Welfare Cost of Business Cycles in Economies with Individual Consumption Risk

Martin Ellison*

Thomas J. Sargent**

 * University of Oxford
** New York University

Contents

Welfare Cost of Business Cycles in Economies with Individual Consumption Risk

The welfare cost of random consumption fluctuations is known from De Santis (2007) to be increasing in the level of individual consumption risk in the economy. It is also known from Barillas et al. (2009) to increase if agents in the economy care about robustness to model misspecification. In this paper, we combine these two effects and calculate the cost of business cycles in an economy with consumers who face individual consumption risk and who fear model misspecification. We find that individual risk has a greater impact on the cost of business cycles if agents already have a preference for robustness. Correspondingly, we find that endowing agents with concerns about a preference for robustness is more costly if there is already individual risk in the economy. The combined effect exceeds the sum of the individual effects.

Ⅰ. Introduction

Lucas (1987, 2003) used a simple calibrated model to calculate that removing aggregate consumption fluctuations of the size observed in the post World War II US leads to welfare gains equivalent to only 0.1% of the level of consumption each period. The costs are sufficiently small that Lucas (2003) argued for prioritising growth over stabilisation policy.

One source of the small welfare costs in Lucas's calculations is his assumption that the economy is populated by a representative agent. If agents are heterogeneous then business cycles are potentially much more costly. What appears to be important is that heterogenous agents face persistent idiosyncratic shocks that are difficult to insure against. An important contribution in this respect is De Santis (2007), who assumes that consumption is subject to permanent aggregate and idiosyncratic shocks. After calibrating his model, De Santis (2007) finds that removing only 10% of the fluctuations in US aggregate consumption would generate welfare gains equivalent to a 0.5% increase in the level of consumption each period. Intuitively, idiosyncratic shocks increase the welfare costs of business cycles because they buffet the individual marginal utilities of consumption at which aggregate fluctuations are costed. A gain of 0.5% of steady-state consumption from removing only a fraction of aggregate consumption volatility is substantially higher than the 0.1% that Lucas calculated.

Tallarini (2000) said that Lucas used the wrong preferences to measure the welfare costs of business cycles. Tallarini found that business cycles are costly in a representative agent framework when agents have Epstein-Zin preferences that break the link between risk aversion and the intertemporal elasticity of substitution. Under Tallarini's interpretation of his model, costly business cycles do still require agents to have implausibly high levels of risk aversion when the model is calibrated also to generate an empirically reasonable equity premium and risk free rate. However, Barillas et al. (2009) dispose of the high risk aversion parameter needed by Tallarini (2000) by reinterpreting the high levels of risk aversion in his model not as describing attitudes towards risk but instead as reflecting a concern to make decisions that are robust to model misspecification. Building on the work on ambiguity-

averse preferences in Hansen and Sargent (2001), Barillas et al. find that agents with a preference for robustness are willing to sacrifice a lot to live in an economy without aggregate fluctuations. The benefits of removing business cycles exceed half a percent point of steady-state consumption, even when the representative agent has only a mild preference for robustness. The increase in the welfare cost of business cycles arises because concerns about model misspecification cause the representative agent's worst-case model to put more probability weight on bad consumption shocks.

This paper combines the insights of De Santis (2007) and Barillas et al. (2009) and argues that the welfare costs of business cycles are substantially higher than previously thought, and that consequently refining stabilisation policy merits at least equal priority with growth policy. To make the case we take a model in which agents are subject to aggregate and idiosyncratic consumption shocks and ask what happens if agents have a fear of model misspecification. Our framework nests the models of De Santis (2007) and Barillas et al. (2009) as special cases. To analyze the model we apply the higher order small noise expansion methods for robust economies developed by Anderson et al. (2012). Seeing that idiosyncratic risk and a preference for robustness in isolation can each increase the costs of business cycles by an order of magnitude, it might be expected that introducing the two mechanisms together would push up the costs of business cycles a lot relative to those identified by Lucas. In fact, our results go further because the welfare costs of business cycles in our model are even higher than might be expected. The reason is that an agent's fear of model misspecification extends to its understanding of aggregate and idiosyncratic consumption risks. Agents with a preference for robustness assign additional probability weight both to negative aggregate and to negative idiosyncratic consumption shocks, and make decisions to guard against worst-case scenarios in which negative aggregate and idiosyncratic shocks coincide.

The assumption that agents fear model misspecification also helps address a potential criticism of De Santis (2007), namely, that his analysis relies on non-logarithmic preferences to generate an impact of idiosyncratic risk on the costs of aggregate fluctuations. It is true in his model that if utility is logarithmic in consumption then the marginal utility of aggregate consumption is independent of the

level of idiosyncratic consumption, and the welfare costs of business cycles do not depend on the degree of idiosyncratic risk in the economy. Only if risk aversion exceeds that implied by logarithmic utility will the costs of business cycles in De Santis (2007) differ from those calculated by Lucas. This is no longer a problem when agents are assumed to have a fear of model misspecification. In our model, idiosyncratic risk affects the welfare gains of removing aggregate business cycles even under logarithmic preferences.

The paper is structured as follows. Our model with idiosyncratic risk and a preference for robustness is presented in Section Ⅱ, and solved in Section Ⅲ and the appendix. The accuracy and appropriateness of the small noise expansion method of Anderson et al. (2012) is discussed in Section Ⅳ. Section Ⅴ introduces the worst case distribution and Section Ⅵ the scaling used when calculating detection error probabilities. The results of the paper appear in Section Ⅶ, after which a final Section Ⅷ concludes.

Ⅱ. Economic environment

The agents in our economy fear model misspecification. Their preferences over individual consumption levels C_t^i are ordered by the value function:

$$U_t = (1-\beta)\,V(C_t^i) - \frac{1}{\sigma}log\,E_t\,\exp(-\,\sigma\beta U_{t+1}) \tag{1}$$

which is a version of the recursive multiplier preferences of Hansen and Sargent (2001) and Hansen et al. (2006). The preference for robustness manifests itself in the parameter $\sigma \geq 0$, which is increasing in the agent's fear of model misspecification. If $\sigma \rightarrow 0$ then the value function takes the standard expected utility form $U_t = (1-\beta)\,V(C_t^i) + \beta E_t U_{t+1}$ and the agent does not fear model misspecification. Whilst we interpret the value function in terms of a preference for robustness, it is equally valid to follow Tallarini (2000) and take it as an expression of Epstein-Zin preferences with the agent's coefficient of relative risk aversion not constrained to

be equal to the intertemporal elasticity of substitution.[1] The two interpretations have identical value functions, but we prefer to think in terms of robustness because Barillas et al. (2009) find equity premia and risk free rates observed in the US require implausibly high levels of risk aversion in Epstein-Zin preferences, but are compatible with mild fears of misspecification in Hansen-Sargent multiplier preferences.

The logarithm of individual consumption C_t^i has aggregate and idiosyncratic components that follow geometric random walk processes:

$$\log C_t^i = \log C_t + \log \delta_t^i \tag{2}$$

$$\triangle \log C_t = \sqrt{\epsilon}\, w_{1t} \tag{3}$$

$$\triangle \log \delta_t^i = \sqrt{\epsilon}\, w_{2t} \tag{4}$$

where $w_{1t} \sim N(0, \tau_1^2)$ and $w_{2t} \sim N(0, \tau_2^2)$. $\sqrt{\epsilon}$ is a scaling factor that is useful when applying a small noise expansion to solve the model. The specification is taken from De Santis (2007).[2] De Santis (2007) discusses empirical evidence in support of permanent aggregate shocks and idiosyncratic consumption shocks. Note that the processes for aggregate and idiosyncratic consumption are assumed to be independent, so the environment abstracts from endogenous relationships between aggregate and idiosyncratic uncertainty of the type explored by Atkeson and Phelan (1994) or Beaudry and Pages (2001).

III. Solution

The value function $W^E(c_t^i)$ satisfies the recursion (1) given the stochastic pro-

---

[1] Tallarini locked the intertemporal rate of substitution to unity, then allowed the risk aversion coefficient to deviate from one.

[2] De Santis's consumption process is identical with one that Constantinides and Duffie (1996) used in their reverse engineering exercise. We alter preferences relative to those assumed by Constantinides and Duffie, so the prices reverse engineered by them will not prevail in our setting.

cesses (2)-(4).

Aside from the special cases to be discussed in Section Ⅳ, the value function recursion does not have a simple analytic solution. We therefore employ the small noise expansion method of Anderson et al. (2012), and obtain a recursive representation of the solution that applies when consumption shocks are small. We seek an approximate solution of the value function recursion that takes the form:

$$W^E(c_t^i) = W^0(c_t^i) + h(c_t^i) \tag{5}$$

where $W^0(c_t^i)$ satisfies the value function in the absence of noise from aggregate and idiosyncratic shocks. $h(c_t^i)$ is a correction term in $\sqrt{\epsilon}$ and the moments of (w_{1t}, w_{2t}) that captures the costs of the small noise generated by aggregate and idiosyncratic consumption shocks. The correction term is at most of order $\sqrt{\epsilon}$ and is written in terms of the logarithm of individual consumption c_t^i.

1. Solution with no noise

The process determining the logarithm of individual consumption can be written in the form $c_{t+1}^i = c_t^i + \sqrt{\epsilon}\,(w_{1t+1} + w_{2t+1})$, so to remove noise from the model we set $\sqrt{\epsilon} = 0$. The consumption process is then $c_{t+1}^i = c_t^i$, and the value function with no noise satisfying the value function recursion is:

$$W^0(c_t^i) = V(c_t^i) \tag{6}$$

2. Approximate solution with small noise

The solution with small noise satisfies the value function for small values of the scaling factor $\sqrt{\epsilon}$ on aggregate and idiosyncratic consumption shocks. An approximate small noise solution of the form (5) solving the value function recursion (1) is therefore:

$$W^0(c_t^i) + h(c_t^i) = (1-\beta)\, W^0(c_t^i) - \frac{1}{\sigma} \log E_t \exp(-\sigma\beta(W^0(c_{t+1}^i) + h(c_{t+1}^i))) \quad (7)$$

for small $\sqrt{\epsilon}$. The small noise expansion method of Anderson et al. (2012) starts by expressing the right hand side of (7) in terms of powers of $\sqrt{\epsilon}$, and proceeds by using the method of undetermined coefficients recursively to determine the correction term $h(c_t^i)$ to any desired order of $\sqrt{\epsilon}$. We describe how to obtain the expansion in the appendix to this paper. As we indicate there, it is useful to follow Anderson et al. (2012) and write the second order approximate solution in the following compact form:

$$W^E(c_t^i) = W^0(c_t^i) + \epsilon\, [\,W_g(c_t^i)\ W_n(c_t^i)\,] \begin{bmatrix} \sigma \\ 1 \end{bmatrix} + \epsilon^2\, [\,W_{g2,2}(c_t^i)\ W_{g2,1}(c_t^i)\ W_{n2}(c_t^i)\,] \begin{bmatrix} \sigma^2 \\ \sigma \\ 1 \end{bmatrix} \quad (8)$$

The matrices $W_g(c_t^i)$, $W_n(c_t^i)$, $W_{g2,2}(c_t^i)$, $W_{g2,1}(c_t^i)$ and $W_{n2}(c_t^i)$ are independent of $\sqrt{\epsilon}$ and σ, and depend only on c_t^i and the form of the utility function $V(C_t^i)$.[3] The matrices with g subscripts emerge entirely from the representative consumer's concerns about model misspecification and vanish in the special case that $\sigma = 0$. The matrices with n subscripts emerge entirely from the consumer's concerns about risk and vanish as the risk aversion parameter γ approaches 1.[4]

---

3) These matrices are obtained by partitioning expressions (18) and (19) in the appendix.

4) The g subscript refers to 'game', indicating that the term $-\frac{1}{\sigma}\log E_t \exp(-\sigma\beta U_{t+1})$ in the value function recursion (1) is itself the indirect utility function of a minimization problem in which a malevolent agent minimizes expected utility plus $-\sigma^{-1}$ times entropy of a distorted distribution relative to the approximating model; the g subscript refers to a zero-sum game between the agent and the malevolent alter ego that chooses the probability distortion.

IV. Accuracy

The purpose of this section is to check the accuracy of the second order small noise approximation against two special cases of our economic environment that permit a simple analytic solution. The first is a simplified version of De Santis (2007) in which the agent has constant relative risk aversion preferences but no fear of model misspecification. The second is a variant of Barillas et al. (2009), with a representative agent combining logarithmic preferences with a fear of model misspecification. For both cases the agent is subject to permanent shocks to aggregate and idiosyncratic consumption. In this section we also take the opportunity to explain why second order expansions are needed to capture the welfare costs of business cycles in a model with aggregate and idiosyncratic risk.

1. CRRA preferences and no fear of misspecification

An agent with constant relative risk aversion preferences and no fear of model misspecification has a value function of the form:

$$U_t = (1 - \beta) \frac{(C_t^i)^{1-\gamma}}{1-\gamma} + \beta E_t U_{t+1} \tag{9}$$

which nests in the general framework of Section Ⅲ when the preference for robustness parameter σ tends to zero. γ is the coefficient of relative risk aversion. This specification of preferences permits a simple analytic solution when consumption follows a random walk process. De Santis (2007) gives details in his Appendix B. The analytic solution is:

$$U_t = \frac{1-\beta}{1-\beta e^{0.5(1-\gamma)^2 \epsilon (\tau_1^2 + \tau_2^2)}} \frac{(C_t^i)^{1-\gamma}}{1-\gamma} \tag{10}$$

The second order approximate solution is derived by following the steps in Section Ⅲ and the appendix and is in the general form of equation (8). The

matrices $W_n(c_t^i)$ and $W_{n2}(c_t^i)$ are functions of the second, third and fourth derivatives of the no noise solution, whereas the matrices $W_g(c_t^i)$, $W_{g2,2}(c_t^i)$ and $W_{g2,1}(c_t^i)$ are all zero because the agent does not fear model misspecification.

We are not interested $per\,se$ in whether the second order small noise expansion is a good approximation to the analytic solution. Rather, we want to know whether the dependency of the costs of business cycles on the level of idiosyncratic risk is well described by the approximate solution. The true cost of business cycles in this model is $U_t - U_t^*$, where U_t^* is the value that welfare (10) takes if aggregate fluctuations are removed. The cost of business cycles approximated by the small noise expansion is correspondingly $W^E(c_t^i) - W^{E*}(c_t^i)$, with $W^E(c_t^i)$ and $W^{E*}(c_t^i)$ being the values of welfare in (8) with and without aggregate shocks. The question is whether this approximation is close to the truth. The free parameters in the model are the discount rate, the coefficient of relative risk aversion, and the variances of aggregate and idiosyncratic consumption shocks. We set the discount factor at 0.95, the coefficient of relative risk aversion to 2, and the variance of aggregate consumption shocks at $(2.9\%)^2$ following De Santis (2007).

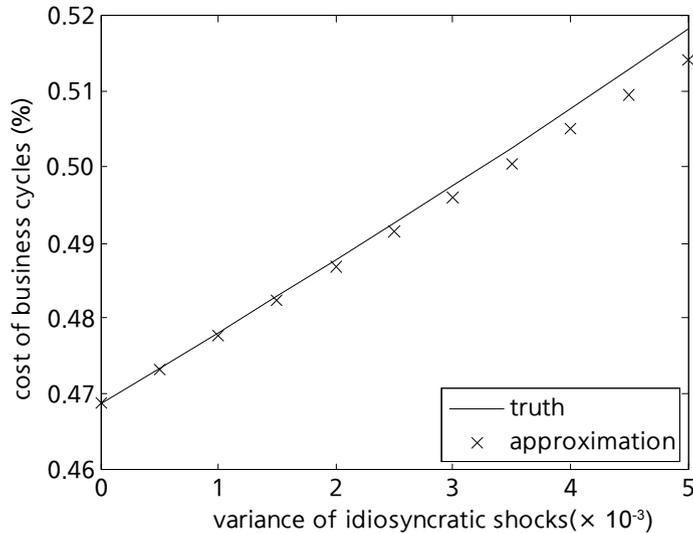## Figure 1: Accuracy of the small noise approximation

Figure 1 plots the true and approximated costs of business cycles as functions of the variance of idiosyncratic consumption shocks. The accuracy of the second order approximation is good, with a maximum discrepancy between truth and approximation of about 0.8%. The second order approximation begins to understate the true costs of business cycles when there is more noise in the economy, or as the agent becomes more risk averse. To increase accuracy, the order of the small noise expansion can always be increased. A sufficiently high order expansion is arbitrarily accurate for any level of noise, provided that the agent has CRRA preferences and no fear of misspecification.

2. Logarithmic preferences and fear of misspecification

A representative agent with logarithmic preferences and a fear of model miss-pecification has a value function of the form:

$$U_t = (1-\beta)c_t^i - \frac{1}{\sigma}\log E_t \exp(-\sigma\beta U_{t+1}) \qquad (11)$$

which is again a special case of the general framework in Section Ⅲ. Barillas et al. (2009) show in their Section Ⅳ-2 that the specification has a simple analytic solution under aggregate consumption shocks. It is straightforward to extend their calculations to include idiosyncratic shocks and obtain:

$$U_t = -\frac{1}{1-\beta}\left(\sigma\beta^2\epsilon\frac{\tau_1^2 + \tau_2^2}{2}\right) + c_t^i \qquad (12)$$

The approximate solution is constructed from the derivatives of the no-noise solution. With logarithmic preferences the no-noise solution is $W^0(c_t^i) = c_t^i$, so the first derivative is unity but all higher derivatives are zero. The only non-zero matrix in the second order expansion (8) is therefore $W_g(c_t^i)$, and the approximate solution is:

$$W^E(c_t^i) = c_t^i - \frac{\epsilon}{1-\beta}\left(\sigma\beta^2\frac{E_t\left(w_{1t}+w_{2t}\right)^2}{2}\right) \tag{13}$$

Equations (12) and (13) are identical, so the approximation is exact when the representative agent has logarithmic preferences and a fear of model misspecification. The absence of terms in ϵ^2 from the analytic solution (12) means that it is unnecessary to approximate to second order in this particular case. A first-order approximation is already exact.

3. Importance of the second order approximation

The computational burden of calculating the small noise expansion to second order is not trivial. Given that a first-order expansion is exact in the special case of Section Ⅳ-2, it is legitimate to ask whether expanding to first-order might also suffice in the general case. The answer is no. The problem is that risk enters the first-order solution only linearly in the term in $E_t\left(w_{1t+1}+w_{2t+1}\right)^2 = \tau_1^2 + \tau_2^2$ in (21), which makes the approximated cost of business cycles independent of idiosyncratic risk in the first-order expansion. A first-order expansion therefore misses the insight of De Santis (2007) that the cost of business cycles depends on idiosyncratic volatility. In the Barillas et al. (2009) model of Section Ⅳ-2, this does not matter because the cost of business cycles is independent of idiosyncratic volatility when preferences are logarithmic. However, it does matter in the model of Section Ⅳ-1 when the agent has non-logarithmic preferences. If Figure 1 is re-drawn for the first-order expansion then the ×'s of the approximation lie on a horizontal line.

The first-order approximation is even more misleading if Figure 1 is plotted in terms of the relative cost of business cycles. With the absolute cost of business cycles constant and welfare decreasing in the variance of idiosyncratic shocks, the first-order approximation predicts a negative relationship between the relative cost of business cycles and idiosyncratic risk. In reality the relationship is positive. The second order expansion does not suffer these problems because it includes a term in $E_t\left(w_{1t+1}+w_{2t+1}\right)^4 = 3\tau_1^4 + 6\tau_1^2\tau_2^2 + 3\tau_2^4$ in (22) that captures the impact of

idiosyncratic risk on the cost of aggregate fluctuations.

V. Worst-case distribution

The agents fearing misspecification in our economy distrust their approximating model for the joint distribution of $z_{t+1} \equiv (\log C_{t+1}, \log \delta_{t+1}^i)$. Their worst-case conditional joint distribution of z_{t+1} is $\hat{f}(z_{t+1}) = m(z_{t+1})f(z_{t+1})$, where $f(z_{t+1})$ is the conditional density of z_{t+1} under the approximating model and m_{t+1} satisfies:

$$m_{t+1} \propto \exp(-\sigma \beta U_{t+1}) \tag{14}$$

where the factor of proportionality makes m_{t+1} have conditional mean one, i.e. it is the reciprocal of the conditional mean of $\exp(-\sigma \beta U_{t+1})$. U_{t+1} is an approximation to the value function, in our case taken from the small-noise expansion (8).

VI. Detection error probabilities

The difference between the worst-case and approximating distributions allows Hansen et al. (2002), Anderson et al. (2003), and Hansen and Sargent (2008, ch. 9) to interpret the robustness parameter σ in terms of detection error probabilities. A higher fear of model misspecification is captured by supposing that the agent would employ more stringent standards when trying to detect whether their model is misspecified. Formally, the detection error probability approach calibrates σ in terms of the probability that the agent will make an error in a model selection test that pits the approximating model against the worst-case model associated with a given $\sigma > 0$. Lower detection error probabilities indicate larger fears of misspecification.

The non-stationary consumption process induces non-stationarity also in detection error probabilities if σ is held constant in our model. To ensure that detection error probabilities are stationary, we follow Maenhout (2004) and scale the period t robustness parameter by:

$$\sigma_t = \frac{\sigma}{W^0(c_t^i)} \tag{15}$$

This guarantees homotheticity of the value function (1) when preferences are CRRA, as seen by the correction term $h(c_t^i)$ becoming linear in $W^0(c_t^i)$ in equations (21) and (22). With $W^E(c_t^i) = W^0(c_t^i)(1+\overline{h})$, the difference between the distributions under the worst-case and approximating models depends on:

$$m_{t+1} \propto \exp\left(-\sigma\beta\frac{W^0(c_{t+1}^i)}{W^0(c_t^i)}(1+\overline{h})\right) \tag{16}$$

which is independent of c_t^i for CRRA preferences. Detection error probabilities are similarly independent of the level of consumption and hence stationary.

VII. Results

1. Calibration

The calibration of the model in Table 1 is taken from De Santis (2007). Shock variances match the time series properties of individual consumption in the US. Together with the discount factor and the coefficient of relative risk aversion, they also generate plausible risk-free rates in an economy with no fear of model misspecification.

Table 1: Calibration

Parameter	Symbol	Value
Standard deviation of aggregate consumption shocks	$\sqrt{\epsilon}\,\tau_1$	2.9%
Standard deviation of idiosyncratic consumption shocks	$\sqrt{\epsilon}\,\tau_2$	10%
Discount factor	β	0.95
Coefficient of relative risk aversion	γ	1, 1.25, 1.5
Logarithm of initial individual consumption	c_0^i	1

We calibrate the agent's preference for robustness to capture very mild fears of model misspecification. The values of the robustness parameter in Table 2 generate detection error probabilities of 50%, 45% and 40% in a sample of 100 years of annual data, for different calibrations of the coefficient of relative risk aversion parameter γ. A detection error probability of 50% indicates that the agent has no fear of model misspecification.

Table 2: Calibrated values of the initial robustness parameter $\sigma_0 = \dfrac{\sigma}{W^0(c_0^i)}$

		Detection error probability		
		50%	45%	40%
	1	0	0.23	0.48
γ	1.25	0	0.30	0.60
	1.5	0	0.36	0.75

2. Welfare cost of business cycles

The welfare cost of random consumption fluctuations in the calibrated model is presented in Table 3. The table decomposes the cost into a contribution from idiosyncratic consumption risk and a contribution from the agent's fear of model misspecification. The decomposition starts from the baseline cost of business cycles when there is no idiosyncratic risk or fear of misspecification. It is comparable to the Lucas calculation, obtained from the small noise expansion (8) by setting to zero both the variance of idiosyncratic shocks and the preference for robustness.[5]

The contribution of idiosyncratic consumption risk is isolated in the next row, which shows the additional cost of business cycles when there is idiosyncratic risk but no fear of misspecification. This is the case studied by De Santis (2007). The contribution of fears of model misspecification is in the next row, which is the additional cost of business cycles with a fear of misspecification but no idiosyncratic risk, the robust economy of Barillas et al. (2009). The cost of business cycles in Table 3 exceeds the sum of the baseline cost and the individual contributions from idiosyncratic risk and fear of misspecification. To complete the decomposition we therefore need an additional row, which measures the additional increase in the cost of business cycles when idiosyncratic risk and fear of misspecification are introduced jointly rather than separately.

That the combined effects of idiosyncratic consumption risk and a fear of model misspecification exceed their individual contributions is the central result of our paper. We find that introducing idiosyncratic risk increases the cost of business cycles more if agents have a preference for robustness. Symmetrically, endowing agents with a fear of model misspecification is more costly if the economy has idiosyncratic risk. The combination effect is quantitatively important, with a contribution of +0.123% to a baseline cost of 0.309% representing a significant increase in the absolute cost of business cycles when $\gamma = 1.5$. If the data sample used in the calculation of detection error probabilities is restricted to less than 100 years then the costs become higher.

Table 3: Decomposition of the welfare costs of business cycles

γ		1	
Detection error probability	50%	45%	40%
Cost of business cycles	0%	0.175%	0.364%
= Baseline cost	= 0	= 0	= 0
+ Contribution of idiosyncratic risk	+ 0	+ 0	+ 0
+ Contribution of robustness	+ 0	+ 0.175	+ 0.364
+ Joint contribution of idiosyncratic risk and robustness	+ 0	+ 0	+ 0

---

5) The cost of business cycles in Table 3 exceed those calculated by Lucas (1987, 2003) because aggregate consumption shocks in our model are permanent rather than temporary.

γ		1.25	
Detection error probability	50%	45%	40%
Cost of business cycles	0.178%	0.353%	0.542%
= Baseline cost	= 0.176	= 0.176	= 0.176
+ Contribution of idiosyncratic risk	+ 0.002	+ 0.002	+ 0.002
+ Contribution of robustness	+ 0	+ 0.157	+ 0.314
+ Joint contribution of idiosyncratic risk and robustness	+ 0	+ 0.018	+ 0.050
γ		1.5	
Detection error probability	50%	45%	40%
Cost of business cycles	0.324%	0.501%	0.718%
= Baseline cost	= 0.309	= 0.309	= 0.309
+ Contribution of idiosyncratic risk	+ 0.015	+ 0.015	+ 0.015
+ Contribution of robustness	+ 0	+ 0.130	+ 0.271
+ Joint contribution of idiosyncratic risk and robustness	+ 0	+ 0.048	+ 0.123

3. Intuition

The intuition for our result is that there is more for the agent to fear when they are subject to idiosyncratic risk. One part of this is because individual consumption becomes more volatile, so the agent fears misspecification not only with respect to aggregate risk but also with respect to idiosyncratic risk. Another part is because the agent fears that aggregate and idiosyncratic risks may be interrelated. For example, it could be that aggregate and idiosyncratic shocks are correlated, or that their variances co-move over time. A positive correlation between aggregate and idiosyncratic shocks is something to fear because it implies that the agent is more likely to suffer negative idiosyncratic shocks when the aggregate economy is doing badly. Co-movement between variances is similarly fear-inducing, in that it increases the likelihood that periods of idiosyncratic risk will occur against a backdrop of heightened aggregate risk. Figure 2 plots contours of the normalised difference between the worst-case density and the density under the approximating model, for a calibration of $\gamma = 1.25$ and detection error probability of 40%.[6]

---

6) The difference in densities is normalised by dividing values of $\hat{f}(z_{t+1}) - f(z_{t+1})$ by the constant $\int (\hat{f}(z_{t+1}) - f(z_{t+1})) dz_{t+1}$. This makes contours in Figure 2 independent of the fineness of the z_{t+1} grid.

Figure 2: Difference in densities for $\gamma = 1.25$
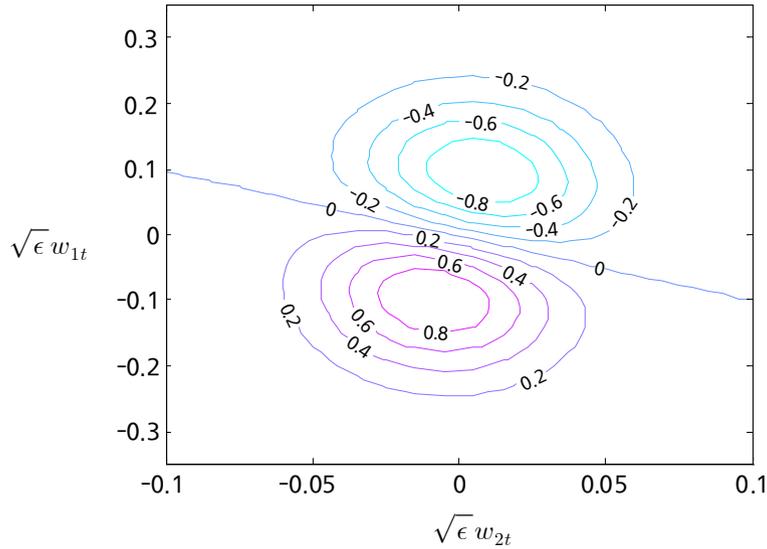and 40% detection error probability

Table 4 summarises the moments of the worst-case conditional consumption density for different levels of risk aversion and detection error probability. The worst-case density differs from the density under the approximating model in three respects. Firstly, it is centered on negative values of the aggregate consumption shock $\sqrt{\epsilon}\,w_{1t}$ and the idiosyncratic consumption shock $\sqrt{\epsilon}\,w_{2t}$. Secondly, it has slightly higher variance. Thirdly, the worst-case density has greater probability mass than the density under the approximating model in the region where both aggregate and idiosyncratic shocks are negative. The characteristics of the worst-case density confirm the intuition that an agent with a preference for robustness fears negative shocks to aggregate and idiosyncratic consumption, especially if they occur at the same time. Table 4 shows the degree to which these effects attenuate as the agent becomes less risk averse and less concerned with robustness. The covariance between aggregate and idiosyncratic shocks in the worst-case density tends to zero in the limit of either logarithmic utility or no fear of model misspecification.

Table 4: Moments of worst-case density

γ	1		
Detection error probability	50%	45%	40%
$\overline{w}_{1t}$	0	-0.184×10^{-3}	-0.383×10^{-3}
$\overline{w}_{2t}$	0	-2.185×10^{-3}	-4.560×10^{-3}
τ_1	2.9%	2.9%	2.9%
τ_2	10%	10%	10%
$\rho_{w1,w2}$	0	0	0
γ	1.25		
Detection error probability	50%	45%	40%
$\overline{w}_{1t}$	0	-0.190×10^{-3}	-0.385×10^{-3}
$\overline{w}_{2t}$	0	-2.263×10^{-3}	-4.588×10^{-3}
τ_1	2.9%	2.90006%	2.90014%
τ_2	10%	10.0056%	10.0058%
$\rho_{w1,w2}$	0	0.166×10^{-3}	0.342×10^{-3}
γ	1.5		
Detection error probability	50%	45%	40%
$\overline{w}_{1t}$	0	-0.184×10^{-3}	-0.396×10^{-3}
$\overline{w}_{2t}$	0	-2.189×10^{-3}	-4.705×10^{-3}
τ_1	2.9%	2.90013%	2.90030%
τ_2	10%	10.0056%	10.0124%
$\rho_{w1,w2}$	0	0.326×10^{-3}	0.724×10^{-3}

VIII. Conclusions

The message that the welfare cost of business cycles is higher in an economy with individual who face idiosyncratic consumption risk and are concerned about statistical model misspecification will come as no surprise to readers of De Santis (2007) and Barillas et al. (2009). What might be surprising is that the total effects of idiosyncratic risk and concerns about robustness exceed the sum of their parts. Our calibration of the quantitative magnitude of the total effects offers additional support to the view that aggregate fluctuations are substantially more costly to individual consumers than calculated by Lucas (1987, 2003). Relative to Lucas's (2003) numbers, our calibration raises stabilisation policies and other determinants of macroeconomic fluctuations as macroeconomic priorities.

Appendix A: The small noise expansion

In what follows, we outline how we apply a small noise expansion to obtain an approximate solution of value function recursion (1) in our economic environment. For details of the general case, see Anderson et al. (2012).

The first step in expressing the right hand side of (7) in powers of $\sqrt{\epsilon}$ is to expand the exponential expression over which an agent forms expectations through the operator E_t. The Maclaurin expansion for small $\sqrt{\epsilon}$ of this expression is:

$$e^{-\sigma\beta(W^0(c_{t+1}^i) + h(c_{t+1}^i))} = e^{-\sigma\beta(W^0(c_t^i) + h(c_t^i))} \left(1 + \sum_{n=1}^{\infty} \frac{\epsilon^{n/2}\mu_n'}{n!}(w_{1t+1} + w_{2t+1})^n\right),$$

$$\mu_n' = \kappa_n + \sum_{m=1}^{n-1} \binom{n-1}{m-1}\kappa_m\mu_{n-m}', \quad \kappa_n = -\sigma\beta(D^n W^0(c_t^i) + D^n h(c_t^i)) \tag{17}$$

which is of the same form as the moment generating function of a random variable, with μ_n' and κ_n analogous to the random variable's non-central moments and cumulants. The second step involves taking expectations of (14), applying logarithms, multiplying by $-\sigma^{-1}$, and then making a Maclaurin expansion of the resulting expression for small $\sqrt{\epsilon}$. When taking expectations of (14), the operator E_t only applies to the expression $(w_{1t+1} + w_{2t+1})^n$, because μ_n' is a deterministic function of individual consumption c_t^i, which the agent knows at time t. The result of the second step is an expression for the right hand side of (7) in terms of powers of $\sqrt{\epsilon}$, as required:

$$W^0(c_t^i) + \beta h(c_t^i) - \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k\sigma}\left(\sum_{n=1}^{\infty} \frac{\epsilon^{n/2}\mu_n'}{n!} E_t(w_{1t+1} + w_{2t+1})^n\right)^k \tag{18}$$

We are now in a position to use the method of undetermined coefficients to identify the correction term $h(c_t^i)$ in the small noise approximation. Substituting expression (15) into the right hand side of (7) implies that the correction term satisfies:

$$(1-\beta)h(c_t^i) = -\sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k\sigma} \left(\sum_{n=1}^{\infty} \frac{\epsilon^{n/2}\mu'_n}{n!} E_t (w_{1t+1} + w_{2t+1})^n \right)^k \tag{19}$$

The way that $\sqrt{\epsilon}$ enters the right hand side of (16) directly through the expression $\epsilon^{n/2}$ and indirectly through the dependency of μ'_n on $\kappa_{n\ldots1}$ and the derivatives $D^{n\ldots1} W^0(c_t^i) + D^{n\ldots1} h(c_t^i)$ suggests that the correction term may be written as a linear function of powers of $\sqrt{\epsilon}$. A candidate is:

$$h(c_t^i) = \sum_{n=1}^{\infty} \epsilon^{n/2} h^n(c_t^i) \tag{20}$$

where the functions $h^n(c_t^i)$ are independent of $\sqrt{\epsilon}$. We verify the form of the candidate solution and identify its functions by applying the method of undetermined coefficients to equations (16) and (17). For odd values of n, we immediately have that $h^n(c_t^i) = 0$ because the coefficient on $\epsilon^{n/2}$ when n is odd always includes a term $E_t (w_{1t} + w_{2t})^p$ for some odd $p \leq n$. The expected value $E_t (w_{1t} + w_{2t})^p$ is in turn given by the central moments of the bivariate distribution of w_{1t} and w_{2t}. In our model the shocks have normal distributions, which by definition have odd central moments that are zero. It follows that $h^n(c_t^i) = 0$ whenever n is odd. The method of undetermined coefficients applied to terms in ϵ in (16) and (17) gives:

$$h^2(c_t^i) = \frac{E_t (w_{1t+1} + w_{2t+1})^2}{2\sigma(1-\beta)} \left(\sigma\beta D^2 W^0(c_t^i) - (\sigma\beta D W^0(c_t^i))^2 \right) \tag{21}$$

There are no derivatives of $h(c_t^i)$ in the first-order correction term (18) because $h(c_t^i)$ is at most of order ϵ. The second order correction term is defined by the undetermined coefficients on ϵ^2 in (16) and (17):

$$h^4(c_t^i) = \frac{E_t(w_{1t+1} + w_{2t+1})^2}{2\sigma(1-\beta)} \left(\sigma\beta D^2 h^2(c_t^i) - 2\sigma^2\beta^2 DW^0(c_t^i) Dh^2(c_t^i) \right)$$

$$+ \frac{E_t(w_{1t+1} + w_{2t+1})^4}{24\sigma(1-\beta)} \begin{pmatrix} \sigma\beta D^4 W^0(c_t^i) - 4\sigma^2\beta^2 D^3 W^0(c_t^i) DW^0(c_t^i) \\ -3(\sigma\beta D^2 W^0(c_t^i))^2 + 6\sigma^3\beta^3 D^2 W^0(c_t^i)(DW^0(c_t^i))^2 \\ -(\sigma\beta DW^0(c_t^i))^4 \end{pmatrix}$$

$$+ \frac{(E_t(w_{1t+1} + w_{2t+1})^2)^2}{8\sigma(1-\beta)} \left(\sigma\beta D^2 W^0(c_t^i) - (\sigma\beta DW^0(c_t^i))^2 \right)^2 \tag{22}$$

where the derivatives $Dh^2(c_t^i)$ and $D^2 h^2(c_t^i)$ can be calculated from equation (21).

The small noise expansion can be constructed to any desired order of ϵ by recursively calculating the higher order correction terms. For our purposes, a second order expansion is necessary but sufficient. To isolate the role of the robustness in our paper, it is useful to follow Anderson et al. (2012) and write the second order approximate solution in compact form:

$$W^E(c_t^i) = W^0(c_t^i) + \epsilon[W_g(c_t^i) \ W_n(c_t^i)] \begin{bmatrix} \sigma \\ 1 \end{bmatrix} + \epsilon^2[W_{g2,2}(c_t^i) \ W_{g2,1}(c_t^i) \ W_{n2}(c_t^i)] \begin{bmatrix} \sigma^2 \\ \sigma \\ 1 \end{bmatrix} \tag{23}$$

The matrices $W_g(c_t^i)$, $W_n(c_t^i)$, $W_{g2,2}(c_t^i)$, $W_{g2,1}(c_t^i)$ and $W_{n2}(c_t^i)$ are independent of $\sqrt{\epsilon}$ and σ, and depend only on c_t^i and the form of the utility function $V(C_t^i)$. They are obtained by partitioning equations (21) and (22).

References

Anderson, E. W., L. P. Hansen, and T. J. Sargent (2003), "A quartet of semi-groups for model specification, robustness, prices of risk, and model detection," *Journal of the European Economic Association* 1(1), pp. 68-123.

Anderson, E. W., L. P. Hansen, and T. J. Sargent (2012), "Small noise methods for risksensitive/robust economies," *Journal of Economic Dynamics and Control* 36, pp. 468-500.

Atkeson, A. and C. Phelan (1994), "Reconsidering the Costs of Business Cycles with Incomplete Markets," *NBER Macroeconomic Annual*, Cambridge, MA: MIT Press.

Barillas, B., L. P. Hansen, and T. J. Sargent (2009), "Doubts or variability?," *Journal of Economic Theory* 144, pp. 2388-2418.

Beaudry, P. and C. Pages (2001), "The Cost of Business Cycles and the Stabilization Value of Unemployment Insurance," *European Economic Review* 45(8), pp. 1545-1572.

Constantinides, G. M. and D. Duffie (1996), "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy* 104(2), pp. 219-240.

De Santis, M. (2007), "Individual Consumption Risk and the Welfare Cost of Business Cycles," *American Economic Review* 97(4), pp. 1488-1505.

Hansen, L. P. and T. J. Sargent (2001), "Robust Control and Model Uncertainty," *American Economic Review* 91(2), pp. 60-66.

Hansen, L. P. and T. J. Sargent (2008), "Robustness, Princeton, New Jersey," Princeton University Press.

Hansen, L. P., T. J. Sargent, G. A. Turmuhambetova, and N. Williams (2006), "Robust Control, Min-Max Expected Utility, and Model Misspecification," *Journal of Economic Theory* 128, pp. 45-90.

Hansen, L. P., T. J. Sargent, and N. E. Wang (2002), "Robust permanent income and pricing with filtering," *Macroeconomic Dynamics* 6, pp. 40-84.

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

Lucas, R. E. (2003), "Macroeconomic Priorities," *American Economic Review* 93(1), pp. 1-14.

Maenhout, L. J. (2004), "Robust Portfolio Rules and Asset Pricing," *The Review of Financial Studies* 17(4), pp. 951-983.

Tallarini, T. D. (2000), "Risk-sensitive real business cycles," *Journal of Monetary Economics* 45(3), pp. 507-532.

<Abstract in Korean>

Martin Ellison*, Thomas J. Sargent**

De Santis(2007)에 의하면 개별가계의 소비위험(consumption risk)이 클수록 총소비 변동에 따른 복지비용이 증가하게 된다. 또한 가계가 모형 식별오류에 대한 강건성(robustness to model misspecification)을 선호하는 경우 총소비 변동에 따른 복지비용이 늘어날 수 있다(Barillas et al., 2009). 본고에서는 이러한 두 논점을 집합적으로 고려하여 경기변동의 복지비용을 이론적으로 계산하였다. 개별 가계는 소비위험에 노출되어 있고 모형 식별오류에 대한 강건성을 선호하는 것으로 모형에서 가정하였다.

분석결과, 개별 가계의 소비위험과 가계의 모형 식별오류에 대한 강건성 선호는 서로 상승효과를 일으켜 경기변동에 따른 복지비용이 개별효과의 합보다 더 크게 증폭되는 것으로 나타났다. 이는 모형 식별오류에 대한 강건성을 선호하는 가계는 전체 소비위험과 개별 소비위험을 동시에 고려할 뿐만 아니라 두 소비위험이 동시에 부정적인 방향으로 발생할 가능성에 높은 가중치를 두기 때문이다.

 * University of Oxford
** New York University