Welfare Cost of Business Cycles with Idiosyncratic Consumption Risk and a Preference for Robustness

By MARTIN ELLISON AND THOMAS J. SARGENT*

The welfare cost of random consumption fluctuations is known from De Santis (2007) to be increasing in the level of uninsured idiosyncratic consumption risk. It is known from Barillas, Hansen and Sargent (2009) to increase if agents care about robustness to model misspecification. We calculate the cost of business cycles in an economy where agents face idiosyncratic consumption risk and fear model misspecification, finding that idiosyncratic risk has a greater impact on the cost of business cycles if agents already fear model misspecification. Correspondingly, endowing agents with fears about misspecification is more costly when there is already idiosyncratic risk.

JEL: E32, E63, D81

Keywords: Cost of Business Cycles, Idiosyncratic Risk, Model Uncertainty, Robustness

Robert E. Lucas (1987, 2003) used a calibrated representative agent model to calculate that removing aggregate consumption fluctuations of the size observed in post World War II US data leads to welfare gains equivalent to only 0.1% of the level of consumption each period. Subsequently, researchers have considered variations on the Lucas calculation and identified two modifications that can lead to larger welfare costs of business cycles: changes in the stochastic structure of consumption and changes in the preferences that agents use to value stochastic consumption streams.

The stochastic nature of consumption was changed by Massimiliano De Santis (2007), who calculated the welfare costs of business cycles when individual consumption is subject to both aggregate and idiosyncratic shocks. He found there are large gains to reducing the aggregate component of consumption fluctuations. For example, removing only 10% of the variation induced by aggregate shocks can generate a welfare gain equivalent to 0.5% of the level of consumption each period. Idiosyncratic shocks increase the welfare costs of business cycles because they affect the individual marginal utilities of consumption at which aggregate fluctuations are costed.

How individuals value stochastic streams was the subject of Francisco Barillas, Lars Peter Hansen and Thomas J. Sargent (2009), who extended the Lucas calculation by assuming that the representative agent has a preference for robustness.¹ They showed

* Ellison: University of Oxford, Department of Economics, Manor Road Building, Manor Road, Oxford, OX1 2UQ, United Kingdom, martin.ellison@economics.ox.ac.uk. Sargent: New York University and Hoover Institution, Department of Economics, New York University, 19 W. Fourth Street, NY, NY 10012-1119, United States of America, thomas.sargent@nyu.edu. We thank three anonymous referees and Evan Anderson for helpful comments, and Pascal Paul for excellent research assistance. Funding from the Bank of Korea, the British Academy, and the Oxford Fidelity Research Fund is also gratefully acknowledged.

¹Thomas D. Tallarini (2000) similarly argued that Lucas used the wrong preferences to measure the welfare costs

that with a preference for robustness to misspecification of the consumption process, the welfare gain of removing aggregate fluctuations in consumption can easily exceed 0.5%. Building on the work on ambiguity-averse preferences in Lars Peter Hansen and Thomas J. Sargent (2001), the reasoning is that agents with a preference for robustness are willing to sacrifice a lot to live in an economy without fear of misspecification. 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.

Our paper combines the insights of De Santis (2007) and Barillas, Hansen and Sargent (2009) and argues that the welfare costs of business cycles are important at lower levels of risk aversion and fear of model misspecification than previously thought. 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 fear model misspecification. Our framework nests the models of De Santis and Barillas, Hansen and Sargent as special cases. To analyze the model we apply the higher order small noise expansion methods for robust economies developed by Evan W. Anderson, Lars Peter Hansen and Thomas J. Sargent (2012).

Recognising that idiosyncratic risk and a fear of model misspecification are in isolation only weak mechanisms for generating welfare costs in economies with low levels of risk aversion and very mild fears of model misspecification, it might be expected that introducing the two mechanisms together would not push up the costs of business cycles very much relative to those identified by Lucas. This is not the case; our results go further because the welfare costs of business cycles in our model embody interaction effects between the two mechanisms. 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 occur simultaneously.

The assumption that agents fear model misspecification addresses a criticism of De Santis (2007), namely, that his analysis relies on highly non-logarithmic preferences to generate an impact of idiosyncratic risk on the costs of aggregate fluctuations. In his baseline calibration the coefficient of relative risk aversion is 2, and most of his headline results are drawn from calculations in which the coefficient of relative risk aversion is set at 4. If the utility function is instead logarithmic in consumption, then the marginal utility of aggregate consumption is completely independent of the level of idiosyncratic consumption, and the welfare costs of business cycles in his model are unaffected by the amount of idiosyncratic risk in the economy. Only if risk aversion is significantly in excess of that implied by logarithmic utility will the costs of business cycles in De Santis differ from those calculated by Lucas. This is no longer the case when agents fear model

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 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.

misspecification. In our model, idiosyncratic risk affects the welfare gains of removing aggregate business cycles even under logarithmic preferences.

The introduction of idiosyncratic risk alternatively speaks to concerns over how to interpret the claim of Barillas, Hansen and Sargent (2009) that they only endow agents with mild fears of model misspecification. Their calculations for the welfare costs of business cycles assume that agents entertain a relatively large set of models as measured by detection error probabilities based on 235 quarterly observations. Introducing idiosyncratic risk, though, implies that business cycles are costly even when the set of models being considered is small. The fears of model misspecification needed to make aggregate fluctuations costly are therefore even milder when there is idiosyncratic consumption risk in the economy.

A second issue with De Santis (2007) is that the process he assumes for consumption is problematic empirically. Whilst he goes to some trouble to justify his random walk processes in aggregate and idiosyncratic consumption by appealing to a no-trade theorem, the idea that all variation in consumption growth can be attributed to a random walk process for individual consumption does not fit with empirical studies. For example, Richard Blundell, Luigi Pistaferri and Ian Preston (2008) explicitly focus on the extent of consumption smoothing that different individuals achieve in the face of permanent income shocks. A key result is that the extent of insurance is quite high for a large fraction of the population. Similarly, the work of Jonathan Heathcote, Kjetil Storesletten and Giovanni L. Violante (2014) does attribute some of the fluctuations in consumption to a unit root process, but this is by no means the only source of variation. The bottom line is that De Santis is likely overstating the case by assuming that all consumption variation is reflected in unit root processes. If one moves to more empirically reasonable specifications then the costs that De Santis finds are much less impressive. We address this issue in two ways. Firstly, we show that fears of model misspecification render business cycles costly even when the variance of permanent idiosyncratic consumption shocks is smaller than that assumed by De Santis. Secondly, we introduce a fear of model misspecification into the more realistic economic environment of Tom Krebs (2007), where an agent's stochastic consumption stream is subject to the risk of job displacement. Empirical work by Louis S. Jacobsen, Robert J. LaLonde and Daniel G. Sullivan (2009) has shown that job displacement risk is large and has a substantial cyclical component for certain groups of workers. If agents fear model misspecification then the cyclical component of job displacement risk creates large welfare costs of business cycles, even if their utility function is close to logarithmic in consumption.

I. Economic environment

Agents in our economy fear model misspecification. Their preferences over consumption processes are ordered by the value function recursion:

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

where U_t is the value function and $V(c_t^i)$ is the period utility from c_t^i , the logarithm of individual consumption. This is a version of the recursive multiplier preferences of Hansen and Sargent (2001) and Lars Peter Hansen et al. (2006). The preference for robustness to misspecification of the consumption process manifests itself in the parameter $\sigma \geq 0$ that is increasing in the agent's fear of model misspecification. If $\sigma \rightarrow 0$ then the value function recursion 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.² The two interpretations have identical value functions, but we prefer to think in terms of robustness because Barillas, Hansen and Sargent (2009) find equity premia and risk free rates observed in the US data require implausibly high levels of risk aversion in Epstein-Zin preferences, but are compatible with relatively moderate 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:

$$\begin{aligned} (2) \quad & c_t^i = c_t + \delta_t^i, \\ (3) \quad & \Delta c_t = \sqrt{\epsilon} w_{1t}, \\ (4) \quad & \Delta \delta_t^i = \sqrt{\epsilon} w_{2t}, \end{aligned}$$

$$\text{where } \begin{bmatrix} w_{1t} \\ w_{2t} \end{bmatrix} \sim N \left[\begin{pmatrix} g - \tau_1^2/2 \\ -\tau_2^2/2 \end{pmatrix}; \begin{pmatrix} \tau_1^2 & 0 \\ 0 & \tau_2^2 \end{pmatrix} \right],$$

and aggregate consumption trends at the rate g . The means of w_{1t} and w_{2t} are adjusted by $-\tau_1^2/2$ and $-\tau_2^2/2$ so that changes in τ_1^2 and τ_2^2 affect only the variance of consumption growth, not its trend growth rate. $\sqrt{\epsilon}$ is a scaling factor that is useful when applying the small noise expansion. The specification is taken from De Santis (2007), who presents empirical evidence in support of permanent aggregate 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 Andrew Atkeson and Christopher Phelan (1994) or Paul Beaudry and Carmen Pages (2001).

II. Solution

The value function U_t satisfies recursion (1) given the stochastic processes (2)-(4). Aside from the special cases to be discussed below, the value function recursion does not

²Tallarini locked the intertemporal rate of substitution to unity, then allowed the risk aversion coefficient to deviate from one.

have a simple analytic solution. We therefore employ the small noise expansion method of Anderson, Hansen and Sargent (2012) to obtain a recursive representation that applies when the trend and shocks to consumption are small. We seek an approximate solution $W^E(c_t^i)$ that takes the form:

$$(5) \quad W^E(c_t^i) = W^0(c_t^i) + h(c_t^i),$$

where $W^0(c_t^i)$ satisfies the recursion in the absence of the trend in aggregate consumption and the 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 trend and the noise generated by the shocks. The correction term depends on the logarithm of individual consumption c_t^i and is at most of order $\sqrt{\epsilon}$.

A. Approximate solution

The solution with no trend and noise $W^0(c_t^i)$ is obtained by noting that processes (2)-(4) for the logarithm of individual consumption can be combined as $c_{t+1}^i = c_t^i + \sqrt{\epsilon}(w_{1t} + w_{2t})$. To remove the trend and shocks we set $\sqrt{\epsilon} = 0$. The consumption process becomes $c_{t+1}^i = c_t^i$ and the solution to recursion (1) without trend and shocks is:

$$(6) \quad W^0(c_t^i) = V(c_t^i).$$

The approximate solution $W^E(c_t^i)$ satisfies recursion (1) for small values of the scaling factor $\sqrt{\epsilon}$ on the trend and noise from shocks. The small noise approximation therefore solves:

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

for small $\sqrt{\epsilon}$. The small noise expansion method of Anderson, Hansen and Sargent (2012) starts by expressing the second term on the right side of (7) in terms of powers of $\sqrt{\epsilon}$, and proceeds by using the method of undetermined coefficients recursively to find the correction term $h(c_t^i)$ to any desired order of $\sqrt{\epsilon}$. We describe how to obtain the expansion in Appendix A.

B. Accuracy

The accuracy of the second order small noise approximation can be checked against two special cases of our economic environment that have simple analytic solutions. 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, Hansen and Sargent (2009), with a representative agent combining logarithmic preferences with a fear of model misspecification. The two special cases correspond to simple parameter restrictions on the model, the first to the preference for robustness

parameter σ tending to zero and the second to the coefficient of risk aversion tending to one within a CRRA period utility function.

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 U_t that satisfies:

$$(8) \quad U_t = (1 - \beta) \frac{(e^{c_t^i})^{1-\gamma}}{1 - \gamma} + \beta E_t U_{t+1},$$

which is a special case of recursion (1) when the preference for robustness parameter σ tends to zero and the period utility $V(c_t^i)$ is CRRA. $e^{c_t^i}$ is the level of individual consumption and 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:

$$(9) \quad U_t = \frac{1 - \beta}{1 - \beta e^{\sqrt{\epsilon}(1-\gamma)(g - (\tau_1^2 + \tau_2^2)/2) + (1-\gamma)^2 \epsilon (\tau_1^2 + \tau_2^2)/2}} \frac{(e^{c_t^i})^{1-\gamma}}{1 - \gamma}.$$

The second order approximate solution is derived by following the steps in our Appendix A. 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 in percentage terms is $\log(U_t) - \log(U_t^*)$, where U_t^* is the value that welfare (9) takes if aggregate fluctuations are removed. The cost of business cycles approximated by the small noise expansion is correspondingly $\log(W^E(c_t^i)) - \log(W^{E*}(c_t^i))$, with $W^E(c_t^i)$ and $W^{E*}(c_t^i)$ being the values of welfare in (5) with and without aggregate shocks. Our question is whether this approximation is good. 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 1.5, and the standard deviation of aggregate consumption shocks at 2.9% following De Santis (2007).

Figure 1 plots the true and approximated welfare 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 0.0005%.³ The second order approximation begins to understate the true costs of business cycles when there is more noise in the economy. To increase accuracy, the order of the small noise expansion can always be increased. A sufficiently high order

³The true value of $\log(U_t) - \log(U_t^*)$ from equation (9) is 0.5067% for $\tau_2^2 = 0$ on the far left of Figure 1. The approximate value $\log(W^E(c_t^i)) - \log(W^{E*}(c_t^i))$ from equation (5) is 0.5062%, hence the maximum discrepancy is 0.0005%.

FIGURE 1. ACCURACY OF THE SMALL NOISE APPROXIMATION.

expansion is arbitrarily accurate for any level of noise, provided that the agent has CRRA preferences and no fear of misspecification.

The computational burden of calculating the small noise expansion to second order is not trivial and it is reasonable to ask whether expanding to first-order might be sufficient. The answer is no. The issue is that risk enters the first-order solution only linearly in the term in $E_t(w_{1t+1} + w_{2t+1})^2 = \tau_1^2 + \tau_2^2 + \text{terms independent of risk}$ in the first order term (A6), 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. If Figure 1 were to be re-drawn for the first-order expansion, then the $\times$'s of the approximation would lie on a horizontal line. The first-order approximation is even more misleading if Figure 1 were to be 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 asserts 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 $E_t(w_{1t+1} + w_{2t+1})^4 = 3\tau_1^4 + 6\tau_1^2\tau_2^2 + 3\tau_2^4 + \text{terms independent of risk}$ that captures the impact of idiosyncratic risk on the cost of aggregate fluctuations.

LOGARITHMIC PREFERENCES AND FEAR OF MISSPECIFICATION. — A representative agent with logarithmic period utility and a fear of model misspecification has a value function

that satisfies:

$$(10) \quad U_t = (1 - \beta)c_t^i - \frac{1}{\sigma} \log E_t \exp(-\sigma \beta U_{t+1}),$$

which is again a special case of the general framework in Section 2 with the period utility equal to c_t^i , the logarithm of individual consumption. Barillas, Hansen and Sargent (2009) show in their Section 4.1 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:

$$(11) \quad U_t = \frac{\beta}{1 - \beta} \left(\sqrt{\epsilon} \left(g - \frac{\tau_1^2 + \tau_2^2}{2} \right) - \sigma \beta \epsilon \frac{\tau_1^2 + \tau_2^2}{2} \right) + c_t^i.$$

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 approximate solution is then identical to (11), 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 means that it is unnecessary to approximate to second order in this particular case. A first-order approximation is already exact.

III. Calibration

The calibration of the model in Table 1 is taken from De Santis (2007), whose calculations are based on data from 1929-1998.⁴ Shock variances match the time series properties of individual consumption in the US. The initial value of consumption is set at \$77.4 billion, the value of Personal Consumption Expenditures reported by the Bureau of Economic Analysis for 1929.

We calibrate the agent's preference for robustness to embody very mild fears of model misspecification at different levels of risk aversion. A higher fear of model misspecification is captured by supposing that an agent would employ more stringent standards when trying to detect whether its approximating model is misspecified.⁵ The values of the robustness parameter in Table 2 generate detection error probabilities of 50%, 45% and 40% when testing based on a sample period of 69 years, the same length as the data used by De Santis (2007). A detection error probability of 50% indicates that the agent has no fear of model misspecification.

⁴Lucas (1997, 2003) uses data from the post WWII period to calibrate $\sqrt{\epsilon} \tau_1$ to about 2.0%. We retain the calibration of De Santis but acknowledge that the welfare cost of business cycles would be lower if we were to adopt the Lucas calibration. Agents necessarily find that smaller aggregate fluctuations are less costly.

⁵Formally, 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. See Lars Peter Hansen, Thomas J. Sargent and Neng E. Wang (2002), Evan W. Anderson, Lars Peter Hansen and Thomas J. Sargent (2003), and Lars Peter Hansen and Thomas J. Sargent (2008, ch. 9) for more details.

TABLE 1—CALIBRATION.

Parameter	Symbol	Value
Mean consumption growth	g	1.89%
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
Initial individual consumption	c_0^i	$\log(77.4)$

TABLE 2— CALIBRATED VALUES OF THE ROBUSTNESS PARAMETER σ .

		Detection error probability		
		50%	45%	40%
γ	1	0	0.3	0.6
	1.25	0	1.1	2.0
	1.5	0	3.5	6.0

IV. Results

A. Consumption equivalent cost of business cycles

The welfare cost of business cycles calculated from the small noise expansion is $\log(W^E(c_t^i)) - \log(W^{E*}(c_t^i))$, where $W^E(c_t^i)$ and $W^{E*}(c_t^i)$ respectively measure welfare with and without aggregate shocks. To translate welfare costs into the consumption equivalent metric preferred by Lucas (1987), we define Δ as the percentage increase in initial consumption needed to compensate the agent for the presence of aggregate risk. Mathematically, Δ solves:

$$(12) \quad W^E(c_t^i(1 + \Delta)) = W^{E*}(c_t^i).$$

The consumption equivalent cost of business cycles is presented in Table 3. The table decomposes the cost into a contribution from idiosyncratic consumption risk and a contribution from an 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 (5) by setting to zero both the variance of idiosyncratic shocks and the pref-

erence for robustness.⁶ 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 fear 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 economy of Barillas, Hansen and Sargent (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 operate jointly rather than separately.

TABLE 3—DECOMPOSITION OF THE CONSUMPTION EQUIVALENT COSTS OF BUSINESS CYCLES

<i>detection error probability</i>	50%	45%	40%
γ	1	1	1
Cost of business cycles	0.802%	1.032%	1.262%
= Baseline cost	= 0.802	= 0.802	= 0.802
+ Contribution of idiosyncratic risk	+0	+0	+0
+ Contribution of fear of model misspecification	+0	+0.230	+0.460
+ Joint contribution	+0	+0	+0
γ	1.25	1.25	1.25
Cost of business cycles	0.945%	1.224%	1.476%
= Baseline cost	= 0.918	= 0.918	= 0.918
+ Contribution of idiosyncratic risk	+0.027	+0.027	+0.027
+ Contribution of fear of model misspecification	+0	+0.227	+0.413
+ Joint contribution	+0	+0.053	+0.117
γ	1.5	1.5	1.5
Cost of business cycles	1.088%	1.428%	1.706%
= Baseline cost	= 1.020	= 1.020	= 1.020
+ Contribution of idiosyncratic risk	+0.068	+0.068	+0.068
+ Contribution of fear of model misspecification	+0	+0.233	+0.402
+ Joint contribution	+0	+0.107	+0.217

That the combined effects of idiosyncratic consumption risk and fear of model mis-

⁶The cost of business cycles in Table 3 exceeds those calculated by Lucas (1987, 2003) because aggregate consumption shocks in our model are permanent rather than temporary.

specification exceed their individual contributions is a central result of our paper. We find that introducing idiosyncratic risk significantly increases the cost of business cycles, even if agents are only slightly risk averse and have only a mild fear of model misspecification. For example, the consumption equivalent cost of adding idiosyncratic risk when $\gamma = 1.5$ is only 0.068% if the agent does not fear model misspecification, but this rises to 0.402% under a fear of model misspecification. We therefore argue that idiosyncratic risk can be costly even when risk aversion is much lower than that assumed by De Santis (2007). Symmetrically, endowing agents with a fear of model misspecification is more costly in an economy subject to idiosyncratic consumption risk. The Barillas, Hansen and Sargent (2009) mechanism by which fears of model misspecification make business cycles costly is therefore amplified. We see that even a very mild fear of misspecification is sufficient to induce significant losses in consumption equivalent terms, for example when $\gamma = 1.5$ the consumption equivalent cost of fearing model misspecification rises from 0.327% to 0.493% if there is idiosyncratic risk in the economy. If the data sample used in the calculation of detection error probabilities were less than 69 years then the costs would be higher.

B. Interpretation

The force behind our result is that there is more for an agent to fear when it is subject to idiosyncratic risk. One part of this is that, as individual consumption becomes more volatile, the agent fears misspecification not only of aggregate risk but also of idiosyncratic risk. Another aspect is that an agent might fear 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 because it increases the likelihood that periods of idiosyncratic risk will occur against a backdrop of heightened aggregate risk.

Precisely how fears of model misspecification translate into welfare losses is captured by the worst-case scenarios that agents construct to value alternative consumption streams. An agent fearing misspecification distrusts its approximating model for the joint distribution of $z_{t+1} \equiv (c_{t+1}, \delta_{t+1}^i)$. A 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:

$$(13) \quad m_{t+1} \propto \exp(-\sigma \beta U_{t+1}),$$

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 (5). For a calibration of $\gamma = 1.25$ and detection error probability of 40%, Figure 2 plots contours of the normalised difference between the worst-case density and the density under the approxi-

imating model.⁷ The worst-case density has more probability mass than the density under the approximating model in the bottom left of the figure where the contours are positive. The opposite is true in the top right where the contours are negative.

FIGURE 2. DIFFERENCE IN DENSITIES FOR $\gamma = 1.25$ AND 40% DETECTION ERROR PROBABILITY.

Note: The worst-case density has more (less) probability mass than the density under the approximating density where the contours are positive (negative).

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 variances. 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, and especially fears that these shocks may occur at the same time. Table 4 shows

⁷The 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.

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

<i>detection error probability</i>	50%	45%	40%
γ	1	1	1
$\hat{w}_{1t}$	0	-0.240×10^{-3}	-0.479×10^{-3}
$\hat{w}_{2t}$	0	-2.850×10^{-3}	-5.700×10^{-3}
$\hat{\tau}_1$	2.9%	2.9%	2.9%
$\hat{\tau}_2$	10%	10%	10%
$\rho_{\hat{w}_1, \hat{w}_2}$	0	0	0
γ	1.25	1.25	1.25
$\hat{w}_{1t}$	0	-0.284×10^{-3}	-0.525×10^{-3}
$\hat{w}_{2t}$	0	-3.381×10^{-3}	-6.245×10^{-3}
$\hat{\tau}_1$	2.9%	2.90010%	2.90019%
$\hat{\tau}_2$	10%	10.0042%	10.0080%
$\rho_{\hat{w}_1, \hat{w}_2}$	0	0.250×10^{-3}	0.469×10^{-3}
γ	1.5	1.5	1.5
$\hat{w}_{1t}$	0	-0.299×10^{-3}	-0.532×10^{-3}
$\hat{w}_{2t}$	0	-3.557×10^{-3}	-6.315×10^{-3}
$\hat{\tau}_1$	2.9%	2.90022%	2.90041%
$\hat{\tau}_2$	10%	10.0092%	10.0171%
$\rho_{\hat{w}_1, \hat{w}_2}$	0	0.539×10^{-3}	0.995×10^{-3}

V. Alternative consumption processes

The unit root processes for aggregate and idiosyncratic consumption in De Santis (2007) are key to generating sizeable welfare costs of aggregate fluctuations. Unfortunately, the empirical evidence for uninsurable unit roots in consumption is shaky. Blundell, Pistaferri, and Preston (2008) find that many individuals are able to insure against permanent income shocks, whereas Heathcote, Storesletten and Violante (2014) attribute only some of the fluctuations in consumption to unit root processes.

To investigate the sensitivity of our results to assumptions about consumption processes, we look at the consumption equivalent cost of business cycles when the unit roots are re-

tained but the variance of idiosyncratic shocks is decreased by a factor of 2. Our aim is to produce a decomposition of welfare costs that can be compared directly to the numbers for $\gamma = 1.5$ and a detection error probability of 40% in the bottom right corner of Table 3. In doing the calculations, we therefore increase the preference for robustness parameter σ from 6.0 to 10.0. This reflects the greater difficulty an agent has distinguishing between models when the variance of idiosyncratic consumption shocks is low, and ensures that the detection error probability remains at 40%. The consumption equivalent cost of business cycles with a lower variance of idiosyncratic consumption shocks is 1.977%, which decomposes into a baseline cost of 1.020%, a contribution from idiosyncratic risk of 0.033%, a contribution from fear of model misspecification of 0.676% and a joint contribution of 0.247%. Although the magnitudes of the costs do fall slightly, we see that fear of model misspecification is again instrumental in making idiosyncratic risk increase the welfare cost of business cycles. Alternatively, introducing a fear of model misspecification is considerably more costly in an economy that is subject to idiosyncratic consumption shocks.

The consumption processes considered so far do not allow for a link between idiosyncratic risk and the aggregate state of the economy, a specification that may substantially understate the cost of business cycles. This is potentially important, as Kjetil Storesletten, Chris I. Telmer and Amir Yaron (2004) find that the cross-sectional income variation in the US is almost twice as high in contractions as in expansions. De Santis (2007) allows for such a link through a correlation between aggregate consumption shocks and the variance of idiosyncratic consumption shocks, but empirical work by Jacobsen, LaLonde and Sullivan (1993) suggests that it is the increased risk and cost of job displacement that causes cross-sectional heterogeneity to rise during contractions. Specifically, job displacement probabilities and long-term earnings losses of job displacement go up in contractions. Indeed, in less formal discussions, the costs of recessions are often attributed to the increase in earnings losses of displaced workers during recessions, Steven J. Davis and Till von Wachter (2011), or the earnings losses for some sub-group of the population, Philip Oreopoulos, Till von Wachter and Andrew Heisz (2014). In Krebs (2007) it has been shown that the welfare costs of business cycles become quite substantial once this particular feature is incorporated into the model.

We introduce job displacement risk by replacing the shock to the idiosyncratic component of consumption in (4) with the process of Krebs (2007):

$$(14) \quad w_{2t} = \begin{cases} -d_H & \text{with prob } \pi p_H \\ \frac{p_H d_H}{1-p_H} & \text{with prob } \pi (1-p_H) \\ -d_L & \text{with prob } (1-\pi) p_L \\ \frac{p_L d_L}{1-p_L} & \text{with prob } (1-\pi)(1-p_L) \end{cases}.$$

The aggregate state of the economy is assumed to follow a stationary i.i.d. process, such that with probability π it is in expansion and with probability $1-\pi$ it is in contraction. The probabilities of job displacement in expansions and contractions are p_H and p_L respectively, with empirical evidence suggesting that $p_H < p_L$. If a worker is displaced

then its individual consumption level falls by d_H if the economy is in an expansion and d_L if the economy is in a contraction, where it is assumed that $d_H < d_L$ since earnings losses are typically lower if a worker is displaced in an expansion than in a contraction. The increase in individual consumption of non-displaced workers is normalised so that the expected value of idiosyncratic consumption is not trending during either expansions and contractions.

Krebs (2007) measures the welfare cost of business cycles by the amount of consumption that agents would be willing to forego in order to live in an economy where both the probability of job displacement and the consumption losses on displacement are constant over the business cycle. To ensure comparability across different economies, the shock to idiosyncratic consumption when there is no cyclical variation in job displacement risk is modelled as:

$$(15) \quad \bar{w}_{2t} = \begin{cases} -\bar{d} & \text{with prob } \bar{p} \\ \frac{\bar{p}\bar{d}}{1-\bar{p}} & \text{with prob } 1 - \bar{p} \end{cases}$$

where $\bar{p} = \pi p_H + (1 - \pi)p_L$ and $\bar{d} = \pi d_H + (1 - \pi)d_L$. The consumption equivalent cost of business cycles is the amount of extra consumption needed to compensate an agent for idiosyncratic consumption shocks following (14) rather than (15).

The i.i.d. nature of this formulation means that job displacement risk only affects the model through changes in the moments of (w_{1t}, w_{2t}) . The small noise expansion of Anderson, Hansen and Sargent (2012) hence remains valid and the calculations in Appendix A still apply. We adopt the numerical calibration of job displacement risk in Krebs (2007). The probabilities of job displacement in expansions and contractions are set at $p_H = 0.03$ and $p_L = 0.05$ respectively, reflecting the lower likelihood of job displacement in an expansion than in a contraction. That workers lose 9% of their earnings if displaced during an expansion motivates setting $d_H = 0.09$. Similarly, job displacement in a contraction leads to a 21% loss of earnings and we have $d_L = 0.21$. Expansions and contractions are assumed to occur with the same frequency, so $\pi = 0.5$. The numerical calibration implies that $\bar{p} = 0.04$ and $\bar{d} = 0.15$ in the model without cyclical variation in job displacement risk. The rest of the model is calibrated as in Section 4, with the coefficient of relative risk aversion taking the value 1.5 and the preference for robustness parameter σ chosen as before to generate detection error probabilities of 50%, 45% and 40% when testing based on a sample period of 69 years.

The welfare cost of a cyclical component in job displacement risk is decomposed at the top of Table 5. The story is similar to that in the previous model with no link between idiosyncratic risk and the aggregate economy, although now the consumption equivalent costs are larger. Looking at the final column for a detection error probability of 40%, we see that the cyclical component of job displacement risk raises the costs of business cycles by 0.153% if agents do not fear model misspecification, but by 0.446% if they do. Alternatively, endowing agents with a fear of model misspecification increases the costs of business cycles by only 0.615% if there is no cyclical component to job displacement risk, as opposed to 0.908% if there is.

TABLE 5—DECOMPOSITION OF THE CONSUMPTION EQUIVALENT COSTS WITH JOB DISPLACEMENT RISK

<i>detection error probability</i>	50%	45%	40%
Cost of business cycles	1.175%	1.597%	2.083%
= Baseline cost	= 1.022	= 1.022	= 1.022
+ Contrib. of cyclical displacement risk	+0.153	+0.153	+0.153
+ Contrib. of fear of model misspec.	+0	+0.287	+0.615
+ Joint contribution	+0	+0.134	+0.293
$\hat{\pi}$	0.5	0.49991	0.49970
$\hat{p}_H$	0.03	0.03001	0.03002
$\hat{p}_L$	0.05	0.05436	0.05953
$E(\hat{w}_{1t} \mid \text{displaced in expansion})$	0	-0.333×10^{-3}	-0.699×10^{-3}
$E(\hat{w}_{1t} \mid \text{displaced in contraction})$	0	-0.355×10^{-3}	-0.743×10^{-3}
$E(\hat{w}_{1t} \mid \neg \text{displaced in expansion})$	0	-0.319×10^{-3}	-0.666×10^{-3}
$E(\hat{w}_{1t} \mid \neg \text{displaced in contraction})$	0	-0.317×10^{-3}	-0.664×10^{-3}

The middle of Table 5 summarises the worst-case scenarios used by agents to value their consumption streams. The worst-case distribution has agents placing less probability mass $\hat{\pi}$ on the economy being in expansion, and more probability mass $\hat{p}_H$ and $\hat{p}_L$ on the risk of job displacement. Agents especially guard against job displacement risk when the economy is in contraction, reflecting the greater loss of earnings they would suffer if displaced during a contraction. The bottom of Table 5 shows how fears of model misspecification inject correlation between the aggregate consumption shock $\hat{w}_{1t}$ and the cyclical component of job displacement risk. The worst-case distribution of $\hat{w}_{1t}$ is always negative, but this is especially true if the worker is displaced during a contraction. The agent therefore most fears that job displacement is likely to occur in both a contraction and when there is a negative aggregate consumption shock. The resulting correlation mirrors what we found between aggregate and idiosyncratic shocks in the worst-case distribution of the previous model.

VI. Concluding remarks

The message that the welfare cost of business cycles is higher in an economy in which individuals face idiosyncratic consumption risk and are concerned about statistical model misspecification should come as no surprise to readers of De Santis (2007) and Barillas, Hansen and Sargent (2009). What might be more surprising is that the total effect of idiosyncratic risk and concerns about robustness exceeds 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). The calculations of Lucas (2003) make no reference to whether business cycle fluctuations represent efficient allocations, but one conclusion from his estimates is that business cycles are not a worry because costs are small even if all the fluctuations are departures from efficiency. We show that the cost of fluctuations is large, and hence it is important to understand the extent to which they are efficient.

REFERENCES

- [1] Anderson, Evan W., Lars Peter Hansen and Thomas J. Sargent, 2003, "A quartet of semigroups for model specification, robustness, prices of risk, and model detection." *Journal of the European Economic Association* 1(1), 68–123.
- [2] Anderson, Evan W., Lars Peter Hansen and Thomas J. Sargent, 2012, "Small noise methods for risk-sensitive/robust economies." *Journal of Economic Dynamics and Control* 36, 468-500.
- [3] Atkeson, Andrew, and Christopher Phelan. 1994. "Reconsidering the Costs of Business Cycles with Incomplete Markets." In *NBER Macroeconomic Annual 1994*, ed. Stanley Fischer and Julio J. Rotemberg, 187-207. Cambridge, MA: MIT Press.
- [4] Barillas, Francisco, Lars Peter Hansen and Thomas J. Sargent, 2009, "Doubts or variability?" *Journal of Economic Theory* 144: 2388–2418.
- [5] Beaudry, Paul, and Carmen Pages, 2001. "The Cost of Business Cycles and the Stabilization Value of Unemployment Insurance." *European Economic Review*, 45(8): 1545-1572.
- [6] Blundell, Richard, Luigi Pistaferri and Ian Preston, 2008, "Consumption Inequality and Partial Insurance." *American Economic Review* 98(5): 1887–1921.
- [7] Davis, Steven J. and Till von Wachter, 2011, Recessions and the Cost of Job Loss, *Brookings Papers on Economic Activity* 43(2): 1-72.
- [8] De Santis, Massimiliano. 2007, "Individual Consumption Risk and the Welfare Cost of Business Cycles." *American Economic Review* 97(4): 1488-1505.
- [9] Hansen, Lars Peter and Thomas J. Sargent, 2001, "Robust Control and Model Uncertainty." *American Economic Review* 91(2): 60-66.
- [10] Hansen, Lars Peter and Thomas J. Sargent, 2008, *Robustness*, Princeton, New Jersey, Princeton University Press.
- [11] Hansen, Lars Peter, Thomas J. Sargent, Gauhar Turmuhambetova and Noah Williams, 2006, "Robust Control, Min-Max Expected Utility, and Model Misspecification." *Journal of Economic Theory* 128: 45-90.

- [12] Hansen, Lare Peter, Thomas J. Sargent and Neng E. Wang, 2002, “Robust permanent income and pricing with filtering.” *Macroeconomic Dynamics* 6, 40–84.
- [13] Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante, 2014, “Consumption and Labor Supply with Partial Insurance: An Analytical Framework.” *American Economic Review*, 104(7): 2075-2126.
- [14] Jacobson, Louis S., Robert J. LaLonde and Daniel G. Sullivan, 1993, “Earnings Losses of Displaced Workers.” *American Economic Review*, 83(4): 685-709.
- [15] Krebs, Tom. 2007. “Job Displacement Risk and the Cost of Business Cycles.” *American Economic Review*, 97(3): 664-686.
- [16] Lucas, Robert E., Jr. 1987, *Models of Business Cycles (Yrjö Jahnsson Lectures)*. London: Blackwell Publishing
- [17] Lucas, Robert E., Jr. 2003, “Macroeconomic Priorities.” *American Economic Review*, 93(1): 1-14.
- [18] Oreopoulos, Philip, Till von Wachter and Andrew Heisz, 2012, “The Short- and Long-Term Career Effects of Graduating in a Recession: Hysteresis and Heterogeneity in the Market for College Graduates.” *American Economic Journal: Applied Economics* 4(1), 1-29.
- [19] Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, 2004, “Cyclical Dynamics in Idiosyncratic Labour Market Risk.” *Journal of Political Economy*, 112(3): 695-717.
- [20] Tallarini, Thomas D., Jr. 2000, ”Risk-sensitive real business cycles.” *Journal of Monetary Economics*, 45(3): 507-532.

THE SMALL NOISE EXPANSION

Here, 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, Hansen and Sargent (2012).

The first step in expressing the right 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:

$$(A1) \quad e^{-\sigma\beta(W^0(c_{t+1}^i)+h(c_{t+1}^i))} = e^{-\sigma\beta(W^0(c_t^i)+h(c_t^i))} \left(1 + \sum_{n=1}^{\infty} \frac{\epsilon^{n/2} \mu'_n}{n!} (w_{1t+1} + w_{2t+1})^n \right),$$

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

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 (A1), applying logarithms, multiplying by $-\sigma^{-1}$, and then making a Maclaurin expansion of the resulting expression for small $\sqrt{\epsilon}$. When taking expectations of (A1), the operator E_t applies only to the expression $(w_{1t+1} + w_{2t+1})^n$, because μ'_n is a deterministic function of the logarithm 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 side of (7) in terms of powers of $\sqrt{\epsilon}$, as required:

$$(A2) \quad \beta W^0(c_t^i) + \beta h(c_t^i) - \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k\sigma} \left(\sum_{n=1}^{\infty} \frac{\epsilon^{n/2} \mu'_n}{n!} E_t(w_{1t+1} + w_{2t+1})^n \right)^k.$$

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 (A2) into the right side of (7) implies that the correction term satisfies:

$$(A3) \quad (1 - \beta)h(c_t^i) = - \sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k\sigma} \left(\sum_{n=1}^{\infty} \frac{\epsilon^{n/2} \mu'_n}{n!} E_t(w_{1t+1} + w_{2t+1})^n \right)^k.$$

The way that $\sqrt{\epsilon}$ enters the right side of (A3) directly through the expression $\epsilon^{n/2}$ and indirectly through the dependency of μ'_n on $\kappa_n, \kappa_{n-1}, \dots, \kappa_1$ and the derivatives $D^n(W^0(c_t^i) + h(c_t^i))$, $D^{n-1}(W^0(c_t^i) + h(c_t^i))$, $\dots$, $D(W^0(c_t^i) + h(c_t^i))$ suggests that the correction term may be written as a linear function of powers of $\sqrt{\epsilon}$. A candidate is:

$$(A4) \quad h(c_t^i) = \sum_{n=1}^{\infty} \epsilon^{n/2} h^n(c_t^i),$$

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 (A3) and (A4). The expected value $E_t(w_{1t} + w_{2t})^n$ is given by the moments of the bivariate normal distribution of w_{1t} and w_{2t} . The method of undetermined coefficients applied to terms in $\sqrt{\epsilon}$ in (A3) and (A4) gives:

$$(A5) \quad h^1(c_t^i) = \frac{E_t(w_{1t+1} + w_{2t+1})}{1 - \beta} \beta D W^0(c_t^i).$$

There are no derivatives of $h(c_t^i)$ in the first correction term (A5) because $h(c_t^i)$ is at most of order $\sqrt{\epsilon}$. The second correction term is defined by the undetermined coefficients on

ϵ in (A3) and (A4):

$$\begin{aligned}
 \text{(A6)} \quad h^2(c_t^i) = & \frac{E_t(w_{1t+1} + w_{2t+1})}{1 - \beta} \beta D h^1(c_t^i) \\
 & + \frac{E_t(w_{1t+1} + w_{2t+1})^2}{2\sigma(1 - \beta)} \left(\sigma \beta D^2 W^0(c_t^i) - (\sigma \beta D W^0(c_t^i))^2 \right) \\
 & + \frac{(E_t(w_{1t+1} + w_{2t+1}))^2}{2\sigma(1 - \beta)} (\sigma \beta D W^0(c_t^i))^2,
 \end{aligned}$$

where the derivative $D h^1(c_t^i)$ can be calculated from equation (A5). This completes the derivation of the first-order correction term $\sqrt{\epsilon} h^1(c_t^i) + \epsilon h^2(c_t^i)$. An expansion of second order or higher is obtained by recursive solution and application of the method of undetermined coefficients to (A3) and (A4).