow if the state today is $s = 1, 2$. They are given by

$$p_{s1} = \beta(\phi_{s1} - v_1)\lambda_1^{-\alpha}$$

and

$$p_{s2} = \beta(\phi_{s2} + v_2)\lambda_2^{-\alpha}.$$

Then the prices of the console paying a coupon of x consumption units every period forever, given that the current state is $s = 1, 2$, are

$$q_{\infty,1} = p_{11}(x + q_{\infty,1}) + p_{12}(x + q_{\infty,2})$$

and

$$q_{\infty,2} = p_{21}(x + q_{\infty,1}) + p_{22}(x + q_{\infty,2}).$$

Solving these two equations, we obtain the prices of the infinitely-long bonds explicitly. They are

$$q_{\infty,1} = \frac{[p_{12}(1+p_{21}) + p_{11}(1-p_{22})]x}{(1-p_{22})(1-p_{11}) - p_{12}p_{21}}$$

if today's state is 1 and

$$q_{\infty,2} = \frac{[p_{21}(1+p_{12}) + p_{22}(1-p_{11})]x}{(1-p_{22})(1-p_{11}) - p_{12}p_{21}}$$

if today's state is 2.

Then the two conditional expected returns become

$$R^b_{\infty,1} = \frac{(q_{\infty,1}+x)\phi_{11} + (q_{\infty,2}+x)\phi_{12}}{q_{\infty,1}}$$
$$= \frac{1 - p_{11} - p_{22} + (p_{12}+p_{11})\phi_{11} + (p_{21}+p_{22})\phi_{12}}{p_{11} + p_{12} + p_{12}p_{21} - p_{11}p_{22}}$$

and

$$R^b_{\infty,2} = \frac{(q_{\infty,1}+x)\phi_{21} + (q_{\infty,2}+x)\phi_{22}}{q_{\infty,2}}$$
$$= \frac{1 - p_{11} - p_{22} + (p_{12}+p_{11})\phi_{21} + (p_{21}+p_{22})\phi_{22}}{p_{21} + p_{22} + p_{12}p_{21} - p_{11}p_{22}}.$$

The implied unconditional return is given by

$$R^b_\infty = \pi E R^b_{\infty,1} + (1-\pi) E R^b_{\infty,2},$$

where π is the invariant probability.

In the data, the volatility of the real 3-month U.S. Treasury Bond return over the period Jan-1948 to Oct-2007 is 2.3%.[5] The average real return on the long-term bond (30-year Treasury bond) over the period Jan-1977 to Oct-2007 is 3.6%. In Table 1, we display the volatility of the short-term bond return and the average return on the long-term bond for different values of α, β, and a. These values all deliver the same equity premium and the same return on the short-run riskless bond: 6.18% and 0.8%, respectively.

β	α	a	σ_b	R^b_∞
1.124	18.351	0.000	8.38%	4.97%
1.110	18.065	0.012	8.18%	4.93%
1.090	17.640	0.029	7.89%	4.88%
1.070	17.194	0.048	7.60%	4.82%
1.050	16.722	0.069	7.30%	4.77%
1.030	16.221	0.091	6.99%	4.71%
1.010	15.683	0.115	6.68%	4.64%
0.990	15.101	0.141	6.35%	4.57%
0.970	14.464	0.171	6.01%	4.49%
0.950	13.752	0.204	5.65%	4.39%
0.930	12.938	0.242	5.26%	4.28%
0.910	11.965	0.288	4.82%	4.15%
0.890	10.701	0.348	4.29%	3.97%
0.873	8.955	0.430	3.59%	3.68%

Table 1: Volatilities and returns

As we can see in the table, the data favors high levels of ambiguity. For $a = 0.43$, the model's long-term real return is 3.68% and the data is 3.6%. Our volatility for the short-term riskfree real return is still high, but since the volatility is decreasing in a, the data suggests a high value for ambiguity (with lower values for the subjective discount factor and the risk aversion). This reassures us that ambiguity aversion plays a role in determining asset prices. However, in the next section, where we study the welfare costs of consumption fluctuations, we will consider the whole range of values for a from 0 to 0.43.

[5] The data on returns and consumer price index are from the Federal Reserve Economic Data, available online at the Federal Reserve Bank of St. Louis's web page: http://research.stlouisfed.org/fred2/.

4 Eliminating consumption fluctuations

As in the previous section, we begin with the simple example and then proceed to the case we use for the quantitative analysis.

4.1 A simple example – *iid* shocks

We first calculate the costs of consumption fluctuations when shocks are *iid*. The present discounted utility when the representative agent consumes the stream of dividends given by the tree if the dividend today is d is given recursively by

$$V(d) = \frac{d^{1-\alpha}}{1-\alpha} + \beta \left\{ \min_{v \in [-a,a]} [(\phi - v)V(\lambda_1 d) + (1 - \phi + v)V(\lambda_2 d)] \right\}.$$

The solution for $V(d)$ is

$$V(d) = Ad^{1-\alpha},$$

where

$$A = \frac{1}{(1-\alpha)\{1 - \beta[(\phi - a)\lambda_1^{1-\alpha} + (1 - \phi + a)\lambda_2^{1-\alpha}]\}}.$$

As shown in the previous section, $v = a$ since $V(d)$ is increasing in d.

Eliminating consumption fluctuations will deliver the present value of total utility corresponding to consuming the expected value of the dividend every period. This utility is given by:

$$\sum_{t=0}^{\infty} \beta^t \frac{[d(\phi\lambda_1 + (1-\phi)\lambda_2)^t]^{1-\alpha}}{1-\alpha} = \frac{d^{1-\alpha}}{(1-\alpha)[1 - \beta(\phi\lambda_1 + (1-\phi)\lambda_2)^{1-\alpha}]}.$$

The costs of consumption variability are given by γ where γ solves:

$$\frac{(1-\gamma)^{1-\alpha}}{1 - \beta[\phi\lambda_1 + (1-\phi)\lambda_2]^{1-\alpha}} = \frac{1}{1 - \beta[(\phi - a)\lambda_1^{1-\alpha} + (1 - \phi + a)\lambda_2^{1-\alpha}]}. \tag{1}$$

Figure 3 below plots γ as a function of a for $\alpha = 2$ and $\beta = 0.95$. The figure shows how the costs of consumption variability are increasing in a.

Turning to the roles of risk aversion and the discount factor, Figure 4 below looks at the special case $a = 0$ (no ambiguity aversion) and plots γ as a function of α, for several different values of β.

As β increases the curve moves up, since the welfare gains are higher the more the consumer cares about future consumption. As expected, as α increases so do the welfare costs of business fluctuations.

Figure 5 below looks at the case with ambiguity: it plots the same function but for $a = 0.1$. When there is ambiguity and for a β that is sufficiently high, the real costs of

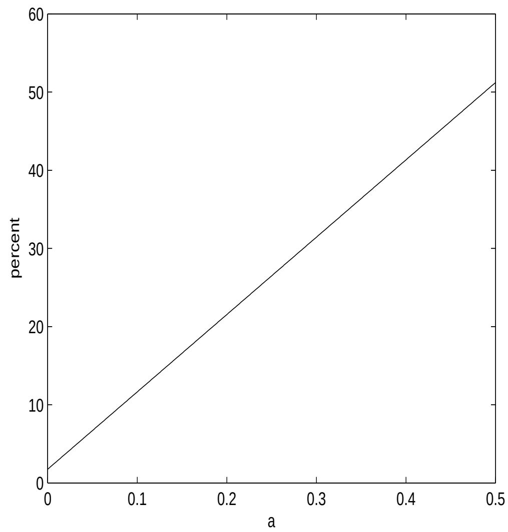

Figure 3: The role of ambiguity (a) for the costs of fluctuations

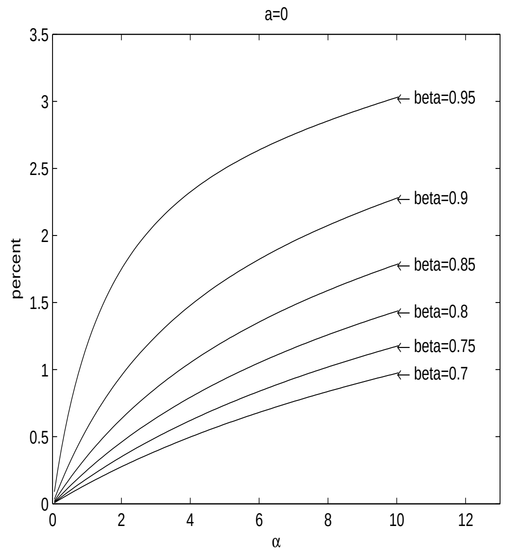

Figure 4: The role of risk aversion (α) for the costs of fluctuations without ambiguity

business fluctuations, strikingly enough, *decrease* as α increases.

What explains the finding that in an economy with ambiguity, with more risk aversion, the costs of fluctuations may decrease? Looking more in detail into the costs of consumption fluctuations, one realizes that there are two kinds of gain from removing these fluctuations: (1) the gain from moving to non-fluctuating consumption and (2) the gain from not being pessimistic about the good state.

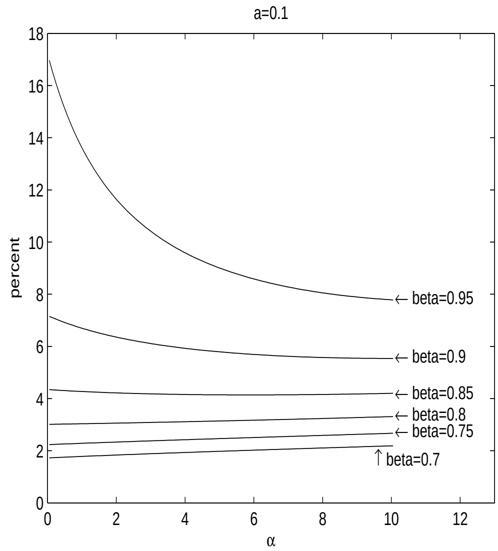

Figure 5: The role of risk aversion (α) for the costs of fluctuations with ambiguity

To analyze this distinction we also construct $\hat{\gamma}$ as a function of α, where $\hat{\gamma}$ isolates the gains from moving to non-fluctuating consumption; i.e., $\hat{\gamma}$ solves

$$\frac{(1-\hat{\gamma})^{1-\alpha}}{1-\beta[(\phi-a)\lambda_1+(1-\phi+a)\lambda_2]^{1-\alpha}} = \frac{1}{1-\beta[(\phi-a)\lambda_1^{1-\alpha}+(1-\phi+a)\lambda_2^{1-\alpha}]}. \quad (2)$$

On the left-hand side of this equation, we have expected consumption growth using pessimistic beliefs, as opposed to neutral beliefs ($a = 0$), thus eliminating the second source of gains from removing business cycles in this economy.[6]

Figure 6 below plots γ and $\hat{\gamma}$ for $\beta = .65$ and $\beta = .95$, respectively.

The difference between the solid line and the dotted line is the gain from removing pessimism. It is decreasing in γ. The reason why the removal of pessimism gives a lower

[6] To prove that the gain can be decreasing in γ, consider the derivative of γ with respect to α at $\alpha = 0$. It is given by

$$\frac{d\gamma}{d\alpha} = -\left(\frac{1-\beta\bar{\lambda}}{1-\beta\bar{\lambda}_a}\right)\left[\ln\left(\frac{1-\beta\bar{\lambda}}{1-\beta\bar{\lambda}_a}\right)+\frac{\beta\bar{\lambda}\ln(\bar{\lambda})}{1-\beta\bar{\lambda}}-\frac{\beta[(\phi-a)\lambda_1\ln(\lambda_1)+(1-\phi+a)\lambda_2\ln(\lambda_2)]}{1-\beta\bar{\lambda}_a}\right],$$

where $\bar{\lambda} = \phi\lambda_1 + (1-\phi)\lambda_2$ and $\bar{\lambda}_a = (\phi-a)\lambda_1 + (1-\phi+a)\lambda_2$.

If $a = 0$, $\bar{\lambda} = \bar{\lambda}_a$ and the sign of this derivative is given by the difference between $\lambda_1 \ln \lambda_1 + \lambda_2 \ln \lambda_2$ and $\lambda_1 \ln \bar{\lambda} + \lambda_2 \ln \ln \bar{\lambda}$. This sign is positive since $\lambda_1 > \lambda_2$ and it does not depend on the value of β.

However, if $a > 0$, then $\bar{\lambda} > \bar{\lambda}_a$ and the sign of the derivative depends on the value of β. For a high value of β, the second term in the sum above is high relative to the first and third terms of the sum, implying that the sign of the derivative is negative. But for a lower value of β, the effect of β on the denominators of the three terms of the sum decreases the value of the sum, and for a β sufficiently low, the sign of the derivative is reversed and becomes positive.

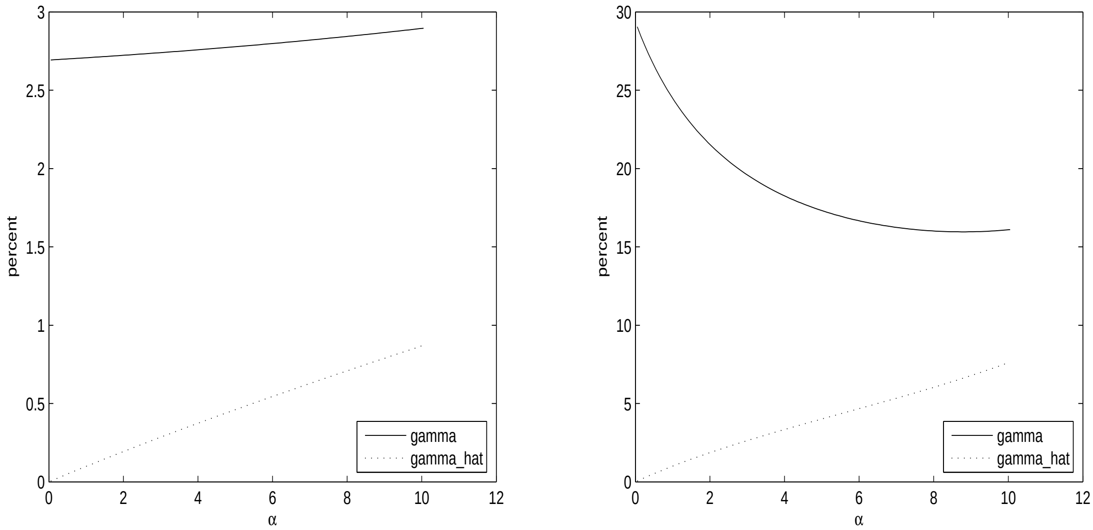

Figure 6: The costs of removing fluctuations and the costs of removing pessimism

welfare increase for lower values of α is that when pessimism is removed, probability mass is in effect moved from the bad state to the good state. What is the utility gain from such an experiment? It depends on the curvature of the utility function: the higher is the curvature, the lower is marginal utility for higher consumption values, so the smaller is the effect. That is, the utility gain from consuming more rather than less, and thus from moving probability mass from the low to the high state, decreases in α. The maximum effect is clearly obtained when $u(c)$ is linear.

In sum, the effect of removing consumption fluctuations is increasing in α while the effect of removing pessimism is decreasing in α. Thus, as α increases, which of these two effects dominates is not a priori clear. The figures above reveal that the answer can depend on the value of β. The reason why β can determine the relative strengths between these two effects – those of removing pessimism and fluctuations – appears to be related to the fact that although both appear in the future, the removal of pessimism relies on increasing average *growth rates*, as opposed to levels. When growth rates are affected, the effect builds up over time, and discounting plays a more important role.

4.2 Serially correlated shocks and a CRRA utility function

Calculating the utility of the deterministic growth path when the shocks are serially correlated is more involving. To this end, we will now introduce some notation.

Let the transition probabilities be given by

$$\Phi \equiv \begin{pmatrix} \phi_{11} & \phi_{12} \\ \phi_{21} & \phi_{22} \end{pmatrix},$$

let

$$\Lambda \equiv \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix},$$

and let

$$\lambda \equiv \begin{pmatrix} \lambda_1 \\ \lambda_2 \end{pmatrix}.$$

Consider the expression

$$\lambda_t^e \equiv (\Phi\Lambda)^{t-1}\,\Phi\lambda. \tag{3}$$

The first row of this expression is the *expected growth rate between now and t periods from now if the state now is state 1* and the second row is the *expected growth rate between now and t periods from now if the state now is state 2*. Denote these $\lambda_t^e|_1$ and $\lambda_t^e|_2$, respectively. That is,

$$\lambda_t^e \equiv \begin{pmatrix} \lambda_t^e|_1 \\ \lambda_t^e|_2 \end{pmatrix}.$$

The utility of the deterministic growth path, where growth is equal to the expected value beginning in state s, is

$$u(d) + \beta u(d\lambda_1^e|_s) + \beta^2 u(d\lambda_2^e|_s) + \beta^3 u(d\lambda_3^e|_s) + \dots,$$

which when we have CRRA utility equals

$$\frac{d^{1-\alpha}}{1-\alpha}\left(1 + \beta\left(\lambda_1^e|_s\right)^{1-\alpha} + \beta^2\left(\lambda_2^e|_s\right)^{1-\alpha} + \beta^3\left(\lambda_3^e|_s\right)^{1-\alpha} + \dots\right). \tag{4}$$

The rate at which λ_t^e grows in the limit is given by the greatest (in absolute value) eigenvalue of $\Phi\Lambda$[7].

The present value of total utility when the dividend is d and the shock is s, is given by

$$V_s(d) = A_s d^{1-\alpha},$$

for $s = 1, 2$, and where

$$A_1 = \frac{1 + \beta\lambda_2^{1-\alpha}(\phi_{12} - \phi_{22})}{(1-\alpha)[1 - \beta(\phi_{22} + v_2)\lambda_2^{1-\alpha} - \beta(\phi_{11} - v_1)\lambda_1^{1-\alpha} - \beta^2(\phi_{12} + v_1 - \phi_{22} - v_2)(\lambda_1\lambda_2)^{1-\alpha}]},$$

[7]Therefore, the sum in (4) exists if this eigenvalue raised to the power of $1-\alpha$ times β is less than 1.

and

$$A_2 = \frac{1+\beta\lambda_1^{1-\alpha}(\phi_{21}-\phi_{11})}{(1-\alpha)[1-\beta(\phi_{22}+v_2)\lambda_2^{1-\alpha}-\beta(\phi_{11}-v_1)\lambda_1^{1-\alpha}-\beta^2(\phi_{12}+v_1-\phi_{22}-v_2)(\lambda_1\lambda_2)^{1-\alpha}]}.$$

Since $v_1 = v_2 = a$,

$$A_1 = \frac{1+\beta\lambda_2^{1-\alpha}(\phi_{12}-\phi_{22})}{(1-\alpha)[1-\beta(\phi_{22}+a)\lambda_2^{1-\alpha}-\beta(\phi_{11}-a)\lambda_1^{1-\alpha}-\beta^2(\phi_{12}-\phi_{22})(\lambda_1\lambda_2)^{1-\alpha}]},$$

and

$$A_2 = \frac{1+\beta\lambda_1^{1-\alpha}(\phi_{21}-\phi_{11})}{(1-\alpha)[1-\beta(\phi_{22}+a)\lambda_2^{1-\alpha}-\beta(\phi_{11}-a)\lambda_1^{1-\alpha}-\beta^2(\phi_{12}-\phi_{22})(\lambda_1\lambda_2)^{1-\alpha}]}.$$

Thus, the welfare cost starting from state 1 is given by the γ_1 solving

$$A_1 = \frac{(1-\gamma_1)^{1-\alpha}}{1-\alpha}\left(1+\beta\left(\lambda_1^e|_1\right)^{1-\alpha}+\beta^2\left(\lambda_2^e|_1\right)^{1-\alpha}+\beta^3\left(\lambda_3^e|_1\right)^{1-\alpha}+\dots\right). \qquad (5)$$

Similarly, the welfare cost starting from state 2 is given by the γ_2 solving

$$A_2 = \frac{(1-\gamma_2)^{1-\alpha}}{1-\alpha}\left(1+\beta\left(\lambda_1^e|_2\right)^{1-\alpha}+\beta^2\left(\lambda_2^e|_2\right)^{1-\alpha}+\beta^3\left(\lambda_3^e|_2\right)^{1-\alpha}+\dots\right). \qquad (6)$$

To solve this numerically, it is helpful to define

$$\hat{\lambda}_t^e \equiv \left(\beta^{\frac{1}{1-\alpha}}\Phi\Lambda\right)^{t-1}\Phi\lambda\beta^{\frac{1}{1-\alpha}}$$

and then solve the implied new equations corresponding to the previous two:

$$A_1 = \frac{(1-\gamma_1)^{1-\alpha}}{1-\alpha}\left(1+\left(\hat{\lambda}_1^e|_1\right)^{1-\alpha}+\left(\hat{\lambda}_2^e|_1\right)^{1-\alpha}+\left(\hat{\lambda}_3^e|_1\right)^{1-\alpha}+\dots\right) \qquad (7)$$

and

$$A_2 = \frac{(1-\gamma_2)^{1-\alpha}}{1-\alpha}\left(1+\left(\hat{\lambda}_1^e|_2\right)^{1-\alpha}+\left(\hat{\lambda}_2^e|_2\right)^{1-\alpha}+\left(\hat{\lambda}_3^e|_2\right)^{1-\alpha}+\dots\right) \qquad (8)$$

for γ_1 and γ_2, respectively. The latter sums are easier to compute whenever they are finite, since then they consist of terms that go to zero as t goes to infinity. (The object to be raised to the power $1-\alpha$ may go to infinity when $\alpha > 1$, but raising to the power then implies that the terms go to zero.)

Figure 7 below plots the costs of business cycles for $\beta = 0.9$ and $\alpha = 2$ as a function of a; i.e., it shows a "comparative-statics" exercise with respect to the ambiguity parameter only.

Clearly, also in the economy with a realistic shock process, more ambiguity aversion increases the costs of business cycles. By eliminating fluctuations (if that is possible),

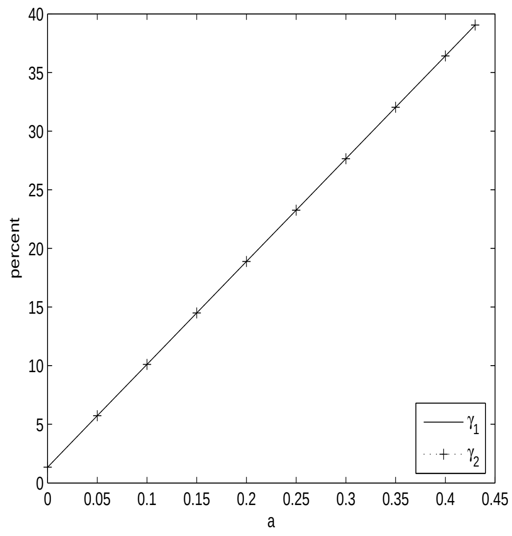

Figure 7: The quantitative costs of removing fluctuations as a function of a

the government would eliminate the first-order negative effect on utility that consumers experience from random consumption.

We continue with comparative statics with respect to various parameters and then finally describe the welfare costs when the parameters are selected to match the asset prices.

Figure 8 below plots the costs of business cycles for $\beta = 0.95$, and $a = 0.1$ as a function of α. As argued in the previous subsection, for high values of β (such as 0.95), consumption fluctuations hurt less in total the more risk-averse is the consumer, since the removal of pessimism is less beneficial for high levels of risk aversion, an effect which relies strongly on low discounting.

Finally, Figure 9 plots the costs of business cycles for $\alpha = 2$ and $a = 0.2$ as a function of β. Consumers who put a higher weight on future consumption clearly suffer more from consumption fluctuations.

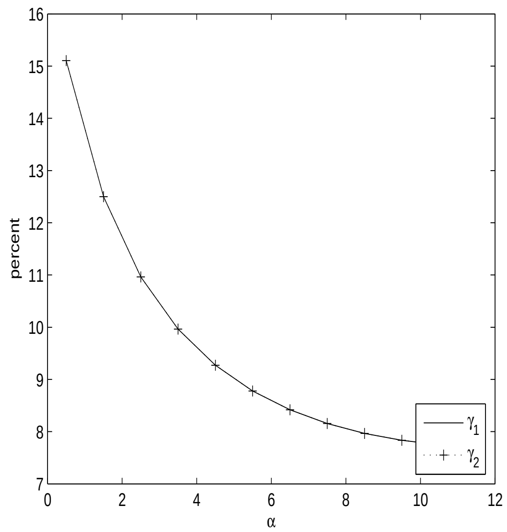

Figure 8: The quantitative costs of removing fluctuations as a function of α

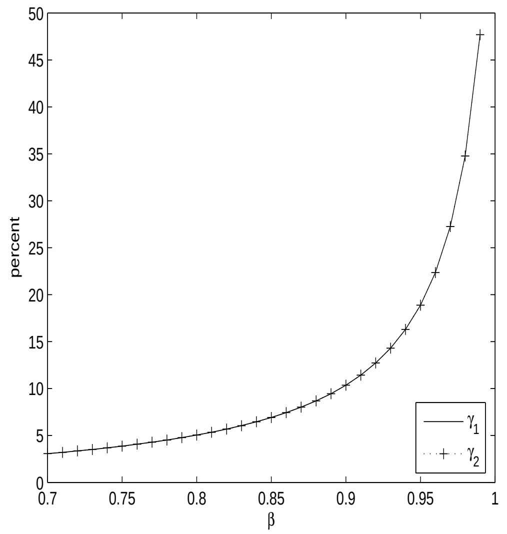

Figure 9: The quantitative costs of removing fluctuations as a function of β

4.3 The costs of fluctuations for calibrations that match asset prices

We now look at the costs of fluctuations when the asset prices match the data for the period considered by Mehra and Prescott, i.e., 1889-1978. Mehra and Prescott found that the average real return on short-term riskless securities was 0.8%, and the average annual real return on the Standard and Poor's 500 Composite Stock Index was 6.98%. This implies that the average equity premium over that period was 6.18%.

As was discussed briefly above, matching these data can be accomplished in different ways, and each of these is associated with a different cost. Table 2 below illustrates that the welfare costs – or, rather, the *potential* welfare costs – of cycles are huge.

β	α	a	R^b	$ER^e - R^b$	γ_1	γ_2
1.124	18.350	0.000	0.8%	6.18%	8.03%	7.99%
1.120	18.270	0.003	0.8%	6.18%	8.11%	8.07%
1.100	17.855	0.020	0.8%	6.18%	8.51%	8.47%
1.070	17.194	0.048	0.8%	6.18%	9.13%	9.10%
1.020	15.956	0.102	0.8%	6.18%	10.29%	10.26%
1.000	15.398	0.127	0.8%	6.18%	10.82%	10.80%
0.950	13.752	0.204	0.8%	6.18%	12.46%	12.44%
0.930	12.938	0.242	0.8%	6.18%	13.33%	13.31%
0.910	11.965	0.288	0.8%	6.18%	14.45%	14.43%
0.890	10.701	0.348	0.8%	6.18%	16.06%	16.05%
0.873	8.954	0.430	0.8%	6.18%	18.72%	18.72%

Table 2: Welfare costs for preferences that match asset prices

The costs (recall that γ_z denotes the cost in state z) differ markedly across the different parameter configurations. Notice that the highest costs appear for higher values of a, i.e., when there is more ambiguity. Looking at the first two columns of the table, we also see that higher values of a are attractive in that they make very high values of risk aversion, α less crucial, and they allow more discounting. Without ambiguity, discounting needs to be negative – a β *higher* than 1 is required – in order to explain the low riskfree rate: to explain a high risk premium, a high α is needed, but a high α makes a consumption unit in the future worth much less than now, given that consumption is increasing on average, thus leading to a high real interest rate, unless β is made very high. A β above unity is consistent with aggregate data, which shows consumption growth during most episodes, but simple introspection (given a consumption path, would you choose a unit of consumption now or a unit later?) does cast some doubt on this possibility.

Lastly, for comparison we show the associated costs when $a = 0$ for the same configurations of α and β; see Table 3. These are high too compared to Lucas's numbers, since α is high (above 10 in all except the last row), but much lower than those above.

β	α	a	R^b	$ER^e - R^b$	γ_1	γ_2
1.124	18.351	0	0.80%	6.18%	8.03%	8.00%
1.120	18.270	0	1.18%	6.15%	7.71%	7.68%
1.100	17.855	0	3.11%	6.00%	6.46%	6.42%
1.070	17.194	0	6.13%	5.75%	5.16%	5.13%
1.020	15.956	0	11.45%	5.28%	3.77%	3.74%
1.000	15.398	0	13.69%	5.06%	3.36%	3.34%
0.950	13.752	0	19.49%	4.40%	2.55%	2.53%
0.930	12.938	0	21.84%	4.07%	2.27%	2.26%
0.910	11.965	0	24.15%	3.67%	2.02%	2.00%
0.890	10.701	0	26.27%	3.15%	1.76%	1.75%
0.873	8.955	0	27.46%	2.45%	1.50%	1.49%

Table 3: Welfare costs for preferences without ambiguity

5 Conclusion

In this paper we have studied asset pricing and evaluated the welfare costs of fluctuations in consumption for an economy where consumers are ambiguity-averse. First, we have shown parameter configurations under which the equity premium is quite large (and the riskfree rate is small); the ability to match these return features comes from the ability of ambiguity aversion to generate *first-order* effects on prices. This sets it apart from risk aversion, which operates through second-order effects. Ambiguity aversion has first-order effects, in essence, because consumers behave as if they have beliefs that the good return outcomes are less likely than they really are.

Second, using the calibrations that deliver realistic asset prices, we have shown that the welfare benefits of eliminating consumption fluctuations need not be as small as those in Lucas's calculations. This is not to say that the benefits are large: the numbers we obtain are, just like Lucas's numbers, upper bounds, and these upper bounds leave open what the costs of stabilization (say, in the form of distortions) might be, and also leave open whether full stabilization even is feasible. Nevertheless, it is valuable to note that these bounds can be as large as 15% in permanent consumption terms when asset prices are matched by the model.

Several questions may arise in response to our hypothesis that there is ambiguity and, in essence, pessimism toward consumption and stock-price fluctuations. For example, wouldn't consumers realize that it is "irrational" to be pessimistic, since predictions based on their beliefs can be found to be systematically off? First, the present model does not, in a strict sense, fall under the behavioral rubric: ambiguity aversion is not only documented in experiments but also rationalized axiomatically based on the experimental evidence, and this axiomatization is employed in the present paper. Second, one may however wonder if ambiguity would not disappear if consumers were able to learn. Marinacci (2002) and Epstein and Schneider (2007, 2008) consider learning and conclude that indeed it is possible that consumers learn and that the pessimism disappears over time. The latter paper, however, also considers learning where "new ambiguity" may appear over time, and in this case complete learning will not occur in general. New ambiguity would be captured by new aggregate events occurring over time, such as technical change (an "IT era" being a particular striking example), financial crises, and political crises, and if such events trigger large consumption fluctuations it would seem reasonable to assume that ambiguity about aggregate events is quantitatively important and a recurring, if not permanent, phenomenon. The present paper is a stark example, perhaps, but could be extended to a setting like that considered by Epstein and Schneider, and we conjecture that ambiguity would still have important effects on the welfare costs of fluctuations in such an economy.

Another question regards heterogeneity: what if some consumers are ambiguity-averse while others are not? Might the ambiguity-averse people tend to lose money on average, given their beliefs, so that those agents with significant wealth in the long run would mainly consist of "standard" consumers? We explore this question in another paper – Alonso and Prado (2007) – and indeed find that ambiguity-averse consumers, by making consistently "bad bets", will see their relative wealth decline over time, and thereby asset prices will be more and more dominated by standard consumers. Note also that these bad bets are not bad in the sense of "crazy portfolios", but simply in the sense of delivering lower return on average by not investing enough in stock. In particular, if ambiguity aversion is large enough, the ambiguity-averse consumers choose to not participate at all in the stock market: the other, standard consumers, hold all the risk (and get all the high returns on average). In order to make the wealth distribution not converge to an extreme outcome, one could consider an overlapping-generations structure, where in each generation of newborns with equal asset wealth some are ambiguity-averse; that way, a significant part of aggregate wealth will always belong to ambiguity-averse consumers, and such a setting would seem quite realistic.

The most important remaining question seems to be about finding alternative ways of identifying the degree of ambiguity aversion, as represented by a in the analysis above. Experiments could perhaps be used but one would seem to need to examine ambiguity event by event. The approach here is to use asset prices to inform us on the value of a; we thus identify combinations of values for the discount factor and for risk aversion that allow us to match the equity premium and the riskfree rate. One approach would be to go further by looking at more aggregate asset-price moments; another would be to try to estimate a by also looking at individual portfolio data.

References

[1] Abel, B. Andrew (2002): “An Exploration of the Effects of Pessimism and Doubt on Asset Returns”, *Journal of Economic Dynamics & Control* 26, 1075-1092.

[2] Aiyagari, S. Rao (1993): “Explaining Financial Market Facts: The Importance of Incomplete Markets and Transaction Costs”, *Federal Reserve Bank of Minneapolis Quartely Review* 17(1), 17-31.

[3] Alonso, Irasema (2007): “Ambiguity in a Two-Country World”, mimeo.

[4] Alonso, Irasema and Jose Mauricio Prado (2007): “Heterogeneity in Ambiguity Aversion, Portfolio Choice, and the Wealth Distribution”, mimeo.

[5] Alvarez, Fernando and Urban J. Jermann (2004): “Using Asset Prices to Measure the Cost of Business Cycles”, *Journal of Political Economy* 112, 1223-1256.

[6] Bansal, Ravi, and Amir Yaron (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles”, *The Journal of Finance* LIX(4), 1481-1509.

[7] Croce M. Mariano (2007): “Welfare Costs and Lon-Run Consumption Risk in a Production Economy”, mimeo.

[8] DeJong, David N. and Marla Ripoll (2007): “Do Self-Control Preferences Help Explain the Puzzling Behavior of Asset Prices?”, *Journal of Monetary Economics* 54(4), 1035-1050.

[9] Epstein, Larry G. and Stanley E. Zin (1989): “Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework”, *Econometrica* 57, 937-969.

[10] Epstein, Larry G. and Martin Schneider (2003): “Recursive Multiple-Priors”, *Journal of Economic Theory* 113, 1-31.

[11] Epstein, Larry G. and Martin Schneider (2007): “Learning under Ambiguity”, *Review of Economic Studies* 74, 1275-1303.

[12] Epstein, Larry G. and Martin Schneider (2008): “Ambiguity, Information Quality and Asset Pricing”, *Journal of Finance* 63, 197-228.

[13] Epstein, Larry G. and Tan Wang (1994): “Intertemporal Asset Pricing under Knightian Uncertainty,” *Econometrica* 62, 283-322.

[14] Gilboa, Itzhak and David Schmeidler (1989): “Subjective Probability and Expected Utility without Additivity”, *Econometrica* 57, 571-587.

[15] Krusell, Per and Anthony A. Smith Jr. (1999): “On the Welfare Effects of Eliminating Business Cycles”, *Review of Economic Dynamics* 2(2), 245-272.

[16] Kocherlakota, Narayana R. (1996): "The Equity Premium: It's Still A Puzzle", *Journal of Economic Literature* 34, 42-71.

[17] Lucas, Robert E. Jr. (1978): "Asset Prices in an exchange economy", *Econometrica* 46, 1429-1445.

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

[19] Lucas, Robert E. Jr. (2003): "Macroeconomic Priorities", *American Economic Review* 93, 1-14.

[20] Marinacci, Massimo (2002): "Learning from ambiguous urns", *Statistical Papers* 43, 145-151.

[21] Mehra, Rajnish and Edward C. Prescott (1985): "The Equity Premium: A Puzzle", *Journal of Monetary Economics* 15, 145-161.

[22] Mukoyama, Toshihiko and Ayşegül Şahin (2006): "Costs of Business Cycles for Unskilled Workers", *Journal of Monetary Economics* 53 (8), 2179-2193.

[23] Otrok, Christopher; Ravikumar, B. and Charles H. Whiteman (2004): "Stochastic Discount factor Models and the Equity Premium Puzzle", mimeo.

[24] Reis, Ricardo (2007): "The Time-Series Properties of Aggregate Consumption: Implications for the Costs of Fluctuations", mimeo.

[25] Storesletten, Kjetil, Chris I. Telmer and Amir Yaron (2001): "The Welfare Costs of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk", *European Economic Review* 45(7), 1311-39.

[26] Tallarini, Thomas D. Jr. (2000): "Risk-Sensitive Real Business Cycles", *Journal of Monetary Economics* 45, 507-532.